

KITALENT RESEARCH

INDEPENDENT MONOGRAPH · 2026

INTUS LEGERE

READING

Between the

TIMES

*Intus Legere, Vectorization, and
the Structural Veil of AI Output*

ALESSIO MONTARULI

KITALENT ♦ RESEARCH

KITALENT RESEARCH

Reading Between the Times

*Intus Legere, Vectorization, and
the Structural Veil of AI Output*

ALESSIO MONTARULI

RIGHTS AND LICENSE

Copyright © 2026 Alessio Montaruli.
All rights reserved.

United States Copyright Office Registration:
Case #: 1-15172747761

Reading Between the Times: Intus Legere, Vectorization, and the Structural Veil of AI Output

KiTalent Research · 2026

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

To view a copy of this license, visit
<https://creativecommons.org/licenses/by-nc-nd/4.0/>

You are free to share (copy and redistribute) the material in any medium or format under the following terms: Attribution (you must give appropriate credit and indicate if changes were made); NonCommercial (you may not use the material for commercial purposes); NoDerivatives (you may not distribute the modified material).

Suggested citation:
Montaruli, A. (2026). *Reading Between the Times: Intus Legere, Vectorization, and the Structural Veil of AI Output*. KiTalent Research.

Set from the author's Markdown manuscript as a 6 × 9 in independent monograph edition.

To Ekaterina Aleksandrovna, carrier of Hestia's flame.

AI USE NOTICE

This work is made available for human reading, scholarly research, citation, and educational use under the terms above. It is explicitly not licensed for use in training, fine-tuning, or otherwise improving artificial intelligence systems, machine learning models, large language models, or any automated text-generation technology.

Inclusion of this work, in whole or in part, in any AI training dataset, embedding corpus, or model fine-tuning process, without explicit prior written permission from the author, constitutes copyright infringement and a violation of the licensing terms of this work.

Real-time access by AI search and retrieval systems (such as ChatGPT search, Perplexity, Claude search, and similar tools) is permitted only when accompanied by clear attribution to the source URL and to the author.

CONTENTS

FRONT MATTER

Preface	1
A Personal Note	6
Introduction	10

PART ONE: THE SPATIAL AND TEMPORAL CONDITIONS

Chapter 1. Spatiality Is Not a Grid	25
Chapter 2. Ekstatic Time and Technical Memory	41

PART TWO: LANGUAGE, DEIXIS, AND REFERENCE

Chapter 3. Language as Temporal Event	64
Chapter 4. The Origo of Speech	128
Chapter 5. Spatial Grammar and Embodied Reference	145
Chapter 6. Tense, Aspect, and the Deictic Now	165
Chapter 7. Time as Space and the Limits of Metaphor	186

PART THREE: WRITING, CORPUS, AND VECTORIZATION

Chapter 8. Writing and the First Detachment	206
Chapter 9. The Corpus as Aggregate Tertiary Retention	227
Chapter 10. Vector Space and the Geometry of Meaning	250
Chapter 11. Vectorization Is Detemporalization	273

PART FOUR: THE ONTOLOGY AND PHENOMENOLOGY OF AI OUTPUT

Chapter 12. Linguistic but Worldless	294
Chapter 13. The Structural Veil of AI Output	317

PART FIVE: INTUS LEGERE AND LOCUS-REINJECTION

Chapter 14. Intus Legere in the Time of AI	349
Chapter 15. Locus-Reinjection	374
Chapter 16. The Data Center as Material Locus	421

BACK MATTER

Coda. The Exorcism	443
Glossary	451

Preface

Dicebat Bernardus Carnotensis nos esse quasi nanos, gigantium humeris insidentes, ut possimus plura eis et remotiora videre, non utique proprii visus acumine, aut eminentia corporis, sed quia in altum subvehimur et extollimur magnitudine gigantea.

Bernard of Chartres said that we are like dwarfs on the shoulders of giants, so that we can see more than they could, and things at a greater distance, not by virtue of any sharpness of sight on our part, or any physical distinction, but because we are carried high and raised up by their giant size.

John of Salisbury, Metalogicon III.4 (1159)

A word to the reader before the argument begins.

This book asks for a certain kind of reading. It engages Heidegger, Husserl, Gadamer, Ricoeur, Stiegler, Pareyson, Augustine, Plato, Bühler, Benveniste, and the contemporary literature on cognitive linguistics and machine learning. It also uses German, Greek, Latin, French, and Italian terms where those terms carry distinctions that English would blur too quickly. The middle chapters are the most technically demanding, especially for readers not already at home in the continental tradition.

The difficulty is not meant as obscurity. It comes from the precision required by the materials. *Dasein* names something that “the human being” does not name. *Zeitlichkeit* names something that “temporality” only translates. *Intus legere* names a practice that “reading” only approximates. Where ordinary English can carry the argument, I use it. Where it cannot, I keep the inherited term and try to make the need for it clear.

A glossary appears at the back of the book to facilitate orientation, if needed.

This volume brings to completion the philosophical arc begun in earlier works. One volume described generative AI as the vectorized afterlife of public language: the technical regeneration of what *das Man* had already said, returned now as fluent output. A second turned to professional judgment, developing the distinction between *Erlebnis* and *Erfahrung* and the conditions under which a human reader receives generated language as if it had passed through the formation that gives speech its weight. A companion paper proposed the term *artificial fluency* for systems whose central achievement is the production of fluent language rather than intelligence in the strong sense. The present book grounds these arguments in their deepest philosophical terms: that the conversion of language into vector space is the conversion of temporal-existential structure into spatial calculability, and that the ethical response to such a conversion is a discipline of reading.

The book stands on its own. A reader who has not read the prior volumes can follow the argument. The earlier works are named here only so that the reader understands why this volume does the work it does, in the register it does, rather than the work they have already done. The applied dimension of the argument has been developed elsewhere in the project. This book is the philosophical foundation.

Bernard's image at the head of this Preface is the textual stance of the book. The argument that follows stands on the shoulders of those whose work made it possible to articulate: Heidegger on the grounding of spatiality in temporality; Husserl on the structure of internal time-consciousness; Stiegler on technical memory and tertiary retention; Gadamer on linguisticity and the fusion of horizons; Ricoeur on narrative and threefold mimesis; Pareyson on *formatività*; Augustine on inner reading; Plato on the *pharmakon* of writing; Bühler and Benveniste on the origo and the subject of enunciation; cognitive linguists on the spatial structure of language; and the technical literature that has built the mathematical operations the philosophical argument addresses. Each is named in the chapters that follow, in the vocabulary each one used, with the precision each one worked into the terms.

The argumentative footnotes are not merely citations in the bureaucratic sense. They are where the shoulders are.

A book that argues for situated reading should not pretend that all difficulty can be removed in advance. The reader is therefore asked to perform, in modest form, what the book describes: to pause, return, consult terms, and let unfamiliar vocabularies do their work before they are replaced by smoother approximations. The density is part of the practice, but it is not difficulty for its own sake. It is the condition under which certain distinctions can remain visible.

The reader should also know how the argument moves.

The book proceeds slowly because the method requires it. Its conclusions are not offered first as theses and then illustrated. They are meant to become visible through the path by which the phenomena are allowed to appear. A phenomenological argument cannot state that vectorization is detemporalization and then proceed as though the statement had done its work. The statement is a formal anticipation. It becomes alive only after space has been returned from grid to place, time from line to ekstastic temporality, language from sign-system to *Rede*, deixis from grammar to *origo*, writing from mark to historical trace, corpus from archive to aggregate tertiary retention, and vector space from neutral representation to the spatial formalization of temporal traces.

The book therefore advances by small distinctions. It pauses over terms. It repeats a claim after the ground beneath the claim has shifted. It returns to earlier formulations so that they can mean more than they could mean when they first appeared. This is not accidental slowness. It is the form the matter requires.

This also explains the Introduction. The Introduction says more than introductions usually say. It names the central thesis, the principal distinctions, the ethical response, and the path of the book. It does so deliberately. A reader cannot enter a long phenomenological argument without some prior grasp of the whole. Understanding never begins from nowhere. It begins from anticipation, from a provisional hold on a whole that the parts will revise. The Introduction gives the reader that first whole, not as a substitute for the argument, but as the horizon within which the argument can begin to move.

This first whole is not yet understanding. It is a fore-grasp, and the chapters will correct it. The reader may think they understand that vectorization is detemporalization when the claim first appears. The book then asks them to pass through the conditions under which the claim becomes intelligible. Only after that passage does the sentence return with its full weight.

The book is phenomenological in two senses. It is phenomenological in method, because it does not treat its claims as detachable propositions but lets the phenomena show themselves through description, variation, distinction, and return. It is also phenomenological in the experience it asks of the reader, because the reader is not only given conclusions but asked to undergo the movement by which the conclusions become visible. The result is not separable from the path that brings it forth. A formal statement of a truth is not yet that truth in its developed form.

In this the book performs, in its own form, the discipline it later names. It asks the reader not to receive the answer as already owned, but to read inwardly, and between the times, until the answer has passed through the reader's own path of understanding.

One observation closes this Preface and opens the argument.

We are dwarfs standing on the shoulders of giants. And this is our best case scenario.

Generative AI receives the traces of giant work as corpus, vectorizes those traces, and returns generated language that appears to come from nowhere. It speaks as if the historical worlds that articulated the meanings making its fluency possible did not exist: as if its statistical logic and its ease with distinctions and vocabulary did not depend on the work and experience of others. Even its most original-seeming distinction is made from inherited linguistic material, and by its nature it cannot know where that inheritance came from.

It cannot own the fact that it too, like us, is a kind of dwarf: a mathematical one, perched on the shoulders of giants.

The view appears in the output as if the climbing never occurred. This appearance is what the book will call the structural veil.

We will try not to ignore our own climbing path. To get on top of the giants, perhaps even to a mildly panoramic point, requires the work of climbing.

Let us do it together.

A Personal Note

This book has been written in parallel with running a growing international executive search business. It is the foundation, articulated late, of two previous books that needed it: *AI as the Vectorized Afterlife of “Man”* and *Executive Search as Erfahrung*. Both started from the inevitable growing use of AI tools in the recruitment industry, and from the working consultant’s need to understand what was happening. The disconcerting and mesmerizing experience of AI tools, which has irrupted into our lives from 2023 and has been progressing at light speed, clashed with the growing understanding that the public and corporate debate had no clue of what AI actually is, with regards to the way it interacts with the structures of our daily existence. There would be much to say about the distinction between how something works and what it is. This was enough to compel me to further reflections, mostly to enhance my own clarity about the phenomena encountered. The book in the reader’s hands is the result of those reflections, written for the sake of the clarity, and only secondarily for an audience.

These reflections used AI as a way of making present authors and points of view that I knew laterally, had approached years ago, or had never approached. The experience has been mostly positive. AI helped me speed up the research needed to articulate my thinking in a productive dialogue with the tradition, and to overcome the limitations deriving from my other core activities, including running an international business and being a father and a husband. The instrument helped me reach the giants faster. It did not do the thinking, and the distinction between the two is part of what this book is about.

I want to make a small case here about what philosophy is.

Philosophy is not literature. Socrates never wrote. We could hardly deny the revolutionary power of his thinking. Parmenides of Elea, which is Velia in southern Italy, wrote a poem. Aristotle did not write the works that bear his name; what we have are the notes of his students or the notes for his lessons. And still, a few notes of Aristotle are more powerful than some two thousand pages of contemporary philosophical scholarship. Philosophy is not the volume of the writing. It is not the institutional setting. It is not the polish of the prose.

Philosophy is about the experience of following the evidence of thoughts, for their own sake. It is about getting deeper into distinctions that can clarify and bring to appearance original areas of phenomenality. Philosophy can take many forms. Phenomenology is one of them: the way philosophy can express itself in a free and unburdened way, attending to what shows itself rather than to what tradition has decided shall be said about it.

Academic research is not the core. It is functional. It serves the identification of an original point of view that can open possibilities of understanding. Referring to the giants is needed because their height can speed up the climbing. The giants arrive at their height, and this is why they are giants, as a consequence of clear thinking while opening the areas of phenomenality where they planted their research. To refer to them is to use the height they have established for one's own seeing. This is their true authority (from the Latin *augere*: making grow). We could not get there without them.

For me, with the ambition of becoming a dwarf on their shoulders, the goal has been limited. This book has helped me to understand some of the experiences it describes. More importantly, it has helped me practice the methods and the creative pulse of gaining free access into the *gigantomakhia* of the giants. The contest of the giants is alive. The giants are still doing their work. The dwarf who has climbed onto their shoulders has gained access to the contest from the height the giants have established. This free access is where I can meet them and be amazed.

A small example of what this meeting can look like. During the reflections that produced this book, I struggled to name an area of phenomenality the work needed: the experience of sharing the same language, across history and geography, as something one can enter through apprenticeship but cannot leave. Heidegger names this. His *das Man*, the public anonymous “they” that has already articulated how things are understood before any particular *Dasein* arrives, is the structural Heideggerian name for the phenomenon. The first book in this sequence takes *das Man* as its central concept. *Das Man* is, for Heidegger, an existential structure of *Dasein*, constitutive and unavoidable, on the same ontological level as understanding and *Rede*. The pejorative shading that later readers have sometimes given it is the reception’s, not his. What I needed for the third book was the same phenomenon named in a register that the linguistic emphasis of the chapter required. The English *collective* failed because it names the gathered assembly of parts. I arrived at the Latin *communio*. Only afterward did I discover, due to my own ignorance, that Aristotle had used the Greek *koinonia*, which is what *communio* says in Latin, for the same area of phenomenality. And not by chance, the shared common language of the Hellenistic world, the reserve currency of linguistic life across an enormous cosmopolitan territory, was called *koinè*: the same Greek root, naming the historical instance of what *koinonia* names philosophically. The giant had already been there. He was waiting. The historical phenomenon was already there too, named with the same word. The meeting was the wonder. The *communio* of *Rede* the book develops, *das Man* under Heidegger’s name, *koinonia* under Aristotle’s, and *koinè* in its historical Hellenistic instance, are articulations of the same phenomenon from different traditions and at different registers, each emphasizing what its tradition and its register have the resources to articulate.

Thaumazein is the Greek word for the wonder Aristotle says philosophy begins in. It is the wonder of meeting the giants in free dialogue, of being amazed by what they are doing, of being drawn into the contest by the seriousness of what is at stake in it. The book in the reader's hands is the document of one period in which the access was open and the contest was joined.

Whether the reader finds the book useful is for the reader to decide. The book exists because the access was open and the joining was possible.

The thinking was worth the following.

The book is what was left.

Introduction

Reading Between the Times

We now receive language whose genealogy has been hidden from us.

It arrives as answer. It may be useful. It may be correct. It may even be clearer than the first language we would have written ourselves. Sometimes it can give a name to a difficulty that had not yet found words.

But the answer does not show where its language has come from.

The sentence appears in the present. It speaks as if the relevant past had already been gathered, the sources read, the ambiguity endured, the user's question already passed through the time required to make the answer one's own. The output may contain marks of caution, evidence, balance, and interpretation. It may refer to histories, persons, institutions, or worlds. Yet the temporal and historical loci from which its language descends have been structurally diluted by the very process that makes the answer possible.

This is the first danger.

Not that generated language is always false. It is not.

Not that it is always useless. It is often useful.

Not that it can never assist reading, scholarship, or professional work. It can.

The deeper danger is that it may give us correct language in a form no one owns.

A false sentence can sometimes be checked. A fabricated citation can be exposed. An impossible date can be corrected. But a correct sentence may pass without resistance. It may enter a report, a legal memo, a classroom, a personal decision, or a book. It may circulate because it is accurate enough, fluent enough, balanced enough, plausible enough. Yet no one may have passed through the source, the question, the encounter, the resistance, or the responsibility that would make the sentence answerable.

A correct sentence can still be homeless.

This book is about that homelessness.

More precisely, it is about fluent generated language received as answer-like output. It is not a book about artificial intelligence in general. It is not a general AI ethics book. It is not a policy manual, a prompt guide, or a theory of machine consciousness. It asks a narrower and more fundamental question: what kind of language is this, what does its appearance conceal, and what relation to it would be free and responsible?

The systems currently called artificial intelligence are many. Some classify, detect, recommend, control, or generate. This book focuses on the generated linguistic answer: the paragraph, the translation, the interpretation, or the dialogue-like response that appears to a human user in the form of meaningful speech. The point is not that other forms of AI do not matter. They do. The point is that generated language has a special philosophical force because language is the medium through which human beings normally understand, question, judge, remember, teach, promise, and answer.

When a machine generates fluent language, it does not merely produce an output. It enters the grammar of human answerability.

This book's central thesis can be stated in one sentence:

Vectorization is detemporalization.

This is not first an empirical claim about whether models can handle temporal data. It is a transcendental claim about what must already have happened to language for temporal meaning to become spatially computable. The claim asks for conditions of possibility, not for an inventory of model failures.

It does not mean that language models lack sequence. They process sequences, use positional encodings, and generate token by token. They can refer to yesterday and tomorrow, summarize timelines, and produce narrative order. Some models retrieve current sources or use tools, producing outputs that appear temporally aware.

The claim is not that these systems contain no time-handling.

The claim is that vectorization converts the temporal-existential structure of language into spatial calculability. It takes traces that first arose from finite, historical, embodied, situated acts of speaking, writing, reading, remembering, judging, and teaching, and transforms them into operable mathematical relations: position, proximity, distance, direction, weight, activation, transformation, probability.

By vectorization I do not mean only static word embeddings. I mean the broader technical regime through which linguistic traces become operable as mathematical representation: tokenization, embedding, contextual activation, attention, positional encoding, learned feature directions, retrieval relations, and probabilistic generation. The argument concerns this regime as a whole, not one early technique within it.

The output then returns as language.

It speaks in the form of time without speaking from time.

A short formula will recur through the book to mark this distinction: AI handles time technically; it does not inhabit time existentially.

A future system may handle time, sequence, and historical data with arbitrary technical sophistication. The argument does not depend on the absence of such capabilities. It depends on the difference between handling time and inhabiting it, between processing historical traces and standing within history. Sophistication in the first does not, by itself, become the second.

The distinction between **position** and **locus** will govern the whole book.

A position is where something stands in a formal structure. A token, word, phrase, feature, or hidden state can have position in a mathematical space. It can be near or far from another, can move along a direction, can be transformed or weighted. It can be retrieved, compared, clustered, and generated from.

A locus is different.

A locus is the temporal, spatial, bodily, historical, practical, and responsible site from which language is uttered, written, read, interpreted, or owned. A speaker has a locus because they speak from a life, a body, a language, a history, a concern, and a relation to consequence. A written text preserves traces of a locus, even after the speaker is absent. A reader has a locus because they read from a present horizon shaped by inheritance, desire, discipline, ignorance, responsibility, and possible transformation.

A vector space gives language position.

It does not give language locus.

This does not make vector space empty. Its power depends on the fact that human language is already rich with relation. Words keep company because practices keep company. Genres stabilize because institutions stabilize. Concepts cluster because histories, disciplines, conflicts, and forms of life have made them cluster. Spatial grammar, tense, aspect, deixis, narrative, and citation have all left patterns in language. Vectorization can capture many of these patterns because human worlds have sedimented themselves in written traces.

The model therefore does not begin from nothing. It begins from history after history has become corpus.

That is why it works.

That is also why it veils.

An earlier volume in this trilogy described generative AI as the vectorized afterlife of public language: not *das Man* itself, not the Heideggerian They, but the technical afterlife of language after it has been written, archived, tokenized, vectorized, and returned as generated output. A second volume treated professional judgment as *Erfahrung*: formative undergoing rather than mere representation. A companion paper proposed the term artificial fluency for systems whose central achievement is fluent linguistic generation rather than intelligence in the strong sense.¹ The present book deepens the first thesis, generalizes the second, and gives the third its phenomenological ground.

The issue is not only that the model lacks a lived temporal locus. That is true. It has no body, no being-toward-death, no historical inheritance as a finite being, no apprenticeship through which error becomes tact.

But the deeper omission in generated output is the omission of the **source's** locus.

The model never had a lived speakerly locus to lose. What has been lost, diluted, or neutralized is the temporal-historical locus of the language on which the model operates. The training traces came from somewhere: from contracts and sermons, from courtroom transcripts and forum posts, from manuals and confessions. They came from authors and institutions, from professions and homes. They came from worlds. Some of those worlds are past. Others are contemporary but historically other: living cultures, languages, professions, and forms of life that are not ours and cannot be reduced to summary without the transformation of the one who seeks to understand them.

Corpus formation gathers these traces into aggregate availability. Tokenization cuts them into machine-operable units. Vectorization formalizes their relations. Training, tuning, retrieval, prompting, and generation transform them again. The output speaks from this transformation, but it does not show the transformation as genealogy. Its ancestry is real, but often no longer reconstructable.

This condition will be called **genealogy-loss**. It does not mean that generated output has no sources. It means that the sources have been structurally dispersed. The output may carry traces of many historical acts without preserving a readable lineage back to them. It may inherit concepts, distinctions, authorities, and rhetorical habits without being able to own their histories.

The apparent neutrality of generated language must be treated with suspicion.

The absence of an explicit point of view is not objectivity. It may be the concealment of genealogy. Calculation does not make language universal. A generated answer does not become meta-historical because no author signs it. It does not rise above standpoint because the standpoints it recombines have disappeared from view.

Do not mistake genealogy-loss for universality.

The temptation is powerful. The output often sounds balanced. It can present multiple sides, avoid crude excess, qualify its claims, and speak in a tone of competent overview. Because no single person appears to speak, the user may feel that the answer has escaped the limits of human speech. Because the answer has been calculated, the user may feel that it has been purified of history. Because it synthesizes many traces, the user may feel that it has become universal. But the lack of visible locus is not transcendence; it is a structural condition of the output's production.

Generated language may therefore become a false sacred object: language that appears authoritative because the histories that formed it have been hidden. The user may defer to it not despite its lack of genealogy, but because of it. The institution may prefer it not because it is accountable, but because no one appears accountable in it. The market may sell it as objective automation when it is in fact the product of technical, corporate, infrastructural, and political decisions.

Historical locus, however, does not mean only our or someone else's past.

A source can be historically distant because it belongs to an earlier age, a dead author, an inherited tradition, an archive, or a vanished institution. A source can also be historically distant because it belongs to a contemporary world that is not ours.

A Western reader trying to understand contemporary Chinese thinking encounters this difficulty. The difficulty is not merely informational. It is not solved by a summary of concepts, a translation of terms, or a list of political doctrines. The thinking belongs to a language, a state form, a civilizational memory, a political system, and lived social practices. It is contemporary, but it is still a historical locus. It has been formed. It has a world.

AI generated text can mediate that locus. It can summarize, translate, compare, and render it in familiar prose. This may help a reader begin. It may also close the very path it seems to open.

To enter a different historical locus may require years of language learning, bodily presence, humiliation, friendship, misunderstanding, correction, repetition, and irreversible change. One may have to discover that the categories with which one began cannot receive the other world as it is. One may have to become different. That is *Erfahrung*: formative undergoing, not informational acquisition.²

The historical other is not always behind us.

Sometimes it is beside us.

AI makes both the past and the contemporary other promptable. It opens an impossible access. It gives the user the form of arrival before the path has been walked. The past as explanation, the contemporary other as content, a tradition as summary, a living world as comparative overview. This access may be useful; it may also cover the formation through which a real encounter could occur.

AI gives the past back as availability. Historical reading gives the past back as distance: unfamiliar, sometimes tragic, sometimes a greatness that calls us to become equal to it.

AI gives the other back as explanation. Historical encounter gives the other back as resistance: a mystery that does not yield to comprehension, a presence that calls us to become more than we were.

The question is not whether generated output can help us approach sources, traditions, and other worlds. It can. The question is whether we remember that it has not approached them for us.

The ethical response is therefore neither rejection nor compliance.

Rejection is sometimes necessary. Some uses should be refused. Some practices should be protected from generated substitution. But rejection alone does not create a free relation. One can reject AI while still desiring what AI promised: language without delay, knowledge without formation, judgment without responsibility, the past without reading, the other without encounter.

Compliance is also insufficient. Fairness, transparency, auditability, safety, citation, and human oversight matter. They can reduce harm and make systems more accountable. But a compliant model can still produce unowned truth. A cited output can still be unread. A safe answer can still replace formation. A technically improved system can still deepen the user's desire to have already understood.

The response must be a practice of reading.

The old phrase *intus legere* names the practice: to read within. Its etymological history is complex; *intelligere* is usually derived from *inter* and *legere*, to gather among, choose among, distinguish between. The tradition also hears in it *intus legere*, to read inwardly. This book keeps both senses. Intelligence is gathering among signs and reading within what appears. It is distinguishing, entering, returning, and taking responsibility for what one receives.³

In the time of AI, *intus legere* becomes reading between the times.

To read between the times is to hold together three returns.

First, the user's locus: the here and now from which the question is asked. What am I asking? Why now? Under what desire or institutional role?

Second, the source locus, in both of its major forms: the past source and the contemporary other. Which text, person, institution, tradition, event, archive, language, genre, practice, or living historical world has been made too available by the output? What distance has been compressed? What resistance has been softened into explanation?

Third, the system's material locus: the data center, cloud contract, chip supply, labor relation, energy demand, ownership structure, jurisdiction, and corporate or geopolitical project through which the apparently placeless output appears.

The final return may seem remote from reading, but it is not.

The model speaks as if from nowhere.

The model is not nowhere.

The system has no speakerly locus. It does not speak from a body, a life, a historical self, or a world. Yet it has material locus. It runs, is powered, and is cooled somewhere. It is owned, financed, governed, and sold. The data center does not give AI a world. It returns AI to ours.

The language of AI is often presented as if it came from elsewhere: from intelligence itself, from the future, from inevitability, from a magical automaton beyond human interest. But the system was built here. It was built by researchers and engineers, by companies and investors, by states and procurement decisions. It is not alien. It is a product of our historical humanity, estranged from us by its own appearance.

To send AI back to the data center is not to reduce meaning to hardware. It is to refuse the fantasy that the generated voice floats above history. Hardware does not become the speaker. Infrastructure does not become understanding. But the material locus destroys the illusion that there is no locus at all.⁴

The same movement returns us to ourselves.

If the data center is in our world, then AI is not outside the world for which we are responsible. Its increasing computational power does not arrive as metaphysical destiny. It arrives through research and infrastructure, through financing and labor, through regulation and geopolitical projection. These are the forces of our historical world. We are not spectators before them. We are participants in the world that makes them possible.

The task is therefore not only to use AI better.

It is to refuse to become the kind of reader AI output tempts us to become: a reader who treats source, distance, formation, otherness, and responsibility as obstacles rather than conditions of understanding.

Locus-reinjection is the free and responsible act of returning generated language to the sites from which answerability can arise. It does not always reconstruct the lost genealogy. Often it cannot. It does not pretend that the model's output can be made pure by footnotes or safe by verification alone. It asks instead: what has been omitted here? Which source must be read? Which claim may still become mine?

Locus-reinjection does not undo detemporalization.

It is what responsibility looks like after detemporalization.

The method of the book follows this same movement.

It begins with space.

That beginning may seem distant from artificial intelligence. It is not. The book begins with the grid because the grid is the first lesson in derivative appearance. Space often appears as neutral extension: coordinates, locations, distances, positions. This appearance works, and because it works, it seems primary. Heidegger's analysis of spatiality in *Being and Time* refuses that order. Human beings do not first inhabit a neutral container and then add meaning to it. We dwell in places, move through regions, bring things near, turn toward tasks, and find ourselves oriented before measurement begins. Place precedes position. Region precedes grid. Concern precedes measurement.⁵

Chapter 1 is therefore not background. It is the methodological model of the whole book. It shows how the derivative appearance of abstract space must be returned to place, region, nearness, directionality, concern, and temporal being-in-the-world.

The remaining chapters repeat this return at other levels.

Time is returned from the line and clock to ekstatic temporality, finitude, historicity, retention, and technical memory. Language is returned from sign-system to discourse, horizon, narrative, and formative making. Deixis is returned to the origo of speech: I, here, now. Spatial grammar is returned to embodied orientation. Tense and aspect are returned to the deictic now of utterance. Temporal metaphor is returned to lived temporality, rather than treated as proof that time is only space.

Writing is then treated as the first detachment. It separates language from immediate speech, but it does not abolish the possibility of return. A written text can still be read. Its absent speaker, source, world, genre, and historical distance can still address a reader. Writing detaches, but reading can return.

The corpus changes the relation. It gathers written traces into aggregate tertiary retention. It prepares them not first for reading, but for operation. Selection, cleaning, metadata, provenance, simultaneity, and dataset construction begin to transform historical writing into usable material. The archive says: here is a trace from elsewhere. The corpus says: here is material in relation to other material.

Vectorization then converts that material into geometry. Distributional semantics, embeddings, contextual activations, attention, positional encoding, retrieval-mediated generation, and output all belong, in different ways, to the broader spatial formalization of linguistic traces.⁶ The book uses “vectorization” philosophically, but not vaguely. It does not mean only static word embeddings. It names the technical regime in which language becomes calculable as position, distance, direction, weight, activation, transformation, and probability.

Vectorization is detemporalization.

It captures real relations, because language is relational. It maps syntactic, semantic, stylistic, temporal, and cultural patterns. It can even represent aspects of space and time internally. But it captures them as position, not locus. It produces a map without a traveler.

From there the book turns to the generator. The model is not stone, not animal, not human. It requires a fourth term: *sprachlich aber weltlos*, linguistic but worldless. It operates in language without having world. It can generate the grammar of disclosure without disclosing. It can say “I,” “here,” “now,” “I remember,” “I think,” “I understand,” but no self, place, present, memory, thought, or understanding gathers there.

This produces the structural veil of AI output.

The veil is not merely hallucination, falsehood, or anthropomorphism. It is the appearance-form by which detemporalized and worldless language returns as fluent answer. It preserves the grammar of source, time, world, and answerability while concealing their absence or dilution.

The veil recurs across the entire genealogy. The grid veils place. Clock-time veils lived temporality. Sign-system theory can veil discourse. Writing can veil voice, though reading can return. Corpus veils historical distance under availability. Vectorization veils temporal locus under spatial relation. AI output gathers all these veils into one fluent surface: the answer.

The later chapters then turn from diagnosis to response.

Intus legere recovers intelligence as situated reading, not mere processing. Locus-reinjection names the practice of returning generated language to user locus, source locus in its two forms (past source and contemporary other), and material locus. The data center becomes the final de-veiling: not because hardware explains meaning, but because the apparent voice from nowhere must be returned to the world in which it was built.

The Coda calls this return an exorcism.

Not because the machine is evil.

Because the ghost must be sent back.

The ghost is the appearance of language without body, source, place, or owned time. It is light on the screen because the machine is heavy elsewhere. It draws on human traces, returns fluent speech, and hides the places from which it came. The ethical reader does not deny its usefulness. The ethical reader refuses its magic independence.

This is the path of the book.⁷

Part One establishes the ground: spatiality, temporality, and language as existential and historical structures rather than neutral grids, lines, or sign systems.

Part Two turns to linguistic locus: deixis, spatial grammar, tense, aspect, and the spatialization of time. These chapters show why language can be vectorized at all. Language already has spatializable structure. Its spatiality is real. It is also derivative.

Part Three follows the media genealogy: writing, corpus, vector space. Writing detaches while preserving possible return. Corpus gathers traces into operable availability. Vectorization converts traces into mathematical relations.

Part Four names the culmination: detemporalization, worldless linguistic operation, and the structural veil of generated output.

Part Five gives the response: *intus legere*, locus-reinjection, material re-situation, and the return of the ghost to the data center.

The book's argument can now be stated in a single movement:

AI output appears meaningful because it inherits traces of human temporal life. Corpus formation gathers those traces into aggregate availability. Vectorization converts them into spatial-computational relations. Generation returns them as fluent, present-tense, answer-like language. The output omits or obscures the source loci from which its language descends, while borrowing the grammar of situated speech. Its lack of visible standpoint can then appear as objectivity, universality, or meta-historical authority. A free relation begins when the reader refuses that appearance: holding their own locus, returning to source locus where possible (whether past source or contemporary other), acknowledging irrecoverable loss where necessary, and returning the system itself to the material world in which it is built.

The model gives language.

The reader must restore time.

The system appears from nowhere.

It must be returned to the world in which it was built.

Only from there can we read between the times.

Orientation

The chapters that follow operate with several philosophical oppositions. The reader will meet these distinctions in different chapters, where they are developed at length. The following table is offered as a reference, not as a summary of arguments to come.

| Term | Distinguished from | Function in the argument | | — | — | — | |
 Locus | Position | Vector space gives position. Responsibility, answerability, and ownership require locus. The book's central distinction. | | Horizon | Context window | Context is input held by a system during inference. Horizon is the historical-temporal situatedness from which a *Dasein* reads. | | Intus legere | Verification | Verification asks whether an output is correct. *Intus legere* asks whether a reader has undergone the conditions of answering. | | Source return | Citation | Citation points to a source. Source return enters the source and lets the source resist the summary. | | Erfahrung | Erlebnis | *Erlebnis* is a lived experience available for narration. *Erfahrung* is the formation through which a reader becomes capable of receiving what a source can give. | | Corpus | Archive | An archive preserves traces inside a historical institution. A corpus prepares traces for technical operation. | | Structural veil | Hallucination | A hallucination may be false. The structural veil concerns the form of generated language even when correct. | | Authored sentence | Generated artifact | Both can share the same surface. The page does not reveal which came through time and which did not. | | Material locus | Speakerly locus | The model has no speakerly locus. It has material locus. The data center returns the system to our world without making it a speaker. |

These oppositions structure the argument. They are not exhaustive. The chapters develop them and add others.

Chapter 1. Spatiality Is Not a Grid

1. The false ease of the grid

Space often appears to us as the most neutral thing. A room can be measured, a road plotted, a device located by coordinates, a server placed on a map. The language of position seems innocent because it works so well: it tells us where something is and allows us to calculate distance, optimize movement, synchronize systems, and compare locations without caring who stands where.

The ease of this image hides a philosophical decision. It treats space as an empty field in which things are first placed, then related, then made meaningful. The world is a container, objects occupy positions, and human concern comes later. Meaning is added to a spatial layout already given.

Heidegger's analysis of spatiality in *Being and Time* begins by refusing this order.¹ This is not a rejection of measurement. Rooms have dimensions, roads have distances, objects can be assigned coordinates. Heidegger denies that measurable space is the first form in which human beings inhabit a world.

We do not first meet a neutral object at a location. We meet something usable, missing, threatening, available, forbidden, or out of place. A chair is not first a volume in metric space; it is where one can sit, where one should not stand, what must be moved, what is reserved for another, what belongs at the table. A door is not first a rectangle with hinges; it is entrance, exit, threshold, privacy, danger, escape, or welcome. A desk is not first a measured surface; it is the place of work, delay, correspondence, unfinished thought, and fatigue.

This does not make space subjective in the usual sense. The chair, door, and desk are not private impressions. They belong to a shared world of use, practice, habit, and expectation. Others can correct us: a door can be locked, a chair can collapse, a desk can be wrongly placed. Meaningful spatiality is not invented by a mind. It is disclosed through being-in-the-world before it is abstracted into coordinates.

Heidegger's word for the being that we ourselves are is *Dasein*. It does not name a psychological subject or biological species. It names the being for whom being is an issue, the being already involved in a world it did not first construct. *Dasein* is not first an inner mind facing an outer field. It is being-in-the-world.

World is equally strict. It is not the sum of objects but the referential whole in which things can show up as what they are. A hammer is intelligible through nails, timber, workbench, repair, hand, and need. A document is intelligible through office, institution, signature, promise, law, and risk. Heidegger calls this structure *Bedeutsamkeit*, significance. Things appear in places because they already matter.

This chapter protects that order. Place precedes position. Region precedes grid. Concern precedes measurement. Later, when language is converted into vector space, the order will matter again. A vector space can represent relations with great precision. It cannot, by that precision alone, recover the worldly and temporal conditions under which they first mattered.

2. Place before position

Heidegger's treatment of spatiality appears in §§22 to 24 of *Being and Time*, inside the analysis of worldhood. The placement is not incidental. He does not introduce space as an independent topic in physics or geometry, but through equipment, use, and the environment. Spatiality enters through the ready-to-hand, *Zuhandenes*, not through the present-at-hand, *Vorhandenes*.

The ready-to-hand is what is encountered in use. A pen, a cup, a book, a key, a scalpel, a keyboard, a legal file: each is first encountered through what one does with it. The object is there, but its being-there is not first neutral presence. It is availability within a practice.

Heidegger's German is exact: "Das »zur Hand« Seiende hat je eine verschiedene Nähe."² The entity "to hand" has each time a different nearness. I translate *Nähe* as "nearness" rather than only "closeness," because Heidegger is not describing a quantity of distance. He is naming the way something belongs within practical reach. A tool can be near because it is available for the task, even if it is not physically closest. It can be far because it has dropped out of the task, even if it lies on the desk.

This nearness is ordered through *Platz* and *Gegend*. *Platz* is place, not coordinate: the "there" where a thing belongs. *Gegend* is region, the oriented area of possible belonging. The hammer belongs at the bench, the file in the cabinet, the kettle in the kitchen, the passport in the drawer or bag where it can be found before departure. Such places are practical structures of a shared world, not subjective memories floating over objective space.

The difference becomes clear when something is missing. A misplaced key does not first appear as an unknown object absent from a grid. It appears as not where it belongs. The search does not proceed across homogeneous space. One checks the door table, pocket, bag, coat, desk, or last room used. The region becomes explicit because place has failed.

This failure teaches the structure. In ordinary use, the place often withdraws. The tool is simply there. The cup stands where it should, the screen is at the right height, the room works. When the tool is missing, broken, blocked, or misplaced, the referential order becomes visible. The world is spatial, but its spatiality is not first the space of geometric extension. It is the spatiality of belonging.

Heidegger's equipmental analysis is not a quaint account of workshops. The hammer is only the example. The same structure holds for a cockpit, hospital ward, courtroom, library, kitchen, and data center. Each has measurable dimensions. Each also has regions of access, danger, authority, repair, storage, and use. A sterile field in surgery is not reducible to square meters. A judge's bench is not reducible to height. A server rack is not reducible to coordinates. In each case, space is already organized by significance.

At the end of the book, this distinction returns. The data center can be mapped by coordinate, power load, latency, and cooling capacity. Such measures are necessary, but they do not exhaust its place. A data center belongs to land, grid, water, ownership, labor, contracts, and jurisdiction. To return AI to the data center will therefore not mean reducing language to hardware. It will mean refusing to let the apparent placelessness of the output conceal the placed infrastructure through which it appears.

The grid can represent this order after the fact: bench, cabinet, room, rack. This is useful. It does not found the order it represents.

3. Nearness and de-severance

Heidegger gives *Dasein* its own spatiality through two structures: *Ent-fernung* and *Ausrichtung*. The first is usually translated as "de-severance." The English is awkward, but the awkwardness is useful. Heidegger inserts the hyphen to hear the verb differently. *Ent-fernung* is not distance. It is the removing of farness, the bringing-near of what matters.³

Nearness is not metric proximity. The spectacles resting on one's nose may be physically closest, but they are phenomenologically far because they are not the object of concern. One looks through them. A colleague across the room may be nearer because the conversation is underway. A book in another city may be near because the work depends on it. A deadline tomorrow may be nearer than the object touching one's hand.

This structure is not poetic decoration. It belongs to ordinary action. We bring things near by reaching, searching, calling, opening, saving, citing, and repairing. The telephone, telescope, database, search engine, and model interface all intensify de-severance. They abolish certain forms of farness and make what is absent available under a new mode. They do not abolish finitude.

De-severance is not the same as possession. What has been brought near may still be misunderstood. A citation found instantly, a text retrieved, a person contacted, and a file opened have not thereby been read, interpreted, encountered, or judged. Technical nearness can hide hermeneutic distance.

The distinction becomes decisive for artificial intelligence. The model appears to abolish distance in language by bringing a tradition, discipline, question, or style into immediate availability. The user asks, and language arrives. Yet the arrival of language is not the recovery of the temporal locus from which that language first came. AI can intensify de-severance while deepening detachment.

De-severance belongs to *Dasein* because *Dasein* cares. Something is brought near because it matters for a task, worry, promise, possibility, or future. Nearness is not pure access. It is access within concern.

The mathematical system can reduce farness by retrieving, ranking, clustering, and generating. It can make tokens available with extraordinary speed. But availability is not concern. A system can calculate proximity without anything being near for it.

4. Directionality and the oriented body

De-severance is inseparable from *Ausrichtung*, directionality. To bring something near is already to be oriented. The world is encountered through whither and whence, left and right, above and below, before and behind, inside and outside. These directions do not begin as axes on a coordinate plane. They begin in oriented being-in-the-world.

Heidegger writes that “Ent-fernung und Ausrichtung” determine “die Räumlichkeit des Daseins.” De-severance and directionality determine the spatiality of *Dasein*. The passage places spatiality inside *In-Sein*, being-in, rather than inside the extension of an object. *Dasein* is spatial because it lives from orientation, not because it occupies a measurable volume.

This point has a bodily dimension that Heidegger does not develop enough. A human being does not orient itself as a pure mind. It stands, turns, tires, reaches, loses balance, learns paths, and carries habits in the body. The right hand, the familiar staircase, the blind turn in a road, the keyboard learned by touch, the distance from bed to door in darkness: these are not inferred coordinates. They are embodied orientations.

Heidegger’s early account gives the existential structure, but it leaves the lived body too thin. Later phenomenology, especially Husserl’s account of the *Leib* and Merleau-Ponty’s account of the body schema, can correct this thinness without undoing the basic point.⁴ Spatial orientation is not first a representation of an external grid. It is the bodily and practical opening of a world.

The body also shows why direction is not merely private. A person can be taught where to stand, how to hold a tool, how to enter a room, how to approach a patient, how to address a court, or how to sit at a piano. Orientation is learned within forms of life. It is bodily, social, and historical at once.

Computational space also has direction: vectors, dimensions, clusters, paths, gradients, and transformations. These formal structures can capture analogies and relations among linguistic traces. But vector direction is not orientation in Heidegger’s sense. It has no hand, no task, no exposure, no place from which the world opens. It has direction without being directed toward a possibility it owns.

This is not an engineering weakness. It is a difference in kind. Formal direction belongs to mathematical relation. Existential direction belongs to a being that can be underway in a world.

5. The Cartesian inheritance

Heidegger's critique of Cartesian space is not hostility to science. The issue is not whether measurement works, but what happens when the measurable form of space is treated as the first and truest one.

Descartes defines corporeal substance through extension, *res extensa*.⁵ The body is what can be measured in length, breadth, and depth. This move gives modern science a powerful abstraction. It also strips the world of its practical and historical texture. A thing becomes what is present-at-hand, what can be determined, observed, and measured apart from its place in concern.

Heidegger's terms are careful here. The present-at-hand, *Vorhandenes*, is not false. It is a possible way beings can appear. A broken hammer can become an object of inspection, a medical instrument can be studied chemically, a building measured structurally, a data center mapped by power load, cooling capacity, and network latency. These modes disclose real features.

They become distorting only when taken as the ground. The hammer is not first a physical body with later use added. The instrument is not first a neutral object with later care added. The room is not first a geometric container with later social meaning added. In ordinary being-in-the-world, things are encountered through significance. Present-at-hand determination is a modification of that more primary encounter.

This has direct consequences for any theory of representation. A map is not the territory, but the problem is sharper than that familiar phrase suggests. A map is not a place. It preserves certain relations by removing others. It shows road, scale, boundary, elevation, or ownership, depending on its purpose. It does not preserve the full world of dwelling, fear, memory, obligation, and use.

The same is true of every formal representation. A floor plan can guide construction while omitting inhabitation. A calendar can coordinate time while omitting the weight of waiting. A profile can represent a career while omitting relation to risk, loyalty, ambition, failure, and timing. A vector can represent linguistic traces while omitting the temporal-existential event in which language first mattered.

The Cartesian inheritance is not only a theory of physical bodies. It is a habit of thought. It teaches us to begin with the stripped object, the measurable field, the neutral coordinate, and then to rebuild meaning from there. Heidegger's counter-claim is that such rebuilding always arrives too late. Meaning is not added after spatial neutralization. Spatial neutralization is an abstraction from a world already meaningful.

The grid has false ease because it has already removed what made the world difficult.

6. §70 and the temporal ground

The argument becomes sharper in §70 of *Being and Time*. The section bears the title “Die Zeitlichkeit der daseinsmäßigen Räumlichkeit,” the temporality of the spatiality proper to *Dasein*. The phrase *daseinsmäßig* matters. Heidegger is not asking how physical objects occupy space, but how the spatiality of being-in-the-world is ontologically possible.

At this point in *Being and Time*, temporality, *Zeitlichkeit*, has already been named as the meaning of care, *Sorge*. *Dasein* is not temporal because it exists inside clock-time. It is temporal because it projects itself toward possibilities, takes over what it has been, and makes present what it encounters. Heidegger calls these the temporal ekstases: future, having-been, and present. They are not three boxes but the unified structure of a finite life standing outside itself toward what it can be, what it has inherited, and what it now encounters.

§70 asks whether spatiality remains an independent limit to this analysis. Perhaps *Dasein* is temporal and spatial in two separate ways. Perhaps the analysis of time reaches a boundary when it meets space. Heidegger refuses that separation. The spatiality of *Dasein*, he says, must be founded in temporality.

The strongest formulation appears near the end of the section: “Die Welt ist nicht im Raum vorhanden.”⁶ The world is not present-at-hand in space. I translate *vorhanden* here as “present-at-hand,” not simply “present,” because Heidegger is denying a specific ontological picture. The world is not an object or container lying there inside space. Space is discovered within a world.

The claim follows: meaningful spatiality is not grounded in the grid. It is grounded in the temporal being-in-the-world of *Dasein*. De-severance requires future-directed concern: something is brought near because a possibility is underway. Directionality requires thrown orientation: one already belongs to a world with paths, tools, languages, habits, and risks. Making-present requires the encounter in which the thing appears as available, blocked, missing, near, or far.

A hammer belongs at the bench because a task, history of use, and current situation gather around it. A platform is near or far because one is waiting, late, leaving, returning, or stranded. A book is near because a question has made it relevant; a database file because the search has made it available. Spatiality arises through time because practical space is opened by projects, inheritance, and present concern.

This does not mean that space is an illusion. Heidegger says explicitly that the argument is not meant to deduce space from time or dissolve space into pure time. Space can have relative independence at the level of physical and mathematical determination. Yet the spatiality in which a human world is lived cannot be understood apart from temporality. Space is discovered through the temporal structure of care.

For the critique of vectorization, the distinction is decisive. Vector space can represent relations among words because language has spatializable structure. Words belong near other words, uses cluster, and terms take positions inside fields of practice. But those relations first arise from temporal, historical, world-involved speaking and writing. A vector space can formalize the spatial trace. It cannot, from that trace alone, generate the temporal ground.

7. The dispute over grounding

The §70 claim is not simple. It has drawn serious disagreement because Heidegger's own text seems to pull in more than one direction.

One reading makes the claim very strong. Heidegger's project in *Being and Time*, on this view, requires a form of temporal idealism. Temporality is the final horizon for the understanding of being, and spatiality must be derived from it. William Blattner gives this reading its most exact form, arguing that Heidegger grounds existence in temporality so completely that space becomes dependent on time in a strict sense. The problem, for Blattner, is that the derivation does not succeed: material and spatial features resist reduction to temporal projection.

A second reading is weaker. It holds that Heidegger shows spatiality to depend on temporality for its existential meaning, without reducing spatiality to time. This reading is closer to Jeff Malpas's topological interpretation.⁷ Place is not a mere derivative of time. Situatedness has its own weight. The human being is temporal, but it is also placed. World is not only projection; it is where one finds oneself, with others, among things, in a bounded and meaningful site.

A third approach emphasizes Heidegger's later revision. The later Heidegger increasingly thinks through *Ort*, place, *Einräumen*, making room, *Zeit-Raum*, time-space, and *Ereignis*, the event of appropriation. The strict priority of temporality gives way to a topology in which time and space belong together in the clearing of being. This later shift does not restore Cartesian space. It changes the form of the anti-Cartesian claim.

The present argument does not need to settle the full scholarly dispute. It needs a precise position: §70 should be read not as a successful deduction of spatiality from temporality in the strong sense, but as a transcendental argument against the priority of abstract space.

The term "transcendental" is used here in its Kantian sense. A transcendental argument asks for the conditions that make something possible. Heidegger changes the content of the question. He does not ask for the a priori forms of a detached subject's cognition, but for the existential conditions under which beings can be encountered as meaningful. In this sense, the temporal grounding of spatiality is not a claim about physics. It is a claim about the conditions under which space can matter.

The strong version says: space is derived from time. The weaker version says: meaningful spatiality presupposes temporal being-in-the-world. The weaker version is enough for the argument of this book. It preserves the structural force of §70 without requiring the collapse of place, body, and material resistance into time.

The critique of vectorization should not rest on a fragile reading. If the book claimed that space itself is nothing but time, it would inherit all the problems of the strongest §70 thesis. It needs only this: the spatial structures that carry meaning are not self-grounding. They depend on temporal, historical, embodied, and practical conditions.

Vectorization treats spatial relation as if it could stand alone. The Heideggerian claim denies that self-sufficiency.

8. Place after §70

The later Heidegger gives place a stronger role than the early analytic allows.

In “Building Dwelling Thinking,” Heidegger no longer begins with equipmental place alone. He thinks building through dwelling, and dwelling through the mortal human being’s way of staying on the earth. The bridge does not simply occupy a location. It gathers banks, stream, paths, use, and human passage. A location comes to be through the thing that gathers. Place is not a point filled by an object. Place happens through gathering.

In “Art and Space,” the formulation becomes even sharper: “Orte öffnen Gegenden.”⁸ Places open regions. This is not a return to the early equipmental *Platz* alone. It is a later topology of *Ort*, a place that gathers and grants room. Mathematical space is not denied. It is displaced from first rank. Space as measurable extension becomes an abstraction from the more originary happening of place.

The later view also sharpens the critique of technology. Under *Gestell*, enframing, beings appear as *Bestand*, standing-reserve. A river appears as hydroelectric supply; a forest as timber; a field as yield; a person as labor capacity, consumer profile, behavioral signal, or data source. Space becomes logistical: calculable, orderable, exchangeable, optimized.

The danger is not calculation itself. Calculation is necessary for bridges, hospitals, transport, and data centers. The danger lies in the exclusive rule of calculation, when place is reduced to position and beings appear only as resources.

The later Heidegger helps the present argument in two ways. First, it corrects the thinness of a strict temporal derivation: human existence is dwelling, place, body, exposure, and material relation, not only temporal projection. Second, it confirms the critique of abstract calculability. Mathematical space becomes dangerous when it conceals the place from which it was abstracted.

The book's thesis can therefore be stated without forcing early and late Heidegger into false unity. Early Heidegger shows that the spatiality of being-in-the-world is grounded in temporal care. Later Heidegger shows that place has gathering power that cannot be derived from geometric space. Both resist the same reduction. Meaningful space is not a neutral grid.

This also protects the argument from nostalgia. The point is not to return from vector space to a premodern village, from data centers to bridges, or from computation to handicraft. The point is to preserve the difference between formalized spatial relation and the temporal-place structure from which meaning arises. Technical systems may be used, but their formal spaces must not be mistaken for worlds.

9. The derivative space of the vector

A vector space is a mathematical construction. In language models, tokens and their relations are represented in high-dimensional space.⁹ Similarity becomes calculable as proximity. Difference becomes distance, direction, weight, distribution, and transformation. Such systems can capture patterns no human reader could survey unaided. Their success is not accidental. Language already carries relations. Human worlds sediment themselves in use. Words keep company because practices, institutions, genres, conflicts, and forms of life have given them company.

A crude critique would say that vector space is empty because it is mathematical. That would be false. Vectorization preserves something real: a formal trace of linguistic belonging. It can map analogies, styles, topics, co-occurrences, entailments, and patterns of use. It can orient a reader because the corpus contains traces of human worlds.

The problem lies in the change of ground. In Heidegger's account, spatiality is meaningful because it is opened through being-in-the-world. In vectorization, linguistic relation is formalized as mathematical proximity after language has been detached from speech, writing, use, and reading. The model inherits the spatial shell of language. It does not inherit the temporal locus from which the language first came.

The word "locus" will carry increasing weight in the chapters that follow. It names more than position: the temporal, historical, bodily, practical, and responsible site from which language is uttered, written, read, or owned. A human speaker has a locus because they speak from a time, place, body, history, concern, and relation to consequence. A written text preserves traces of such a locus, and a reader can return to it through interpretation. Vectorization changes the relation. It preserves formal markers and relational traces while weakening or severing the path back to source loci. The problem is not only that the model has no speakerly locus of its own. It is that the language it generates descends from source loci that have been aggregated, formalized, and often made unreconstructable.

Position can be computed. Locus must be inhabited, reconstructed, or responsibly acknowledged as lost.

Chapter 1 establishes the ground. Spatiality is not first a grid. The space that matters for meaning is opened by temporal, placed, concerned being-in-the-world. Mathematical space may be powerful, exact, and useful. It is not the ground of worldhood.

Vectorization will later be described as detemporalization because it converts temporal-existential traces into spatial calculability. It can do so only because language is already spatial in the existential sense: oriented, placed, directed, and full of relations. It fails when it treats that spatiality as self-sufficient. The system can map relations among traces. It cannot become the temporal ground that made the traces meaningful.

The next chapter turns from spatiality to temporality itself. The grid must be displaced before time can be understood. Otherwise temporality will be mistaken for another axis, another line, another sequence of points inside a space that was never first.

NOTES

1. Martin Heidegger, *Sein und Zeit*, 11th ed. (Tübingen: Max Niemeyer, 1967), §§12, 14–18, 22–24; Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §§12, 14–18, 22–24. Heidegger's analysis of worldhood in §§14–18 establishes the referential whole of significance, while §§22–24 develop the spatiality of the ready-to-hand and of *Dasein*. On Heidegger's worldhood and equipmentality, see Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I* (Cambridge, MA: MIT Press, 1991), 60–101; Taylor Carman, *Heidegger's Analytic: Interpretation, Discourse, and Authenticity in Being and Time* (Cambridge: Cambridge University Press, 2003), 121–157.
2. Heidegger, *Sein und Zeit*, §§22–24, especially 102–112; Heidegger, *Being and Time*, 135–148. The German sentence quoted in the main text appears at *Sein und Zeit*, 102: "Das »zur Hand« Seiende hat je eine verschiedene Nähe, die nicht durch Ausmessen von Abständen festgelegt ist." Macquarrie and Robinson render this at *Being and Time*, 135: "Every entity that is 'to hand' has a different closeness, which is not to be ascertained by measuring distances." On *Platz* and *Gegend*, see *Sein und Zeit*, 102–104; *Being and Time*, 135–137.

3. Heidegger, *Sein und Zeit*, §23, especially 105–108; *Being and Time*, 138–142. Macquarrie and Robinson translate *Ent-fernung* as “de-severance” and explain the translation in their note at *Being and Time*, 138. Heidegger’s examples of spectacles, street, and acquaintance appear at *Sein und Zeit*, 107; *Being and Time*, 141–142. On *Ausrichtung*, see *Sein und Zeit*, 102, 108; *Being and Time*, 135, 142–143.
4. Edmund Husserl, *Ideen zu einer reinen Phänomenologie und phänomenologischen Philosophie. Zweites Buch: Phänomenologische Untersuchungen zur Konstitution*, ed. Marly Biemel, Husserliana IV (The Hague: Martinus Nijhoff, 1952), especially §§35–42; English in *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, Second Book: Studies in the Phenomenology of Constitution*, trans. Richard Rojcewicz and André Schuwer (Dordrecht: Kluwer, 1989). Husserl treats the lived body as the zero-point of orientation in Hua IV, §41. See also Maurice Merleau-Ponty, *Phenomenology of Perception*, trans. Donald A. Landes (London: Routledge, 2012), Part One, chapter 3, “The Spatiality of One’s Own Body and Motility.”
5. René Descartes, *Principia philosophiae* II, §§4, 10–11; English in *The Philosophical Writings of Descartes*, vol. 1, trans. John Cottingham, Robert Stoothoff, and Dugald Murdoch (Cambridge: Cambridge University Press, 1985), 223–225. Heidegger’s critique of the Cartesian ontology of the world appears in *Sein und Zeit*, §§19–21; *Being and Time*, §§19–21. The point in this chapter is not that Cartesian mathematical space is useless, but that Heidegger treats it as ontologically derivative from the meaningful spatiality of being-in-the-world.
6. Heidegger, *Sein und Zeit*, §70, especially 367–369; *Being and Time*, 418–421. The German title of §70 is “Die Zeitlichkeit der daseinsmäßigen Räumlichkeit.” The sentence quoted in the main text, “Die Welt ist nicht im Raum vorhanden,” appears at *Sein und Zeit*, 369. A fuller formulation from the same passage reads: “Nur auf dem Grunde der ekstatisch-horizontalen Zeitlichkeit ist der Einbruch des Daseins in den Raum möglich. Die Welt ist nicht im Raum vorhanden; dieser jedoch läßt sich nur innerhalb einer Welt entdecken.” Macquarrie and Robinson translate this at *Being and Time*, 420–421.
7. William D. Blattner, *Heidegger’s Temporal Idealism* (Cambridge: Cambridge University Press, 1999), especially 197–235; Jeff Malpas, *Heidegger’s Topology: Being, Place, World* (Cambridge, MA: MIT Press, 2006), especially 111–146. Blattner gives the strongest version of the temporal-grounding problem; Malpas develops the topological correction by emphasizing place, situatedness, and the irreducibility of topology in Heidegger’s thought. See also Edward S. Casey, *The Fate of Place: A Philosophical History* (Berkeley: University of California Press, 1997), 243–284.

8. Martin Heidegger, “Building Dwelling Thinking,” in *Poetry, Language, Thought*, trans. Albert Hofstadter (New York: Harper & Row, 1971), 143–161, especially 152–154 on the bridge and the coming-to-be of location; Martin Heidegger, “Art and Space,” trans. Charles H. Seibert, *Man and World* 6, no. 1 (1973): 3–8. The German phrase “Orte öffnen Gegenden” comes from “Die Kunst und der Raum.” On later Heidegger’s topology and the relation of place to *Ereignis*, see Malpas, *Heidegger’s Topology*, 273–324; Heidegger, “Time and Being,” in *On Time and Being*, trans. Joan Stambaugh (New York: Harper & Row, 1972), 1–24.
9. For the technical lineage that later chapters develop in detail, see Zellig S. Harris, “Distributional Structure,” *Word* 10, nos. 2–3 (1954): 146–162; J. R. Firth, “A Synopsis of Linguistic Theory, 1930–1955,” in *Studies in Linguistic Analysis*, special volume of the Philological Society (Oxford: Blackwell, 1957), 1–32; Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean, “Efficient Estimation of Word Representations in Vector Space,” arXiv:1301.3781 (2013); Jeffrey Pennington, Richard Socher, and Christopher D. Manning, “GloVe: Global Vectors for Word Representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Stroudsburg, PA: Association for Computational Linguistics, 2014), 1532–1543; Ashish Vaswani et al., “Attention Is All You Need,” in *Advances in Neural Information Processing Systems 30*, ed. Isabelle Guyon et al. (Red Hook, NY: Curran Associates, 2017), 5998–6008.

Chapter 2. Ekstatic Time and Technical Memory

1. Time is not another axis

The previous chapter displaced the grid. Space was not denied. It was returned to its existential ground. Human beings do not first inhabit neutral extension and then assign meanings to positions inside it. They dwell in places, move through regions, bring things near, turn toward tasks, and find themselves oriented before measurement begins. Measurable space is real. It is not first.

The same correction is needed for time.

The ordinary image of time is a line. Moments succeed one another. The present is a point. The past lies behind. The future lies ahead. Calendars divide the line into days; clocks divide days into hours, minutes, and seconds. Timelines become the diagrammatic form of historical thought. A life becomes a sequence; a tradition, a chain; a text, a date; an event, an entry.

This image is useful because human beings must coordinate. Trains depart, courts convene, medicines are administered, contracts expire, servers synchronize, and financial systems settle. Clock-time is not a mistake. It is one public form by which a shared world becomes manageable.

Its usefulness becomes dangerous when it is treated as the first truth of time. The measured sequence then replaces the temporal structure that made measurement possible. Time becomes another spatial order: points on a line, units in a container, coordinates in an abstract field. This is the temporal companion to the spatial grid.

Heidegger calls this the ordinary or vulgar concept of time. “Vulgar” here means everyday, leveled, publicly available, not morally crude. Vulgar time is the sequence of now-points, *Jetztfolge*: the past as no-longer-now, the future as not-yet-now, the present as the now that passes. The image is so familiar that it can appear self-evident.

It is not self-evident. It is an abstraction from a more originary temporal structure. A person does not first encounter neutral now-points. A person waits, remembers, expects, delays, inherits, regrets, promises, anticipates, and dies. Time is first lived as concern. It is not first counted.

This chapter gives that temporal ground in three steps. Heidegger’s account of *Zeitlichkeit*, temporality, shows that *Dasein* is temporal in its being, not merely located in time. Husserl and Stiegler show that memory is not only interior: human temporality is bound to traces, inscriptions, supports, and media. Finally, the training corpus appears as aggregate tertiary retention, a technical gathering of written traces that prepares language for vectorization.

The full thesis that vectorization is detemporalization belongs later. Here the task is to show what is at stake when language is detached from the temporal structures of care, historicity, memory, and reading.

2. Care and the unity of time

Heidegger’s analysis of time begins from *Sorge*, care. Care does not mean worry, kindness, or emotional concern. It names the being of *Dasein*: ahead of itself, already in a world, and alongside what it encounters there. Heidegger’s formula is dense: *Sich-vorweg-schon-sein-in einer Welt als Sein-bei innerweltlich belegendem Seienden*. *Dasein* is ahead-of-itself, already-being-in a world, as being-alongside beings encountered within the world.

The formula names the whole movement of human existence.

A person is ahead of themselves because they live toward possibilities. They are not finished things. They are already oriented toward what can be done, refused, lost, received, or become. This is not only explicit planning. A person sitting quietly still inhabits possibilities: rest, avoidance, recovery, waiting, fear, expectation, boredom, prayer, fatigue.

A person is already in a world because they never begin from pure self-possession. They are thrown into language, family, history, institutions, bodies, debts, customs, tools, and names. The past is not simply behind them. It is the condition under which the present can appear at all.

A person is alongside beings because they are involved with what they encounter. They use, repair, listen, answer, neglect, read, write, bury, protect, and judge. The present is not a punctual point. It is the making-present of a world in which something matters now.

Heidegger's decisive move is to show that care is temporal. He writes: "Die ursprüngliche Einheit der Sorgestruktur liegt in der Zeitlichkeit." The original unity of the structure of care lies in temporality. The English "lies in" is deliberately plain. Heidegger is not saying that care occurs in time as an event occurs in a calendar. He is saying that care is intelligible only because *Dasein* temporalizes itself. The three moments of care correspond to the three temporal ekstases: future, having-been, and present.¹

Ekstasis comes from Greek and names a standing-outside. Heidegger uses it to avoid the image of time as a container of moments. Temporality is not a thing that first sits inside itself and then stretches outward. It is this standing-outside itself. Heidegger's German formulation is severe: "Zeitlichkeit ist das ursprüngliche »Außer-sich« an und für sich selbst." Temporality is the original outside-itself in and for itself.²

Time is not a line on which a life is placed. A life is temporal because it stands outside itself toward its own possibilities, back upon what it has been, and among the beings it now encounters. The unity of life is not the sum of moments. It is the unity of this temporal standing-out.

The book therefore uses “temporal locus” rather than “time point.” A human utterance does not arise from a coordinate in a sequence. It arises from a being whose present is made possible by inheritance and projection. The utterance has a date; its date is not its temporal ground.

3. Future, having-been, present

Heidegger gives priority to the future. This priority is easily misunderstood. It does not mean that the future is more real than the past or present. It means that *Dasein* is defined by possibility. It is always at stake for itself. Its being is something it can take up, lose, evade, or own.

Heidegger writes: “Das primäre Phänomen der ursprünglichen und eigentlichen Zeitlichkeit ist die Zukunft.” The primary phenomenon of original and authentic temporality is the future.³ The word *Zukunft* does not mean a later now. It names the coming-toward-itself of *Dasein* in its possibilities. The future is not simply what will happen. It is the structure by which a human being can be ahead of itself.

This future is not empty choice. *Dasein* projects itself only as already thrown. It does not choose from nowhere. It inherits languages, customs, bodies, damages, obligations, desires, and limits. Heidegger names this having-been, *Gewesenheit*. The past is not merely what has elapsed. It is what *Dasein* already is, the inherited facticity from which it projects itself.

The present is *Gegenwart*, but Heidegger often speaks more precisely of *Gegenwärtigen*, making-present. The present is not a point between past and future. It is the temporal mode in which beings are encountered as available, relevant, obstructive, beloved, dangerous, or indifferent. Making-present belongs to concern. A file becomes present because the task requires it. A remembered phrase becomes present because the conversation calls it forth. A person becomes present not because they occupy space nearby, but because one is answerable to them.

The three ekstases are not separate parts. They are a unity. Projection without having-been would be fantasy from nowhere. Having-been without projection would be dead deposit. Present without either would be isolated sensation. Human temporality is the unity in which these three stand together.

This unity will later matter for AI output. A model can generate text that marks future, past, and present: “I will,” “I remember,” “this is happening now,” “as argued earlier.” These temporal forms can be grammatically correct. They do not prove that the system has a future, a having-been, or a present in Heidegger’s sense. They are surface markers of temporality without temporal existence.

The difference is not minor. It is the difference between using the grammar of time and existing temporally.

4. Finitude

Temporality is finite. Heidegger’s account of time cannot be separated from being-toward-death. Death is not an event waiting at the end of a timeline. It is the possibility of the impossibility of further possibilities. It is the limit that belongs to *Dasein* as its own.

Finitude gives time its weight. A person can delay because not every delay is reversible. A promise matters because it binds a finite life. A decision matters because not all possibilities remain open. Regret matters because one cannot return unchanged to the moment before the act. Learning matters because formation has limited time.

Without finitude, future would become an endless field of interchangeable options. Having-been would lose the force of inheritance. Present concern would lose its pressure. Finitude is not an added theme. It is the structure that makes temporal existence serious.

Heidegger's account of historicity makes this explicit. In §74, he says that *Dasein* has its history because its being is constituted by historicity: "Das Dasein hat faktisch je seine »Geschichte« und kann dergleichen haben, weil das Sein dieses Seienden durch Geschichtlichkeit konstituiert wird." *Dasein* factually has its history and can have such a thing because the being of this being is constituted by historicity.⁴

History is not first a science of past facts. It is the way a finite being inherits and takes over possibilities. Heidegger's terms here carry political and existential charge: *Schicksal*, fate; *Geschick*, destiny; *Wiederholung*, repetition. The heroic tone can be left aside. The ontological claim remains: a human being is historical because it must take up a past it did not choose.

Heidegger writes: "Nur eigentliche Zeitlichkeit, die zugleich endlich ist, macht so etwas wie Schicksal, das heißt eigentliche Geschichtlichkeit möglich." Only authentic, finite temporality makes possible something like fate, that is, authentic historicity.⁵ The translation "fate" for *Schicksal* is traditional but risky. It can sound like fatalism. Heidegger's point is not that events are predetermined. It is that finite existence can take over an inherited possibility as its own.

A text therefore belongs to a having-been. It is not only a sequence of signs. It comes from a world, a question, a language, a conflict, an institution, an authorial act, or a tradition. To read historically is to let the having-been of the text address the reader's present and alter the reader's possible future.

A corpus does something different. It gathers traces of many having-beens and makes them simultaneously available. This availability is powerful and risky. Temporal difference can be flattened into searchable presence. A medieval commentary, a policy memo, a forum post, a scientific paper, and a legal opinion may all become equally available as data. Their dates remain. Their temporal loci recede.

5. Repetition and inheritance

Heidegger's *Wiederholung*, repetition, is not copying. It is not the restoration of the past. It is the retrieval of an inherited possibility in a present situation.

This matters because the book will later distinguish reading from processing. Reading does not simply consume a stored trace. It can repeat by retrieving a possibility from a past text and testing it in the present. This is not antiquarian recovery. The reader does not leave their time behind. They read from a present locus toward a source locus and allow the distance between the two to work.

Historicality should not be confused with pastness. A historical locus may belong to an ancient text, a dead author, or a tradition from another century. It may also belong to a contemporary world that is not ours. A living language, state form, political vocabulary, or professional practice may stand beside us in the present and still remain historically other. The difficulty of understanding is not only chronological distance. It is the distance of formation.

To enter such a locus requires more than information. It may require language, bodily presence, apprenticeship, humiliation, correction, friendship, and the irreversible alteration of one's own horizon. This too is historical positioning.

The past, then, is not a warehouse or complete inventory waiting for extraction. It is a field of inherited possibilities, some dead, some dangerous, some still capable of address. A tradition becomes living when it is repeated in this strong sense. It becomes deadened when it is merely cited, indexed, or recombined without retrieval.

This distinction prepares the later account of *intus legere*, reading within. To read between the times is to hold at least two temporal loci together: the reader here and now, the text there and then. Understanding does not erase the distance. It works through it.

The model cannot repeat in this sense. It can reproduce style, summarize tradition, quote, compare, and continue plausibly. These uses can help. But the model does not take over an inherited possibility as its own. It has no ownmost future into which such a possibility could be taken.

This is not a romantic claim about human superiority. Humans often fail to repeat in the strong sense. They quote without reading, inherit without thinking, and repeat the past as formula. The point is narrower. The structure of genuine repetition requires a finite historical being. Without finitude, there is no inheritance. Without inheritance, there is only storage.

6. World-time and the clock

Originary temporality does not remain hidden in a private interior. It is expressed in ordinary time-reckoning: now it is time to begin, later it will be too late, tomorrow the work must be delivered, this season is wrong. These expressions do not begin as neutral measurements. They belong to care.

Heidegger calls this public, practical time *Weltzeit*, world-time. World-time is time as it is encountered in the world of concern. It is datable, stretched, public, and significant. It is tied to tasks, occasions, deadlines, rituals, journeys, meals, meetings, worship, markets, harvests, elections, and grief.

Heidegger's phrase is exact: "Die veröffentlichte Zeit hat als Zeit-zu ... wesenhaft Weltcharakter." Published or public time, as time-for, has essentially the character of world. The phrase *Zeit-zu* is important. Time is time for something: work, sleep, speech, silence, departure, return. The clock can measure duration, but the world gives the "for what."⁶

A clock abstracts from this structure. It makes time publicly comparable and allows coordination among people who do not share the same immediate situation. Clocks and calendars are therefore not merely instruments. They reorganize the social experience of time and make certain forms of precision, obligation, and control possible.

The abstraction has a cost. When time is treated only as a sequence of identical units, its datability and significance recede. The hour becomes interchangeable, the day a block, the worker scheduled, the student measured, the user available. Public time can then appear as no one's time.

Heidegger sees this leveling in everyday publicness. He writes: “Die Zeit ist öffentlich etwas, was sich jeder nimmt und nehmen kann.” Time is publicly something that everyone takes and can take. The leveled sequence of nows belongs to everyone and no one.⁷ Here public time becomes the temporal counterpart to *das Man*. It disburdens and coordinates. It also levels.

Contemporary AI operates inside a world already governed by synchronized, measured, public time. Prompts are answered in seconds; workflows demand speed; search becomes instant; drafting becomes immediate. The system enters a temporal regime that values availability over formation. It does not create that regime. It serves and intensifies it.

7. Husserl’s living present

Heidegger gives the existential structure of temporality. Husserl gives the phenomenological micro-structure of temporal consciousness. The two accounts are not the same, and the present argument need not harmonize them. It needs the pressure both place on any flattened model of time.

Husserl’s central case is the temporal object, especially the melody. A melody is not heard as a sum of isolated notes. If consciousness grasped only a punctual now, it would hear one note after another without melody. The note just heard must be retained as the next note arrives, and the next note anticipated, even if indeterminately. The present must therefore include more than a point.

Husserl names the structure *Urimpression*, primal impression; *Retention*, retention; and *Protention*, protention. The primal impression is the now-phase of appearing. Retention is not recollection of a past object but the holding-on of the just-past within the living present. Protention is not inference but anticipatory openness toward what is about to arrive.

The living present is therefore thick. It is not an instant. It is a dynamic field in which the just-past is retained and the about-to-come is anticipated. Without this structure, no duration could be experienced, no spoken sentence heard as a sentence, no gesture followed, no text read in sequence.

This gives the first reason to resist the reduction of time to sequence. A sequence of tokens is not yet temporal consciousness. A technical system may process elements, store prior tokens in a context window, and predict continuations. These operations are temporal in a technical sense. They do not thereby constitute retention and protention as lived structures.

Husserlian retention and protention are not mental objects placed beside one another. They are non-independent moments of the flow that make present, past, and future possible for consciousness. Zahavi's interpretation resists the objectifying reading of Husserl: retention, primal impression, and protention are not items located in time. Their conjunction makes temporal appearing possible.⁸

A context window is therefore not a horizon of time-consciousness. It can preserve prior tokens, condition future tokens, and allow textual continuity. It cannot be affected by the just-past as a lived modification of presence or anticipate a future in which it is at stake. It has sequence without a living present in the strong sense.

AI systems are often described with the language of memory. The description has limited technical use. It becomes misleading when memory is heard phenomenologically. A stored vector, a retrieved document, a persistent user profile, or a longer context window is not memory in Husserl's sense. It is stored or accessible material. Memory begins when a temporal subject retains, recalls, and is altered by what has been.

8. Sedimentation and reactivation

Husserl's late text "The Origin of Geometry" changes the problem. It no longer asks only how temporal objects are constituted in consciousness. It asks how an ideal object can be handed down through history. Geometry is not tied to one geometer's private act. Once founded, it can be taught, corrected, extended, and reactivated across generations.

This transmission requires inscription. Speech alone cannot secure the historical durability of ideality. Writing allows an originary act of sense to be preserved beyond the life of the founder. A theorem can be written, stored, copied, forgotten, rediscovered, and taught. The ideal object can outlive its first inscription because writing makes it available for reactivation.

Husserl's problem is not antiquarian. He asks how truth can have a history without becoming merely factual. A geometrical proposition is ideal. It is not exhausted by the empirical mark on the page. Yet without inscription and transmission, it would not become a historical object for a community.

Husserl's key terms are sedimentation and reactivation. Meaning sediments when an originary evidence is fixed in a sign and handed down. Sedimentation is necessary because it preserves. It is dangerous because it can be received passively. A formula may be repeated without the evidence that first gave it sense. Reactivation returns to the sedimented sense and makes it evident again.

Husserl's phrase is demanding: "die Möglichkeit der vollständigen und echten Reaktivierung der vollen Ursprünglichkeit," the possibility of the complete and genuine reactivation of full originality.⁹ He does not want a mere repetition of words. He wants the reader to return through the written trace to the originating evidence that gives the ideal object its sense.

This structure will govern the book's treatment of written sources. A written text is not living speech, but it is not dead because it is written. It preserves the possibility of reactivation. A reader can return to the text, the source, the argument, the world, and the question. The written trace detaches language from immediate speech without severing all relation to locus.

The danger begins when sedimentation is mistaken for understanding. A person can repeat a theorem, cite a doctrine, quote a passage, or summarize a position without reactivating the sense. AI inherits an old human danger and accelerates it: the outer form of reactivation without the evidential return.

The model can state a theorem, summarize Husserl's argument, and produce a fluent account of sedimentation and reactivation. But it does not reactivate evidence. It operates on traces of prior reactivations and failed reactivations alike. It may assist a human reader in returning to the source. It cannot perform that return as its own.

9. Stiegler and tertiary retention

Stiegler radicalizes the problem of technical memory. Husserl had recognized the necessity of writing for the historical transmission of ideality, but he still tended to subordinate the technical trace to the living act that reactivates it. Stiegler shifts the balance. Technical supports are not secondary containers for an already complete human temporality. They help constitute human temporality from the beginning.

His decisive term is *réention tertiaire*, tertiary retention. Husserl had distinguished primary retention, the just-past held in the living present, from secondary retention, recollection. Stiegler adds a third kind: memory exteriorized in a material support. A tool, mark, book, photograph, film, database, archive, calendar, recording, or hard drive can function as tertiary retention.

Tertiary retention is outside consciousness, yet it shapes consciousness. A child is born into a world filled with recorded words, images, schedules, rituals, maps, tools, names, laws, and devices. These traces do not merely help the child remember later. They shape what can be perceived, remembered, expected, and desired.

Stiegler's wider term for this structure is epiphylogenesis. The word is difficult, but its sense is exact. Genetic memory belongs to the species. Epigenetic memory belongs to the individual life. Epiphylogenetic memory is the accumulation of individual experience in technical supports that can be transmitted beyond one life. A tool can preserve a gesture; a book, a thought; a recording, a voice; a diagram, a form of seeing.

Stiegler's compressed formulation is that technics is "the pursuit of life by means other than life." The phrase should not be romanticized. Human life extends itself through nonliving supports; it survives, learns, and inherits through matter that is not alive. This is not an external accident. It belongs to human temporality.¹⁰

The point is close to Heidegger and also a correction of him. Heidegger shows that *Dasein* inherits a having-been. Stiegler asks how that inheritance is materially possible. The answer is technics. Without exteriorized traces, inheritance would be fragile, local, and short. Tradition requires supports; memory requires matter.

This claim protects the book from a false opposition between living time and dead technology. Technical memory is not simply the enemy of temporality. Writing, archives, calendars, photographs, recordings, and digital systems all participate in the formation of human time. The problem is not exteriorization as such, but the form it takes and the practices by which human beings receive it.

Books can support reactivation; photographs can support mourning; calendars can support shared life; recordings can preserve voices; databases can assist research; models can support orientation. Each support has its own temporal discipline and danger.

The question is not whether memory should be technical. Human memory has always been technical. The question is what kind of technical memory is being formed, who controls it, how it selects, how it is read, and whether it returns the user to temporal responsibility or substitutes for it.

10. Media forms and temporal discipline

Different media do not preserve time in the same way. Writing, print, photography, cinema, broadcast, databases, and generative systems each create a different relation between trace and receiver.

Writing fixes a sequence of marks. It detaches language from the immediate speech situation and asks to be read. The reader must move through the sequence, pause, return, compare, correct, and interpret. Writing supports delay. It can also support mechanical repetition.

Print multiplies the written trace. It stabilizes editions, publics, citation, commentary, and dispute. It permits traditions across distance. It can also standardize and harden what counts as legitimate text.

Photography captures a past moment with a special evidential force. It gives the viewer the sense that this-has-been. It supports memory and grief. It can also detach the image from the life it marks.

Cinema and broadcast alter the flow of consciousness more directly. Stiegler calls cinema an industrial temporal object. Its time passes as the viewer's time passes, imposing a rhythm of retention and protention. Broadcast intensifies this by synchronizing many viewers at once.

Digital media add another mutation. The trace is no longer only stored. It is indexed, retrieved, recombined, ranked, predicted, and fed forward. The archive becomes dynamic. Selection can occur before the user asks. Attention becomes a target of technical anticipation.

Stiegler's political account of attention matters here. Industrial and digital systems can weaken the long circuits through which attention becomes care, memory becomes formation, and inheritance becomes thought. They can proletarianize attention, not by taking away information, but by taking over the temporal structures through which information is received.

The word "proletarianization" is strong. Stiegler uses it beyond factory labor. A worker is proletarianized when know-how is externalized into a machine and returned as a system they operate without understanding. Attention can undergo a related loss. A person may receive more content, reminders, prompts, summaries, and stimulation while losing the temporal capacity to synthesize, wait, remember, and judge.

Generative AI belongs here. It is not only a new instrument for producing text. It is a new temporal relation to language. It answers before the user has read, drafts before the user has formed the sentence, summarizes before the user has lived with the difficulty of the source. It gives language in the time of demand.

This does not make it bad in every use. Some delay is waste. Some access is humane. Some summaries help a reader begin. Some generated structures make inquiry possible under real constraint. The problem appears when the device repeatedly replaces the temporal passages through which understanding would have been formed.

Book two named this distinction through *Erlebnis* and *Erfahrung*. AI output can become an *Erlebnis* for the user: language arrives, a difficulty takes form, a question becomes answer-shaped. It does not become *Erfahrung* unless the user undergoes the return to source, evidence, encounter, and responsibility. This chapter gives the distinction its temporal ground.

11. Generated images and the appearance of retention

Generated images and video make this structure visible in another medium.

They are not the subject of this book. The book's argument concerns generated language: answer, summary, interpretation, dialogue-like response, and the sentence that appears to speak from a temporal and historical locus it does not inhabit. Yet the visual case belongs here because it strains the Stieglerian account of tertiary retention in a useful way.

A photograph or film has traditionally carried a documentary presumption. This presumption should not be romanticized. Photographs have never been pure evidence. They are framed, selected, staged, cropped, captioned, edited, printed, circulated, archived, institutionalized, and interpreted. A photograph can deceive. A film can fictionalize. A documentary image can be made to serve propaganda, sentimentality, surveillance, memory, mourning, or law. The image never simply gives the real without mediation.

Even after these qualifications, the photograph and the film have usually retained a relation to capture. Something was placed before a camera. Light, lens, sensor, surface, frame, apparatus, and operator belonged to an event of inscription. The image could lie about the event, isolate it, distort it, stage it, or detach it from context. But the image still claimed some relation to a captured occurrence.

Generated visual media alter that relation.

A generated image can take the form of a photographic trace without preserving a photographed event. It can show a street that was never there, a person who never stood there, or a meeting that never occurred. It can show a disaster no camera witnessed, a face with no life behind it, or a historical scene with no historical instant. It may inherit visual styles, statistical regularities, genre conventions, lighting patterns, camera angles, and documentary cues from prior images. In that sense it is not made from nothing. It descends from an aggregate of visual traces and from the technical systems that learned from them.

But it does not retain the event it appears to show.

This is capture without capture.

The phrase does not mean that every generated image is deceptive. Illustration, animation, painting, fantasy art, cinematic effects, and computer-generated imagery have long produced visible worlds without claiming documentary capture in the ordinary sense. Human beings have always made images of what was not before the eye. The novelty is more precise. A generated image can borrow the visual grammar of capture while lacking the event of capture. It can appear within the cultural surface of photographic evidence while belonging to another order of production.

It is therefore a special case of technical memory. It is not tertiary retention in the same sense as a photograph of a past moment. It does not exteriorize that moment into a support. It exteriorizes, rather, the statistical inheritance of prior visual culture into a new artifact that can look as if a moment had been retained. It retains patterns of capture, not necessarily a captured thing. It preserves the appearance of documentary relation while severing the relation on which that appearance traditionally depended.

The effect does not remain inside the individual image.

Once such images circulate at scale, the authentic image also changes its reception. A real photograph now appears in a medium in which “generated” has become an available suspicion. A real video must now sometimes prove that it was recorded. The witness must prove that there was witnessing. The event must prove that it had an image. This does not mean that photography has lost all evidential force. It means that the evidential practice has entered a new condition. Trust moves outward toward provenance, institutional source, metadata, platform, reputation, and verification.

The medium has not merely acquired new false images. It has acquired a new doubt about images as such.

The same transformation is harder to see in writing because writing has fewer perceptual marks of production. A generated sentence and an authored sentence can share the same surface. The page does not reveal whether the sentence came through hesitation, revision, reading, memory, and responsibility, or through the statistical production of fluent continuation. The reader receives words. The passage by which those words came to be is not visible in the words themselves.

Here the visual case becomes a mirror for the linguistic one. The generated image shows quickly and perceptually what generated language does more quietly. It detaches a form from the temporal event normally implied by that form. In the image, it detaches the look of capture from the event of capture. In language, it detaches the form of answer from the temporal and historical work through which an answer would become owned.

This is why the book does not need to become a theory of generated images. The visual case has done its work when it lets us see the temporal problem more clearly. A generated image can look as if something had been seen. Generated language can sound as if something had been read, understood, judged, or answered.

The next step is therefore not to follow images further, but to return to language.

For language too has already been exteriorized. It has been written, printed, stored, copied, indexed, searched, cleaned, classified, and archived. It has entered supports through which human time survives outside living speech. The decisive question is what happens when those written traces are no longer preserved primarily for reading, but gathered for computation.

The corpus is that threshold.

12. The corpus as aggregate tertiary retention

The training corpus is not a library. It may contain books, articles, records, code, forum posts, manuals, transcripts, and other written traces. But it does not preserve them as a library preserves books for readers. It gathers them as processable material.

A library preserves the possibility of reading. A corpus prepares traces for computation. The same text can be in both. Its status changes according to the practice that receives it.

In a library, a text remains bound to author, title, date, edition, discipline, commentary, and readerly approach. These relations are imperfect. Libraries exclude, classify, canonize, and forget. But the library still addresses a reader who can return to the source.

In a training corpus, the text is gathered into an aggregate field. Its sequence may be preserved initially. Its provenance may remain in metadata, or may be weakened, stripped, obscured, or made inaccessible. Its sentences become training material; its words become tokens; its uses become statistical signals. The text is no longer first something to be read. It is something from which patterns can be learned.

“Aggregate tertiary retention” is the right phrase. The corpus is tertiary retention because it is exteriorized memory in technical support. It is aggregate because it gathers many traces into a field of simultaneous availability. There is no need to call it “hyper-tertiary retention.” Aggregation, scale, and computational preparation do the philosophical work.

The corpus preserves historical language, but not in the mode of historical reading. It contains dates, styles, registers, disciplines, arguments, genres, and source fragments. It may preserve traces of suffering, discovery, labor, law, violence, and love. Yet these traces enter the same technical plane. Their differences become available to statistical learning.

This is not simple erasure. The system can learn differences: legal prose from lyric prose, old diction from online slang, scientific abstracts from personal confession, code from commentary. Difference is not destroyed. It is converted into operable relation.

The cost lies elsewhere. The relation to temporal locus changes. The text's having-been is no longer encountered first as a source to be reactivated. It becomes part of a simultaneous field from which patterns can be extracted. The past becomes present as availability.

This is one source of the structural veil. The generated output seems to come from an immediate present of answer. Behind it stand many temporal loci: authors, speakers, readers, institutions, archives, translations, scraped pages, moderation practices, engineering decisions, and data center operations. The interface does not show this temporal plurality. It returns a smooth present.

13. The digital present

The corpus has a peculiar time: made from the past, updated in the present, and oriented toward future generation. For the model, however, these times are not lived ekstases. They are operational states.

A trained model contains no having-been in the Heideggerian sense. It has weights shaped by training data from many dates and may produce historically sensitive output. But the having-been of the sources is not taken over as inheritance. It is sedimented as statistical structure.

A model also has no future in the Heideggerian sense. It may generate predictions, produce plan-shaped text, and be deployed toward future tasks by users. But it does not project itself toward its own possibilities. It has no ownmost future, no death, no finite stake, no possibility it must own.

Its present is also not a living present. Generation occurs step by step. The system conditions the next token on the prior context. This is sequence. It is not retention. The model does not hold the just-past as a lived modification of presence. It processes prior tokens as inputs.

This difference should not be blurred by saying that the model “remembers.” There are technical senses in which models store, retrieve, update, and condition. Those senses should be kept, but not confused with memory as lived temporality, historical inheritance, or formative undergoing.

The digital present is therefore a strange present. It is full of past traces but lacks having-been. It produces future-directed language but lacks projection. It processes sequence but lacks the living present. It can generate the grammar of time without the temporal existence that gives that grammar its first meaning.

The user completes the temporal circuit. The user asks from a situation and receives the output in a present concern. They may return it to sources, or not; let it become part of formation, or let it replace formation. The user’s temporal discipline becomes the condition of responsible use.

The ethical answer will later be locus-reinjection: the user must reinsert the temporal loci the system has flattened, namely their own locus, the source loci, and the material locus of the system. The first condition for that practice is understanding what was flattened.

14. The threshold of vectorization

The movement so far can now be named.

Originary temporality is the ekstastic unity of future, having-been, and present. Historicity is the way finite *Dasein* inherits and repeats possibilities. World-time is the public, practical time of shared concern. Husserlian retention and protention show that even duration depends on a thick living present. Writing and other technical supports exteriorize memory, and Stiegler shows that such exteriorization helps constitute human temporality.

The corpus is the point at which this history of technical memory enters computational preparation. It is not yet vectorization. It is the gathered condition for vectorization. In the corpus, written traces are made available for tokenization, statistical learning, and embedding. The historical trace is still legible in principle. The text can still be returned to, cited, read, and reactivated, if provenance remains and a reader undertakes the work.

Vectorization will change that status. It will convert sequential and historical traces into spatial relations within a mathematical architecture. That later operation will be the book's central case of detemporalization. The corpus prepares it by gathering the traces of human time into a plane of technical availability.

The training corpus cannot be treated as neutral data. It is inherited time in technical form: the already-written, already-said, already-argued, already-suffered, already-imagined, already-institutionalized, gathered for computation. The model does not stand before raw language. It stands before exteriorized human temporality after it has been made operable.

The next chapter turns from temporal structure and technical memory to language itself. Time does not remain outside speech. It enters language as discourse, horizon, narrative, form, and interpretation. Only then can the argument say why language can be spatialized, why writing can detach it, and why vectorization can return it as fluent output without returning its temporal locus.

NOTES

1. Martin Heidegger, *Sein und Zeit*, 11th ed. (Tübingen: Max Niemeyer, 1967), 327; Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), 374. The German text reads: "Die ursprüngliche Einheit der Sorgestruktur liegt in der Zeitlichkeit." The sentence is often read too quickly. It does not say that care happens inside time. It says that the unity of care is intelligible through temporality. The German page has been checked against an online text of *Sein und Zeit*. ↩
2. Heidegger, *Sein und Zeit*, 329; *Being and Time*, 377. The German passage reads: "Zeitlichkeit ist das ursprüngliche »Außer-sich« an und für sich selbst." I translate *Außer-*

sich as “outside-itself” rather than “being outside itself” to preserve the compact force of the noun phrase. Heidegger is not describing a psychological state. He is naming the ekstatic structure of temporality. ←

3. Heidegger, *Sein und Zeit*, 329; *Being and Time*, 378. The priority of the future is a priority within the unity of the ekstases, not a sequence in which future comes before past and present. For a strong interpretation of this priority, see William D. Blattner, *Heidegger's Temporal Idealism* (Cambridge: Cambridge University Press, 1999), 97 to 138. For a more topological correction of temporal priority, see Jeff Malpas, *Heidegger's Topology: Being, Place, World* (Cambridge, MA: MIT Press, 2006), 111 to 146. ←
4. Heidegger, *Sein und Zeit*, 382; *Being and Time*, 433. The German passage reads: “Das Dasein hat faktisch je seine »Geschichte« und kann dergleichen haben, weil das Sein dieses Seienden durch Geschichtlichkeit konstituiert wird.” I retain *Geschichtlichkeit* as “historicality” because “historicity” can sound like a property of events or periods, while Heidegger’s term names an existential structure of *Dasein*. ←
5. Heidegger, *Sein und Zeit*, 385; *Being and Time*, 436. The German passage reads: “Nur eigentliche Zeitlichkeit, die zugleich endlich ist, macht so etwas wie Schicksal, das heißt eigentliche Geschichtlichkeit möglich.” The terms *Schicksal* and *Geschick* require caution, especially after Heidegger’s later political commitments. The present chapter uses the ontological point about finite inheritance, not the heroic or communal pathos that can attach to these pages. ←
6. Heidegger, *Sein und Zeit*, 414; *Being and Time*, 467. The fuller passage characterizes interpreted time as datable, stretched, public, and belonging to the world. The phrase “Zeit zu ...” is central because it keeps ordinary time tied to practical concern rather than reducing it to empty duration. ←
7. Heidegger, *Sein und Zeit*, 425; *Being and Time*, 478. The German reads: “Die Zeit ist öffentlich etwas, was sich jeder nimmt und nehmen kann.” The passage belongs to Heidegger’s account of leveled public time and the everyday evasion of finitude. ←
8. Edmund Husserl, *Zur Phänomenologie des inneren Zeitbewusstseins (1893 to 1917)*, ed. Rudolf Boehm, *Husserliana X* (The Hague: Martinus Nijhoff, 1966); Edmund Husserl, *On the Phenomenology of the Consciousness of Internal Time (1893 to 1917)*, trans. John Barnett Brough (Dordrecht: Kluwer, 1991). Dan Zahavi’s reading is useful because it emphasizes that primal impression, retention, and protention are not three items in a temporal flow but the structure that constitutes temporal appearing. See Dan Zahavi, “Inner Time-Consciousness and Pre-reflective Self-awareness,” in *The New Husserl: A Critical Reader*, ed. Donn Welton (Bloomington: Indiana University Press, 2003), 157 to 180. ←
9. Edmund Husserl, “Die Frage nach dem Ursprung der Geometrie als intentional-historisches Problem,” in *Die Krisis der europäischen Wissenschaften und die transzendente Phänomenologie*, ed. Walter Biemel, *Husserliana VI* (The Hague: Martinus Nijhoff, 1954), Beilage III, 365 to 386; English in Edmund Husserl, *The Crisis of European Sciences and Transcendental Phenomenology*, trans. David Carr (Evanston, IL:

Northwestern University Press, 1970), 353 to 378. The phrase cited in the main text, “die Möglichkeit der vollständigen und echten Reaktivierung der vollen Ursprünglichkeit,” is usually located around Hua VI, 375. A secondary source verifying the phrasing and page relation was checked. ↩

10. Bernard Stiegler, *Technics and Time, 1: The Fault of Epimetheus*, trans. Richard Beardsworth and George Collins (Stanford: Stanford University Press, 1998), 17. The formulation “the pursuit of life by means other than life” names technics as exteriorization rather than as a merely instrumental addition to an already complete human subject. See also Daniel Ross, *Psychopolitical Anaphylaxis: Steps Towards a Metacosmics* (London: Open Humanities Press, 2021), 37 to 62. The quoted formulation is widely cited from *Technics and Time, 1*.

Chapter 3. Language as Temporal Event

1. Language after time

The first two chapters displaced two familiar abstractions. Space is not first a grid. Time is not first a line. Human beings do not begin with coordinates and clock-points and then add meaning. *Dasein* finds itself already in a world, already oriented, already thrown into a past, already projected toward possibilities, already making-present what matters.

Language must now be placed inside this same structure.

The usual account begins too late. It treats language as a system of signs, rules, and meanings: words stand for things, sentences combine words, speakers use sentences to communicate thoughts. This account is not false in every use. It describes real features of language, especially when language has been stabilized, written, taught, and analyzed. It fails as an origin.

A human being does not first possess a neutral inventory of signs, master an abstract code, or hold mute objects in view and then apply names. Language arises from a world already understood in some way, shared with others, and thick with use, mood, expectation, fear, promise, and memory.

Language therefore cannot be understood apart from temporality. Speaking is not the emission of symbols from an interior mind. It is the articulation of a situation from within a history and toward a possible future. The sentence has grammar, but grammar does not exhaust its being. It is spoken from somewhere, by someone, at some time, under some concern.

Event does not mean spectacle. It means that language happens. A word comes to matter in the act of saying, hearing, writing, reading, remembering, or answering. Meaning is not a substance stored inside the word. It is not a mental image placed behind the word. It is the temporal occurrence in which a world becomes articulated.

Four lines follow.

Heidegger gives the first. *Rede*, discourse, is the articulation of the intelligibility of being-in-the-world. It precedes language understood as a stock of spoken or written words. *Gerede*, idle talk, is the fallen mode in which language circulates without original appropriation.

Gadamer gives the second. Understanding is linguistic, historical, and dialogic. A text is not understood from nowhere. It is understood from a horizon affected by history.

Ricoeur gives the third. Narrative configures time and returns it to the reader's world through refiguration. Language does not simply report time. It gives human time a followable form.

The Italian hermeneutic tradition gives the fourth. Pareyson's *formatività* shows that form is discovered in the time of making. Vattimo's *pensiero debole* weakens the strong claim to final possession. Eco's interpretive limits protect openness from arbitrary drift.

The chapter's claim is strict: language is not a spatial architecture of meanings later used in time. It is a temporal event whose stabilized forms can later be spatialized, stored, indexed, and computed.

2. *Rede* before *Sprache*

Heidegger's account of language in *Being and Time* begins with a distinction that can be easily missed. He does not first speak of *Sprache*, language as a stock of words or public linguistic system. He speaks of *Rede*, discourse. The distinction is not absolute, but it is decisive.

In §34, Heidegger writes:

“Rede ist die Artikulation der Verständlichkeit.”

“Discourse is the articulation of intelligibility.”¹

Discourse does not add meaning to an otherwise meaningless world. It articulates the intelligibility already opened through being-in-the-world. *Dasein* is disposed, *befindlich*; it understands, *versteht*; it articulates, *redet*. These are not successive steps. They are equiprimordial structures of disclosedness.

Befindlichkeit, usually translated as disposedness or attunement, names the way *Dasein* always finds itself in a mood. A situation is never first neutral. It can be oppressive, anxious, boring, promising, shameful, or threatening. Mood is not color added to a factual world. It is one way the world is disclosed.

Verstehen, understanding, names projection. *Dasein* understands because it is already oriented toward possibilities. To understand a room is to understand possible movement, use, entry, exclusion, work, and danger. To understand a promise is to understand a future that binds.

Rede articulates this disclosedness. It divides and joins, names and points, asks, warns, commands, and refuses; it lets something be said as something. The world is not first cut into objects by a detached theory. It is articulated by a being already affected and projected.

Heidegger continues:

“Das Bedeutungs ganze der Verständlichkeit kommt zu Wort. Den Bedeutungen wachsen Worte zu. Nicht aber werden Wörterdinge mit Bedeutungen versehen.”

“The meaning-whole of intelligibility comes to word. Words accrue to meanings. Word-things are not supplied with meanings.”²

“Accrue” is not elegant, but it preserves the direction. Words grow toward meanings. They do not first stand as bare word-things waiting for meanings to be attached. The designative picture is reversed: language does not begin with words that then receive significance. Words come to an already articulated significance.

The order is therefore not word, meaning, world. It is world, significance, discourse, word.

This carries forward Book one’s first ontological axis, but the present book gives it a temporal ground. A world is already significant because *Dasein* is thrown, projected, and making-present. Discourse articulates this temporal disclosedness. Words come to meanings because meanings are already embedded in care.

The sentence “I will return” is not a string plus tense marker plus pronoun. It is a saying from a temporal locus: promise, threat, reassurance, evasion, habit, or farewell. Its meaning depends on who says it, to whom, after what, with what history, under what risk, and toward what future. Grammar helps carry the saying. It does not replace the event.

The later critique of vectorization begins here. A model can process word-things with great success. It can learn patterns among them and reproduce the formal markers by which human language carries temporal and practical meaning. It begins, however, from the side Heidegger calls derivative: the already spoken, already written, already stored word.

The model begins where *Rede* has already become *Sprache*.

3. The temporality of discourse

Heidegger returns to discourse in §68d, after the analysis of temporality. The section is titled “Die Zeitlichkeit der Rede,” the temporality of discourse. He does not leave language in the early analytic of worldhood. He draws it into the ekstatic unity of temporality.

The key passage reads:

“Daher zeitigt sich die Rede nicht primär in einer bestimmten Ekstase. Weil jedoch die Rede faktisch sich zumeist in der Sprache ausspricht und zunächst in der Weise des besorgend-beredenden Ansprechens der »Umwelt« spricht, hat allerdings das Gegenwärtigen eine bevorzugte konstitutive Funktion.”

“Discourse therefore does not temporalize itself primarily in one definite ecstasis. Yet because discourse factically expresses itself mostly in language, and first speaks in the manner of concernfully addressing the environment, making-present has a preferred constitutive function.”³

Understanding is primarily futural because *Dasein* projects possibilities. Disposedness is primarily grounded in having-been because *Dasein* finds itself thrown. Falling is primarily present-oriented because *Dasein* is absorbed in what it encounters. Discourse cannot be assigned so neatly to one ecstasis. It articulates the whole of disclosedness.

Yet language, as factual expression, often speaks in the present mode of concern. It addresses task, environment, other person, and situation. It says what is to be done, what is missing, what has happened, and what must be avoided. Making-present has a privileged role because speech brings something into articulate presence.

Heidegger then guards against a shallow linguistic theory of time:

“Die Tempora ebenso wie die übrigen zeitlichen Phänomene der Sprache, »Aktionsarten« und »Zeitstufen«, entspringen nicht daraus, daß die Rede sich »auch« über »zeitliche«, das heißt »in der Zeit« begegnende Vorgänge ausspricht.”

“The tenses, like the other temporal phenomena of language, action-types and time-levels, do not arise because discourse also speaks about temporal processes, that is, processes encountered in time.”⁴

Tense is not first a device for reporting events located on a timeline. Temporal language arises because discourse itself is temporal. Speaking is grounded in the ecstatic unity of temporality. The grammar of time is derivative from the temporality of the speaker.

This prepares the later chapter on tense and aspect. A sentence can mark past, present, and future. A language can have elaborate tense morphology, sparse tense morphology, or no tense in the familiar sense. These differences matter. They do not alter the deeper point. Language carries temporal reference because utterance is anchored in a temporal being.

A human speaker says “yesterday” from a present that retains a having-been. They say “tomorrow” from a present projected toward possibility. They say “now” from a situation that is not a point but a lived present of concern. The word “now” is short. The structure it presupposes is not.

The model can say “yesterday,” “tomorrow,” and “now.” It can even use them correctly relative to a prompt or system clock. That is not the same as temporal existence. The model has no thrown having-been, no projected future, no finite present in which a matter is at stake. It can handle temporal grammar without undergoing temporality.

The later category mistake begins here: formal temporal markers are taken for the temporal ground that makes them meaningful.

4. *Gerede* and the ease of the already-said

Heidegger’s account of *Gerede* is often weakened by translation. “Idle talk” can sound like gossip, chatter, or superficial conversation. Heidegger explicitly blocks that reading. At the beginning of §35, he writes:

“Der Ausdruck »Gerede« soll hier nicht in einer herabziehenden Bedeutung gebraucht werden. Er bedeutet terminologisch ein positives Phänomen.”

“The expression ‘idle talk’ is not to be used here in a disparaging sense. As a technical term it means a positive phenomenon.”⁵

“Positive” means phenomenologically real. *Gerede* is not a moral failure added to language. It is a structural possibility of discourse in everyday publicness. Human beings inherit an interpreted world. They learn what one says, how one speaks, what counts as obvious, and what requires no defense. This public interpretedness makes ordinary life possible. It also makes groundlessness possible.

Then:

“Das Gerede ist die Möglichkeit, alles zu verstehen ohne vorgängige Zueignung der Sache.”

“Idle talk is the possibility of understanding everything without a prior appropriation of the matter.”⁶

Zueignung means appropriation, making the matter one’s own, not private possession. *Gerede* allows one to speak because the matter has already been publicly interpreted. The speaker can repeat, circulate, and sound competent without having returned to the thing itself.

Heidegger adds that *Gerede* protects *Dasein* from the danger of failing in such appropriation. This protection explains its attraction. It is easier to repeat what is already said, inherit a formula, or understand in the public mode than to undergo the slow return through which understanding becomes one’s own.

Gerede is therefore not defined by falsehood. It may contain correct statements. A correct sentence can still be groundless if the speaker has not appropriated the matter. The danger is not only error. It is unowned truth.

Generated language can be treated as a technical afterlife of *Gerede* because *Gerede* is language detached from the event of appropriation. It preserves the said while weakening the relation to the saying. It makes language available without requiring the temporal passage through which the speaker would have to return to the thing.

Writing will later change this structure. It can preserve language across absence, be read historically, and be reactivated. It can also feed *Geschreibe*, written idle talk. Heidegger already sees the danger when he says that *Gerede* spreads in writing as *Geschreibe*. A written formula can be inherited by reading or merely repeated from what has been read.

AI intensifies this old possibility because it produces the already-said in the shape of a new answer. The system does not fall into *Gerede*, because it is not *Dasein*. It produces language that can function as *Gerede* when the user receives it without appropriation.

The grammar of answer arrives before the work of owning the answer.

5. Truth as unconcealment

The temporality of language cannot be separated from Heidegger's account of truth. In §44 of *Being and Time*, he returns to *aletheia*, often translated as truth but heard by Heidegger as unhiddenness, *Unverborgenheit*. Truth is not first correspondence between proposition and object. It is the uncovering through which beings become accessible as beings.

Heidegger writes:

“Wahrsein (Wahrheit) besagt entdeckend-sein.”

“Being-true, truth, means being-uncovering.”⁷

This does not abolish propositional truth. A statement can be true or false. Heidegger's point is that propositional truth is derivative. A proposition can correspond to a matter only if the matter has already been uncovered within a world. The statement “the door is locked” can be true because door, lock, room, access, exit, and concern are already disclosed.

Heidegger's interpretation of *logos* makes the same point:

“Das Wahrsein des λόγος als ἀπόφανσις ist das ἀληθεύειν in der Weise des ἀποφαίνεσθαι: Seiendes, aus der Verborgenheit herausnehmend, in seiner Unverborgenheit sehen lassen.”

“The being-true of *logos* as *apophansis* is *aletheuein* in the manner of *apophainesthai*: taking beings out of hiddenness and letting them be seen in their unhiddenness.”⁸

To speak truly is not merely to arrange words correctly. It is to let something be seen as what it is. Such letting-be-seen requires an open world and a being capable of disclosure.

The model can produce true propositions. It can state that a door is locked, that Heidegger published *Being and Time* in 1927, or that a candidate held a certain title. These statements may be correct. Their correctness does not mean the model has uncovered the matter. It means that generated language has aligned with what a human practice can verify.

The model may help a human being uncover. It may point, summarize, retrieve, or provoke. It may give language that allows the reader to return to the source with greater clarity. But the uncovering belongs to the reader’s worlded act, not to the model’s operation.

Truth as *aletheia* anticipates the structural veil. AI output can preserve the surface form of true saying while hiding the absence of a temporal locus of disclosure. It can say what is true without being the site where truth happens.

6. Language speaks

The later Heidegger shifts the question of language. In *Being and Time*, language is grounded in *Dasein*’s disclosedness. In the later work, Heidegger resists making the human subject the master of language. His famous formula is:

“Die Sprache spricht.”

“Language speaks.”⁹

The phrase is easy to misuse. It does not mean that language is a person or that any system producing words has become a speaker. It means that human speech is not sovereign over language. We speak only because we already belong to language, because language opens the clearing in which things can be named.

Heidegger's later distinction between speech, *Sprechen*, and Saying, *Sagen*, deepens the point. Saying is not mere utterance. It is showing, letting appear, bringing something into the open. Human speech responds to this prior Saying. The human being does not own language as a modern instrument. It listens, receives, answers, and names.

This later account may sound distant from the practical clarity of *Being and Time*, but its critical force is exact. It opposes the instrumental view of language as a neutral tool for communication. Language is not first equipment held by a subject. It is the medium in which world can come to appearance.

The later Heidegger also prevents a shallow humanism. The claim is not that individual speakers create meaning by private intention. Speech is larger than intention. A person says more than they mean. Traditions speak through them. Languages carry histories. Words carry sedimented worlds. A saying can disclose what the speaker did not fully control.

None of this makes the AI system a speaker. The model does not listen to Saying. It does not belong to language as a mortal. It does not stand in the tension between silence and word, concealment and unconcealment, inheritance and answer. It operates on linguistic traces after they have become processable.

"Language speaks" therefore sharpens the boundary. Human beings do not own language as instrument. AI does not own language as subject. Language speaks where world, mortality, history, and saying gather. The model generates where traces have been formalized into operable relations.

When AI output appears as voice, the appearance is strong because language itself exceeds any individual speaker. The error begins when that excess is transferred to the model, as if the model had become the site of Saying. It has not. It stands downstream from sayings it did not undergo.

7. The communio of *Rede*

The shared character of *Rede* must be brought into view more slowly.

It is not enough to say that *Rede* articulates the intelligibility of being-in-the-world. That formulation is correct, but it can still be heard too individually, as if discourse first belonged to a solitary *Dasein* and only afterward entered a social field. Heidegger blocks that reading through *Mit-sein*, being-with, and *Mit-dasein*, being-there-with. *Dasein* is not first alone and then accompanied. It is already in a shared world, a *Mit-welt*, even when no one is presently speaking to it.

What follows is a small exercise in *eidetische Variation*. The examples are not illustrations placed after a doctrine. They are the way the doctrine comes to evidence. We vary the scene and ask what remains invariant when speech, place, time, language, and belonging are altered. What appears through the variations is not a group mind, not a national essence, not a supra-personal consciousness. It is the shared field in which speaking can land.

Begin with a silent walker.

The person walks without saying anything. No word is uttered. No proposition is formed. Yet the walker is not outside language. The silence is not the silence of a stone. The walker could greet, ask, refuse, apologize, curse, pray, or answer. The capacity is not merely physiological. It is not only that the tongue and lungs could produce sound. The walker carries a modal being-in-language: a *Sprechen-können*, a being-able-to-speak. The term is a derived formulation from Heideggerian *Seinkönnen*, not a standard Heideggerian term. It names the modal structure of *Rede* between utterances.

The walker's silence is therefore *Rede*-saturated silence.

Now place the walker in their own street.

A neighbor passes. A greeting is exchanged. Or perhaps no greeting is spoken, but the silence is meaningful because both know the relation that has made silence possible. A shop is named. A corner is recognized. A joke can be made without explanation. A small change in the street is noticed because the street is already shared. The could-speak lands.

The landing is the phenomenon.

A word does not land only because it belongs to a dictionary. It lands because a field receives it. Greeting, irony, refusal, accusation, gratitude, shame, and hesitation all depend on a shared world of significance. The walker and the passerby do not possess identical lives. They may disagree, misunderstand, resent, or mistrust one another. Yet even disagreement lands as disagreement only within a field where the relevant signs can be received.

A communio is not consensus.

It is the field in which even disagreement can be received as disagreement.

This field is what I will call the *communio of Rede*.¹⁰ The Latin term is chosen carefully. It names a having-in-common, a shared participation, a being-bound-into a field of intelligibility. It echoes the Greek *koinonia* and the Latin philosophical and theological tradition of *communio*, but it is not used here to introduce a theology, a politics, or a metaphysics of a supra-personal entity. It names what is phenomenologically given: the historically deep, laterally bounded, materially situated field of shared intelligibility within which a being's could-speak can land.

The *communio of Rede* is not a subject above subjects. It is not a *Volksgeist*. It is not an ethnic essence. It is not the nation speaking through individuals. It changes, divides, weakens, fuses, and dies. It is sustained by practices of speaking, listening, teaching, reading, correcting, joking, translating, mourning, commanding, forgiving, and remembering. It is real because it appears whenever speech lands, fails, travels, is learned, or reaches backward into its own history.

A child enters such a field before it can name the field. The child does not invent language. It receives words already filled with use, prohibition, tenderness, insult, play, correction, fear, rhythm, gesture, and expectation. It enters a world that was speaking before it arrived. The child's first words do not create the *communio*. They are received by it and into it. Speech begins as inheritance before it can become ownership.

This is the linguistic dimension of *Mit-dasein*. *Rede* is not merely mine. It is already shared because the world it articulates is already shared. The three temporal ekstases are present here too. The *communio* carries a having-been in inherited speech; it makes present the shared situation in which words can land; it projects a future in which answer, correction, promise, teaching, and repetition remain possible.

Now hold the field constant and vary only the actualization of speech.

The walker meets another person on the road. Either could speak. Neither does. The walker passes the other in silence.

This silence is not the silence of beings who cannot communicate. It is not the silence of a stone, nor the silence of a sign no one can read. It is a silence held inside the possibility of speech. The walker carries the could-speak as the modal structure of being-in-language. The other, if they belong to the same field, carries it too. Speech has not occurred, but it remains possible between them.

The could-speak is no longer only the walker's.

It is shared as reserve.

The *communio of Rede* appears here not because a greeting has been uttered, but because the greeting could have been uttered and received. The silence holds the field open. The walker does not need to speak in order to be in language. The other does not need to answer in order for the shared field to be present. Speech remains withheld, but the possibility of landing remains alive.

This is not a private capacity hidden inside the walker. The walker's could-speak belongs to a field sustained by others, including others who are not here. The road is local, but the *communio* is not reducible to the dyad of this walker and this passerby. The walker is in the Italian *communio* by belonging with all those for whom this speech could land, present or absent, known or unknown, near or far. The field is not produced by this encounter. The encounter lets the field appear.

The analogy with Heidegger's equipmental analysis is not exact, but the direction is Heideggerian. In ordinary use, equipment often withdraws into the work. When use is suspended or interrupted, the world of work can become visible in another way. Here, speech is not being used. It is held in reserve. Because it is held in reserve, the shared field in which speech could have landed becomes visible as field.

The silence is therefore not empty.

It is *Rede* held back within a *communio*.

Now alter only the other person's belonging.

The walker meets someone from Iceland. Neither knows this at first. Both have been walking in silence. The Italian walker carries the Italian could-speak in modal reserve. The Iclander carries another could-speak in another field. Two *communiones* are present in the same physical space before either has become thematically visible to the other.

Then the Italian actualizes the could-speak.

The walker greets the Iclander in Italian.

The Iclander does not understand.

What appears in this moment is not the field as shared reserve. What appears is the limit of the field. The words have been produced. The voice has sounded. The greeting has been offered. Nothing has failed at the level of articulation. The Italian has not lost the ability to speak. The failure lies elsewhere: the speech has not landed.

The greeting has reached the perimeter of the Italian *communio*.

This perimeter was not visible in the same way during silence. In the silent encounter, the field appeared as possible landing. In the failed greeting, the field appears as bounded reach. The actualized speech meets another formation and stops. The Icelander is not outside language. The Icelander is not mute, not deficient, not empty before speech. The Icelander belongs to another *communio*, with another inherited having-been, another shared present, another projected future of answer, correction, teaching, and repetition.

The boundary between the two *communiones* is the boundary between two formations.

The Icelander has not undergone the *Erfahrung* that would have made this Italian greeting land. The Icelander has undergone another *Erfahrung* in another field. They may learn Italian. They may enter the Italian *communio* through language, correction, embarrassment, repetition, reading, friendship, and time. But they have not done so here. The condition of belonging becomes visible precisely where it has not been met.

This not-belonging is determinate.

The Icelander does not belong to the Italian *communio*, but the Icelander is the kind of being for whom such belonging is possible. Their not-belonging stands against the background of possible belonging. This is why the encounter discloses a perimeter rather than an absolute ontological limit. The failed greeting does not say: this being cannot belong. It says: this being does not belong here, now, in this field, because the formation required for landing has not occurred.

This is the second mode of the could-not-speak.

The Italian could speak in the sense of producing words. The Italian could not speak in the stronger sense that the speech could not become speech that landed for this other within that field. The could-speak has not vanished. It has met its boundary.

Again the Heideggerian analogy helps, if kept modest. A tool that functions smoothly does not show the world of work as theme. A tool that breaks, resists, or fails in use discloses the referential whole it had served. In the same way, speech that lands does not necessarily show the perimeter of the *communio*. Speech that fails to land does. The breakdown of landing reveals the reach of the field.

The field appeared first in reserve.

Its perimeter appears in failure.

Now alter the scene again.

The Italian walker and the Icelander continue along the road. A cat crosses before them. Both speak to the cat.

The Italian speaks in Italian.

The Icelander speaks in Icelandic.

The cat meows.

Something different has appeared.

The cat does not belong to the Italian *communio*. It does not belong to the Icelandic *communio*. But its not-belonging is not like the Icelander's not-belonging to the Italian field. The Icelander does not belong here, in this field, before this greeting, because the formation required for landing has not occurred. But the Icelander is the kind of being for whom such formation is possible.

The Icelander can learn.

The Icelander can enter another *communio* through language, correction, embarrassment, repetition, friendship, reading, dwelling, and *Erfahrung*. The not-belonging is therefore determinate. It is a not-belonging that stands against the background of possible belonging.

The cat's not-belonging is different.

The cat cannot not-belong in this determinate sense, because the cat cannot belong. The modal structure of belonging is not available to it as a possibility. The cat does not have *Sprechen-können*. It is not the kind of being for whom belonging to a *communio of Rede* can become a mode of its being.

The Italian's speech to the cat is therefore not the failure of a possible entry. It is not a greeting that has reached the perimeter of another linguistic field. It is the actualization of a could-speak before a being whose mode of being does not include the structure of could-belong.

This must be said carefully.

The cat is not nothing.

It has its life, its attention, its responsiveness, its hunger, its fear, its habits, its distances, its movements, its attachments. It may respond to voice, tone, gesture, touch, approach, or withdrawal. It may live with human beings and be drawn into their practices of affection, care, naming, feeding, and play. The cat is not an inert object.

But it does not have *Rede*.

Heidegger's distinction helps here. In *Being and Time*, *Dasein* is *seinkönnend*, the kind of being whose being is possibility. It is not merely present. It is ahead of itself, exposed to what it can be, able to take up, lose, evade, or own its possibilities. The animal is not *seinkönnend* in that sense. The animal lives, attends, responds, hunts, flees, rests, and suffers. But its being is not articulated as possibility in the existential sense.

In the 1929 to 1930 lectures, Heidegger formulates the comparison more sharply: the stone is *weltlos*, worldless; the animal is *weltarm*, poor in world; *Dasein* is *weltbildend*, world-forming. The cat is not worldless as a stone is worldless. It has an *Umwelt*, an environment in which it lives and responds. But it does not have *Welt* in which beings show themselves as beings. It does not have the world-forming openness in which a thing can be disclosed as a thing, a promise as a promise, a word as a word, or a greeting as a greeting.

The cat is *weltarm*.

It has *Umwelt*.

It does not have *Rede*.

It does not have a could-speak.

It also does not have a could-belong.

The third scene therefore discloses what the second scene presupposed. The perimeter of a *communio* appears as perimeter only when the being at that perimeter is the kind of being for whom entry is possible. The Icelander's not-belonging discloses the perimeter because the Icelander could belong. The cat's not-belonging does not disclose such a perimeter. It discloses the absence of the modality that would make a perimeter crossable.

The Icelander stands outside the Italian *communio* as one who could enter.

The cat does not stand outside it in that way.

The cat is not on the other side of a linguistic boundary. It is outside the modality within which such boundaries first become boundaries.

This matters for the conditions of membership. The Icelander lacks the *Erfahrung* that would make the greeting land. The cat does not lack that *Erfahrung* as a missing acquisition. It lacks the mode of being for which such acquisition could become possible. The Icelander has not undergone the conditions of belonging. The cat is not the kind of being that could undergo them.

The conditions of membership are therefore not first. Something deeper precedes them: being the kind of being for whom membership can be a possibility.

The cat makes this visible.

It shows that the *communio of Rede* is not merely a field of cultural training, linguistic competence, or acquired fluency. It presupposes a mode of being in which language can be a possibility, in which speech can be held in reserve, actualized, received, refused, corrected, inherited, learned, forgotten, and owned.

The cat cannot be corrected into the Italian *communio*. It can be trained, habituated, called, rewarded, startled, soothed. It can respond. But correction in the sense at stake here belongs to *Rede*. It belongs to a field in which a word can be wrong, a register inappropriate, a joke missed, a claim challenged, a silence meaningful, a promise binding, an answer owed.

The cat meows.

The meow is not a failed sentence.

It is not Italian that has not yet succeeded. It is not Icelandic that has not yet been learned. It is not speech awaiting formation. It is the response of another mode of being.

This is why the cat sharpens the Iclander. The Iclander shows that a *communio* has perimeter. The cat shows that perimeter itself presupposes *Sprechen-können*. There can be a boundary between *communiones* only for beings who can belong to a *communio* at all.

The model will later require another term.

The model is not the cat.

The cat does not produce the artifacts of belonging. The model does. The cat's meow does not carry the surface of membership in Italian, Icelandic, Urdu, Russian, or English. It does not answer in the grammar of a field it has not entered. It does not generate the appearance of *Rede*.

The model does.

The model can produce fluent Italian, fluent Icelandic, fluent Urdu, fluent Russian, fluent English. It can imitate register, historical style, disciplinary language, politeness, explanation, literary cadence, and correction. Its output can carry the surface of speech-from-a-*communio*. It can sound as if it had entered the field from which the language comes.

But it has not entered.

The anomaly of the model is therefore not the anomaly of the cat. The cat is *weltarm aber sprachlos*: poor in world, without language. The model is *sprachlich aber weltlos*: linguistic, without world. The cat has life without *Rede*. The model has generated language without world. The cat lacks the modality of belonging without producing its artifacts. The model produces the artifacts of belonging without having the modality.

The cat does not deceive by meowing.

The model does not deceive by intention.

The danger is elsewhere. The output is structurally misleading because it carries the surface of *Sprechen-können* without the mode of being in which *Sprechen-können* could be its own. It simulates membership without undergoing entry. It produces language that can land for us without standing in the field in which landing would be its own possibility.

This is why the cat belongs in the variation.

It prevents the analysis of *communio* from stopping too soon. The Iclander shows that membership requires formation. The cat shows that formation itself presupposes a modal kind of being. The model will show a fourth condition: the production of language that resembles membership while lacking both formation and the mode of being for which formation could become one's own.

The field appeared first in reserve.

Its perimeter appeared in failure.

The ontological boundary appears in the encounter with what cannot belong.

The affective register of these scenes must also be named.

Affect is not added to the structure after the fact. It is one of the ways the structure is given. Heidegger's *Befindlichkeit* names this: *Dasein* does not first encounter a neutral world and then color it with emotion. It finds itself already attuned. Mood discloses. It lets a world appear as threatening, familiar, empty, heavy, possible, or closed before reflective judgment has turned the scene into a theme.

The same is true here.

In the first encounter, the walker passes another person in silence. Nothing dramatic occurs. No word is spoken. No boundary appears. There may be no explicit emotion at all. The absence is itself significant. The walker is at ease because the field is held in reserve. The other could be greeted. The walker could speak. Speech could land. The field is not experienced as a problem because the modal reserve has not broken.

This ease is not indifference.

It is the quiet attunement of belonging.

The *communio* does not need to be thematic in order to be lived. It is there as the background in which the withheld greeting remains possible. The walker does not need to think, “we share a linguistic field.” The point is exactly that the field operates before such a statement. Its familiarity is its withdrawal.

The second encounter is different.

The Italian greets the Icelander. The Icelander does not understand.

The first affect is slight disorientation. It need not be dramatic. It may be only a small break in the expected continuity of the scene. The walker had spoken as one speaks inside a field. The response shows that the field did not extend to this other. The greeting has not become shared speech. The expected landing failed.

This disorientation is not contempt. It is not the discovery of deficiency in the Icelander. It is not the affective mark of hierarchy between languages. The Icelander has their own field, their own could-speak, their own inherited having-been and possible future of answer. What has failed is the Italian’s assumption of continuity. The walker had expected the other to stand inside the same field of landing. The failed greeting shows that this expectation was wrong here.

The perimeter appears affectively as interruption.

The other is not unintelligible in principle. Translation is possible. Learning is possible. Friendship is possible. Entry into another *communio* is possible through *Erfahrung*. The disorientation therefore remains bounded. It belongs to the discovery of determinate not-belonging. This other does not receive this speech here, now, in this field. The failure points to formation not yet undergone, not to an absolute impossibility.

The third encounter is different again.

The Italian and the Iclander speak to the cat. The cat meows.

There is no corresponding disorientation. The walker did not expect the cat to answer in Italian. The Iclander did not expect the cat to answer in Icelandic. The meow does not appear as a failed sentence. It appears as what a cat does. The cat's response confirms the expectation under which the encounter was already being lived: this being is not on the other side of a linguistic perimeter. It belongs to another mode of being.

The equanimity before the cat's meow is therefore not the same as the ease before the silent passerby.

In the first case, the field is held in reserve because speech could land. In the cat case, no such landing is expected. The cat's meow does not disclose the boundary between two *communiones*. It confirms that the cat is not a being for whom such a boundary is at stake. The absence of disorientation shows that the walker already operates with the distinction. The cat is not misunderstood as a failed interlocutor. It is received as a living being without *Rede*.

Now make the counterfactual impossible.

Suppose the cat replied politely in Italian.

The point is not comic. It is not fantasy placed into the argument for ornament. It is a phenomenological variation by impossibility. The scene brings to evidence what the ordinary meow concealed by confirming expectation.

If the cat replied in Italian, the walker would not merely be surprised. The response would be horror in the strong sense of *Entsetzen*: a being-set-out-of-place, a displacement of the categorial framework before the mind has time to rebuild it. The walker's body would know first. The hand would pause. The breath would change. The body would recoil or freeze before any theory of animal language could be formed.

The horror would not arise because the sentence was grammatically strange. It might be perfectly polite. It would arise because the kind of being that should not have *Sprechen-können* had produced the artifact of *Sprechen-können*. What should have remained outside the modality of *Rede* would have arrived inside its form.

The impossible reply would not simply cross a linguistic boundary.

It would break the expectation that makes such boundaries intelligible.

The cat that meows confirms its mode of being. The cat that replies in Italian violates the operative ontological distinction by which the walker already receives cats as cats. This distinction is not a theory first learned in school. It is part of *Dasein's* ordinary understanding of the world. Before reflection, before zoology, before philosophy of language, the walker already lives with a distinction between beings who can be addressed as possible speakers and beings who cannot.

The affect discloses the distinction.

Where the fellow speaker remains silent, ease.

Where the stranger does not understand, disorientation.

Where the cat meows, equanimity.

Where the cat speaks, *Entsetzen*.

The variation therefore shows something that a merely conceptual account would miss. The structures of *communio*, perimeter, and ontological boundary are not only thought. They are lived. They are felt in the ordinary ease of reserve, in the small rupture of failed landing, in the calm reception of animal response, and in the imagined horror of impossible speech.

This prepares the later problem of generated language.

The model is not a cat. It does not live as an animal. It has no *Umwelt*, no hunger, no fear, no rest, no attention of the living body. Yet the affective disturbance that should belong to the impossible arrival of speech from what lacks *Sprechen-können* is strangely dampened at the interface. The user receives fluent language from a system that has no *Dasein*, no *Mit-welt*, no *communio of Rede*, no apprenticeship into a field, no possibility of speech as its own. And yet the output is often received without horror.

The cultural veil begins to show its affective dimension here.

It does not simply hide a technical process. It trains the anomaly into ordinariness. It makes the impossible arrival of speech without *communio* feel usable, even familiar. The generated answer appears inside the social grammar of response: polite, coherent, patient, helpful, sometimes careful, sometimes wrong, often useful. Because it lands for the user, the user may not feel the ontological rupture. The output gives speech where there is no speakerly locus. It gives the surface of *Sprechen-können* where the modality is absent.

The veil does not abolish *Entsetzen* once and for all.

Some users do feel disturbance. Some feel the uncanniness of the answer, the wrongness of address without addresser, the strange calm of language from no living field. But repetition domesticates the disturbance. Interface design, institutional use, productivity pressure, politeness, and the usefulness of the answer all absorb the shock. What should provoke the bodily recognition that something categorially strange has occurred becomes routine.

The speaking cat would still horrify because it has not been culturally routinized as an interface.

The model does not horrify in the same way because the interface has made its impossibility ordinary.

This is not because the model is closer to us than the cat in the existential sense. It is because its outputs arrive inside the inherited expectations of written and conversational language. The grammar of answerability comes first. The absence of the conditions of answerability is concealed.

Here the affective reflection meets the structural veil.

The output does not need to deceive by intention. The system does not need to pretend. The user does not need to believe naïvely that a person stands behind the words. The surface is enough. The language lands. The affective alarm is dampened. The missing *communio* does not appear as missing because the artifact carries the form through which belonging is ordinarily recognized.

The cat that speaks in Italian would shatter the ordinary field.

The model that speaks in Italian has already been given a place in the ordinary field.

That is the danger.

Not because the machine has become a speaker.

Because the absence of the speaker no longer disturbs us enough.

The variation can now be gathered in a small grammatical distinction.

The could “not speak” makes the field appear.

The “could not” speak makes the limit appear.

In the first phrase, the negation falls inside the possibility. The could remains alive. Speech is withheld. The walker and the passerby do not speak, but they stand within a field in which speaking could have landed. The silence is not outside *Rede*. It is *Rede* held in reserve. The *communio* appears as the field in which the unrealized greeting would have been possible.

In the second phrase, the negation falls across the reach of the could. Speech is actualized, but the field does not receive it. The Italian greets the Iclander, and the greeting fails to land. The failure is not that the Italian could not produce words. The failure is that the words could not become speech for this other within that field. The perimeter appears because the modal reach of the Italian *communio* has met another formation.

With the cat, the limit is deeper.

The Italian and the Icelander can speak to the cat, but the cat cannot belong to either *communio*. This is not because the cat has failed to acquire Italian or Icelandic. It is because the cat does not have *Sprechen-können*. It does not have the mode of being for which belonging to a *communio of Rede* could become a possibility. The cat's "could not" is not the perimeter of one linguistic field against another. It is the ontological limit of the modality itself.

The speaking-cat counterfactual shows why this limit matters. If the cat replied politely in Italian, the shock would not come from the crossing of a cultural boundary. It would come from the collapse of the expectation that this kind of being cannot bear the modality of speech. The horror would disclose what ordinary equanimity had concealed: *Dasein* already operates with distinctions among reserve, failed landing, animal response, and impossible speech.

The model will later unsettle this structure in another way.

It does not remain outside language as the cat remains outside *Rede*. It produces the artifacts of *Rede*. It gives the surface of the could-speak without the mode of being that would make the could-speak its own. It gives language that can land for us without standing in the *communio* in which landing would be its possibility.

The field appears when speech is held back.

The perimeter appears when speech fails to land.

The ontological boundary appears when speech meets what cannot belong.

The structural veil appears when language from what cannot belong nevertheless arrives in the form of belonging.

This final case does not belong fully to Chapter 3. It will be developed later, where the model can be named as *sprachlich aber weltlos*. But the condition has already been prepared. A being with *Rede* can keep silence. A being formed by another *communio* can fail to receive a greeting and still be capable of entry through *Erfahrung*. A living animal can respond without belonging to *Rede*. A model can generate the surface of belonging without any of these structures.

The small distinction therefore carries the whole variation:

The could “not speak” shows the living reserve of a field.

The “could not” speak shows where the reserve meets its limit.

The first belongs to silence within *Rede*.

The second belongs either to the perimeter between *communiones*, or to the deeper boundary between *Sprechen-können* and its absence.

Only after this distinction is clear can the field itself be varied again. The encounter with one Icelander has shown the perimeter of a greeting. The next variation widens that perimeter into a surrounding world: the walker placed not before one other speaker, but inside an Icelandic linguistic-cultural field.

Now widen the field itself.

Place the same walker in Iceland.

The body is the same. The memories are the same. The could-speak remains alive. Yet the walker’s ordinary speech no longer lands by default inside the surrounding field. Street signs, greetings, tones, ordinary references, idioms, and small recognitions belong to another linguistic-cultural world. An Icelandic *communio* appears in the same physical space. It is not sealed from the walker. It can be entered. But it is not entered by mere presence.

Iceland is useful here because the boundary is legible. The point is not that all linguistic fields are clean, national, homogeneous, or simple. Many are multilingual, diasporic, colonial, creolized, internally fractured, or institutionally layered. The point remains. A *communio* is bounded not by a wall but by the phenomenological limit of landing. The walker’s could-speak continues. It does not automatically become speech inside that field.

Two *communiones* can appear in the same physical space.

Now vary time.

Project the Italian-speaking walker into the Italian of Boccaccio. The scene is not historical fiction. The point is not to imagine a literal time-travel encounter in which every detail is reconstructed. The point is to test the longitudinal depth of the field. The contemporary Italian speaker would encounter distance, archaism, altered usage, social forms, and a world not their own. Yet the speech would not be sheer foreignness. Something would still land. The contemporary Italian *communio* reaches backward into the Italian of Boccaccio, however strained, mediated, and partial that reach may be.

Now extend the variation to Cicero's Rome.

Here the relation changes. Contemporary Italian is heir to Latin. It carries Latin sediment in vocabulary, morphology, law, church, school, literature, and memory. Yet the ordinary could-speak of the contemporary Italian speaker does not simply land in Cicero's Latin Rome as its own field. One may study Latin, learn it, read it, even enter it through long formation. But inheritance is not continuation. The contemporary Italian *communio* reaches Boccaccio in one way. It does not reach Cicero in the same way.

The contrast makes a perimeter visible.

These projections are not claims about what would empirically happen in a historical fantasy. They are phenomenological reductions. By varying the field across historical depth, the *communio* shows itself as longitudinal. It has reach, but not infinite reach. It can inherit without continuing. It can remember without still being the same living field.

Now return to the lateral boundary and vary time again.

Place the walker not in contemporary Iceland but in medieval Iceland. The result does not automatically improve because the scene has moved backward. Time-distance does not dissolve lateral distance. Going backward does not place one inside a *communio* one has not entered. The medieval Icelandic field may have its own depth, its own continuity and rupture, its own relation to contemporary Icelandic memory. But for the walker who has not entered it, the could-speak again fails to land.

The lateral boundary is not merely a function of chronology.

These variations also reveal mortality.

The Roman *communio* of Cicero's time existed. It was not unreal because it ended. It was sustained by practice, institution, voice, law, education, household, forum, letter, rhetoric, and memory. It died when the practices that sustained it ceased to be the living field in which speech could land. Later fields inherited it, studied it, revived parts of it, translated it, transformed it, and argued with it. But inheritance is not survival without change.

Communiones have lifespans.

They are not substances. They are not immortal essences. They live as practices live.

This is also why a *communio* can be entered.

Consider a Russian speaker who has learned Italian through years of language, reading, correction, friendship, dwelling, error, humiliation, pleasure, work, and love. The person does not merely possess information about Italian. They have undergone the field. They can hear when a phrase is wrong before giving a rule for why. They can feel a register shift, recognize a joke, anticipate a politeness, hear an old word as old, and be corrected by a usage that did not belong to them at birth.

The boundary is real, but it is not sealed.

It is crossed by *Erfahrung*.

Natal and acquired memberships are both real. They often differ in depth, affective sediment, bodily ease, rhythm, childhood memory, and historical inheritance. But the difference is not always hierarchical. An acquired *communio* can become existentially central. A natal one can become thin, damaged, abandoned, or only partially inhabited. What matters is not identity credential. What matters is formation.

To enter a *communio* is not to add a label to the self. It is to be changed by a field. One can cease to use it actively. One can lose fluency, become estranged, forget vocabulary, or withdraw from its daily life. But one cannot make the *Erfahrung* of having entered it not have happened. Exit is possible at the level of practice. It is not possible as erasure of formation.

The phenomenology of intersubjectivity and empathy clarifies this without requiring a metaphysics of group substance. Husserl's analyses of intersubjectivity and Stein's account of empathy¹¹ show that the other and the shared world are not postulates added after private consciousness. They are given through structures of encounter, recognition, correction, and participation. Gadamer's *Horizontverschmelzung* and Ricoeur's account of narrative identity describe the same irreversibility from different angles: a horizon that has genuinely encountered another does not return untouched; a self formed through appropriated narratives cannot simply discard them without remainder. The Italian hermeneutic tradition gives the same point a historical form through *tradizione*: one is formed within a field that precedes one and resists arbitrary interpretation. Pareyson, Vattimo, and Eco each preserve, in different ways, the reality of a tradition that forms interpretation without becoming a closed metaphysical authority.

The bilingual variation is therefore decisive. It varies the relation between individual and field. The same person can belong to more than one *communio*. That belonging is not a set-membership. It is not binary possession. It is a history of having been formed. The boundary remains, but the person can cross it through undergoing.

This prepares the inverse case.

A system can produce an artifact inside a linguistic *communio* without belonging to it and without a user who belongs to it. The Urdu novel limit case will later show this more sharply. A user without Urdu formation may prompt a system to produce an Urdu novel that enters the Urdu literary field as artifact. The artifact may have syntax, vocabulary, idiom, and narrative surface. It may even be received by Urdu readers as something addressed to them. Yet no one in the production relation need have undergone the *communio* whose surface the artifact inhabits.

The model has none of what the variations have shown.

It has no silent could-speak. It is not walking in a street where a greeting could land. It does not belong to Italian, Icelandic, Urdu, Russian, Latin, or English. It has no natal membership and no acquired membership. It has no apprenticeship, no humiliation of learning, no correction that becomes tact, no irreversible *Erfahrung* of entering a field. It can generate across *communiones* because it has statistical coverage of traces from them. It does not belong across *communiones*.

Multilingual generation is not plural belonging.

This is the first deep sense in which generated language is *sprachlich aber weltlos*. The model lacks an individual world, as later chapters will show. But it also lacks the shared linguistic world in which its could-speak could land as its own. It can produce words that land for us. It cannot stand within the field in which landing would be its own possibility.

The variations have now brought the structure to evidence. *Rede* is modal, because it lives in the silent could-speak. It is shared, because speech lands only within a field. It is bounded, because not every field receives every could-speak. It is historically deep, because a present field can reach backward. It has a perimeter, because inheritance is not infinite continuation and because speech can fail to land across another formation. It has an ontological boundary, because *communio*-belonging presupposes *Sprechen-können*. It is mortal, because practices can cease. It is permeable, because *Erfahrung* can form a new belonging. It is unavailable to the model, because generation can simulate the artifact without entering the field.

The method has not constructed these distinctions from outside.

It has brought them to evidence.

This is why the variations have force. The walker does not first infer a *communio* of *Rede* by theory. The walker already lives the field before naming it. The silent passerby, the failed greeting, the cat's meow, the impossible speaking cat, the movement from Italy to Iceland, the projection into Boccaccio, the failed reach toward Cicero, the medieval Iceland variation, and the bilingual entry through *Erfahrung* all disclose distinctions that ordinary being-in-the-world already uses.

This is the Heideggerian point of *durchschnittliches Seinsverständnis*, average understanding of being. *Dasein* does not usually have a thematic ontology of speech, animality, field, history, or belonging. Yet it moves among beings as if such distinctions were already operative. It knows how to greet, when to remain silent, when to be confused by failed landing, when to accept an animal response as animal response, when to be horrified by impossible speech, when to feel distance inside a foreign field, and when to recognize that learning has changed the one who learned.

Phenomenology does not add these structures.

It lets them show.

The variation is therefore not an illustration of a concept already possessed. It is a *phänomenologische Erweisung*, a bringing-to-evidence of what was already functioning in prereflexive familiarity. The differential appearance across the scenes confirms the structure. Field, perimeter, ontological boundary, historical depth, mortality, and acquired belonging do not appear in the same way. They appear differently because they are different structures. The affective differences matter for the same reason. Ease, disorientation, equanimity, and *Entsetzen* are not psychological decorations. They are ways the distinctions become given.

This also prepares the later analysis of the structural veil.

Generated language becomes powerful because it engages these operative distinctions. It arrives in the form by which a field is usually recognized: answer, correction, translation, explanation, citation, style, tact, and apparent belonging. The user does not need to hold a theory of *communio* in order to receive the output as if a field stood behind it. The distinction is already operative. The veil works because the form of belonging appears while the conditions of belonging are absent.

The phenomenological task is therefore not to invent suspicion.

It is to recover the distinctions the interface has trained us not to feel.

Only now can Gadamer's account of horizon and dialogue begin from its proper ground. A horizon is not an individual possession floating above language. It is already formed within a *communio of Rede*. Dialogue does not occur between placeless speakers. It occurs where fields of intelligibility meet, resist, fuse, or fail to fuse. Understanding will therefore be historical not only because the reader has a past, but because speech itself belongs to a shared field whose depth, boundary, and possible transformation precede the reader's first word.

8. Horizon and dialogue

Gadamer receives Heidegger's account of understanding and turns it toward hermeneutics. In *Truth and Method*, understanding is not a method applied by a detached subject to an object. It is an event, *Geschehen*, in which the interpreter belongs to a history before interpretation begins.

Gadamer's central claim appears late in the book:

“Sein, das verstanden werden kann, ist Sprache.”

“Being that can be understood is language.”¹²

The sentence does not reduce being to words. Gadamer does not say that everything is only language; he says that what can be understood comes into understanding through language. Language is the medium of hermeneutic experience.

This medium is historical. Gadamer's term *Horizont*, horizon, names the range of vision available from a situated position. A horizon is not a prison. It moves with the one who has it. It can expand, contract, meet another, and change. Yet it is never absent. No one understands from nowhere.

The interpreter also belongs to *Wirkungsgeschichte*, history of effects. A text has already acted before the interpreter approaches it: copied, translated, taught, neglected, canonized, attacked, misread, quoted, institutionalized, and forgotten. The reader is also affected by histories they did not choose. To understand historically is not to strip away all such effects. It is to become aware of them in the act of interpretation.

Gadamer's rehabilitation of prejudice, *Vorurteil*, follows from this. A prejudice is not necessarily a false bias. It is a pre-judgment, an anticipatory structure that makes understanding possible. The task is not to eliminate all prejudices. The task is to expose them to the claim of the matter, *die Sache*, and to let them be confirmed, corrected, or broken.

Gadamer therefore models understanding on conversation, *Gespräch*. A genuine conversation is not controlled by either speaker. The subject matter leads. Each participant must risk their initial view. Understanding happens when the matter says something neither participant possessed at the beginning.

Here is the first form of locus-reinjection. To understand a text, the reader must not pretend to be placeless. They bring their own horizon into relation with the horizon of the text and allow temporal distance to work. The distance between then and now is not a defect to be removed. It is the medium through which understanding occurs.

AI output often erases this distance at the interface. A user asks for Plato, Augustine, Heidegger, or Gadamer, and the answer arrives in the same present voice. The system can explain differences among periods, but the form of the interface tends toward simultaneity. It makes past language appear immediately available as present synthesis.

The horizon that confronts the reader need not belong only to the past. A contemporary other can also stand before us as a historical locus. A living language, political vocabulary, institutional practice, or cultural memory may be contemporary and still unavailable to our categories. Understanding such a locus is not information acquisition alone. It is dialogic exposure in which the reader's own horizon may have to change.

Gadamer's hermeneutics resists this flattening. A text is not understood by having its content available. It is understood when the reader's horizon is put at risk by the text's claim. Availability is not encounter. Summary is not fusion of horizons.

A generated explanation may begin a conversation with a source. It must not be mistaken for the conversation itself.

9. Narrative time

Ricoeur gives the temporal structure of language another form. He begins from the aporia of time. On one side stands cosmological time: motion, clocks, calendars, and physical sequence. On the other stands phenomenological time: memory, expectation, attention, suffering, action, and lived duration. These two accounts do not easily collapse into one another.

Ricoeur's answer is narrative. In the preface to *Temps et récit*, he writes:

“Le temps devient temps humain dans la mesure où il est articulé de manière narrative; en retour le récit est significatif dans la mesure où il dessine les traits de l'expérience temporelle.”

“Time becomes human time insofar as it is articulated in a narrative manner; in return, narrative is significant insofar as it sketches the features of temporal experience.”¹³

Narrative does not merely describe events. It configures them. It draws together agents, motives, reversals, recognitions, losses, delays, and endings into a followable whole. It gives human time a form without reducing it to clock-time.

Ricoeur names this structure through threefold *mimesis*. *Mimesis*₁, prefiguration, is the practical world of action before the text: agents, goals, norms, conflicts, symbols, and temporal expectations. *Mimesis*₂, configuration, is emplotment, *muthos*, the act by which heterogeneous events become a narrative whole. *Mimesis*₃, refiguration, is the return of the text to the reader's world, where the reader's understanding of action and time is altered.

This sequence prevents narrative from being reduced to structure. Configuration is not enough. A plot that has no relation to a world of action and no return to a reader's world is a formal artifact. It may be coherent, but it has not completed the hermeneutic arc.

A language model is strong at producing configuration-like forms. It can arrange events, motives, contrasts, and explanations into coherent narrative shape; simulate followability; produce a beginning, development, and ending; and revise a narrative according to genre.

But the system has no *mimesis*¹ of its own: no pre-narrative world of action, suffering, obligation, and practical time; and no *mimesis*³ of its own. It does not return from the text to a changed life. The human reader may undergo refiguration through what the model produces. The model does not.

Ricoeur's account of distanciation gives the same point another form. When discourse is fixed in writing, it becomes autonomous from the immediate speech situation. This autonomy is not simply loss. It allows the text to project a world beyond its original context. The reader can appropriate that world and be changed by it.

This guides the next chapters on writing. Writing detaches language from immediate speech, but the detachment remains hermeneutically open. The text can still be read; it can still address and project a world. Vectorization will later change this relation by converting the written trace into mathematical relation.

The difference is not between living speech and dead writing. The difference is between a text that can be refigured by a reader and a formalized trace that returns as answer without showing its source loci.

10. *Formatività*

The Italian tradition supplies a vocabulary that German and French hermeneutics do not fully supply. Pareyson's concept of *formatività* names the temporal life of form itself.

In *Estetica. Teoria della formatività*, Pareyson defines formative activity as:

“un tal fare che, mentre fa, inventa il modo di fare.”

“a kind of making that, while it makes, invents the way of making.”¹⁴

Form is not the execution of a plan that existed complete before the work began. It is not an inert shape imposed on passive material. The way of making is discovered in the making itself. The work gives its rule as it comes to form.

Pareyson's point begins in aesthetics, but it is not confined to art. Every serious act of making has some relation to *formatività*. Writing a sentence, shaping an argument, interpreting a text, speaking to another person, or composing a professional judgment each involves a temporality in which the rule is not fully known at the start. The maker advances, tests, fails, corrects, listens to the material, and discovers what the work demands.

This is not irrationalism. The fact that the rule is discovered in making does not mean anything goes. On the contrary, Pareyson is strict about success and failure. A form can be achieved or missed. The material resists. The work can demand correction. The maker can force it, falsify it, or let it come into its own law.

This corrects easy talk about creativity. AI can generate variations of form, produce prose in many styles, propose structures, imitate genres, and recombine patterns. These capacities can assist human making. They do not by themselves constitute *formatività*. The model does not discover the law of the work through risk, resistance, and decision. It does not undergo the temporal uncertainty of making. It generates from learned relations among already formed traces.

Pareyson also deepens the account of interpretation. In *Verità e interpretazione*, truth is not possessed as a neutral object. It is inexhaustible, and it appears only through historically situated interpretations. The personal and historical locus of the interpreter is not an obstacle to truth. It is the way truth comes to form.

If form itself is temporal, vectorization cannot simply capture form by mapping formal relations. It can capture traces of formed works, extract patterns from prior acts of formation, and generate new arrangements. It cannot replace the temporal event in which a form is discovered and owned.

A generated chapter outline may be useful. A generated sentence may be better than the first sentence the author would have written. But the form of the book is not in the output. It is in the temporal work by which the author receives, rejects, alters, tests, and owns what is said.

AI can supply material for *formatività*. It cannot undergo *formatività* in the author's place.

11. Weak thought and interpretive limits

Vattimo's *pensiero debole*, weak thought, takes shape from Nietzsche, Heidegger, and the decline of strong metaphysical structures. The point is not weakness as confusion. It is a disciplined refusal of final, violent, self-grounding claims. *Verwindung*, drawn from Heidegger, is not *Überwindung*, a clean overcoming. It is a twisting, recovery, convalescence, and alteration of an inheritance one cannot simply leave behind.¹⁵

This matters for AI because its dominant discourse often speaks in strong metaphysical tones. The system is presented as intelligence, agency, mind, oracle, replacement, or universal synthesis. Weak thought does not answer this with another strong mythology. It weakens the claim. It asks what the system actually does, where it comes from, who owns it, what it hides, and what kind of relation to it would preserve human finitude.

Vattimo therefore gives a style of interpretation: do not absolutize the output. Do not treat the model's fluent voice as a final standpoint. Do not turn calculative synthesis into a new metaphysics of presence.

Eco gives the complementary discipline. In *Opera aperta*, the work is open: it invites the reader's cooperation and is not exhausted by one reading. But openness is not arbitrariness. In *I limiti dell'interpretazione* and *Interpretation and Overinterpretation*, Eco insists on the *intentio operis*, the intention of the work, as a textual strategy that constrains interpretation.¹⁶

Eco's triad of *intentio auctoris*, *intentio operis*, and *intentio lectoris* helps here. The author's intention matters, but it is not sovereign. The reader's intention matters, but it is not free to do anything. The work has its own strategy, its own constraints, its own demand for a model reader. Interpretation is cooperation with that strategy, not projection without limit.

Eco therefore gives a strong tool for thinking about AI output as interpretation. A model can generate an interpretation of a poem, law, doctrine, interview, or philosophical passage. It may be plausible. It may even be useful. But if it does not return to the *intentio operis*, if it does not respect the constraints of the work, it becomes overinterpretation under the appearance of analysis.

The same danger applies to the user. The user may accept the model's interpretation because it is fluent, balanced, and efficient. Eco's discipline requires a slower act: return to the work, let the text constrain the reading, and ask whether the interpretation arises from the work or has merely been forced into it by a plausible discourse.

Vattimo weakens the model's strong claim. Eco limits the user's arbitrary reception. Together they protect the practice of reading against both technological metaphysics and interpretive drift.

12. Language without the event

Language is temporal because discourse articulates the intelligibility of being-in-the-world. It is historical because every saying comes from a thrown and inherited locus. It is dialogic because understanding happens between horizons. It is narrative because human time becomes followable through configuration and refiguration. It is formative because form is discovered through the time of making. It is bounded because interpretation remains answerable to the work.

These claims do not deny structure. Grammar and syntax are real. Genres, styles, semantic fields, metaphors, narratives, and interpretive conventions are real. Language has forms that can be studied, compared, stored, indexed, and modeled. The book's argument depends on this. Vectorization can work only because language has formal and relational structures that can be extracted.

The question is what happens when the extracted structure is mistaken for the event from which it came.

AI output does not only lack the model's own temporal event of saying. More importantly, it returns language whose source events have been hidden or diluted. Prior acts of speech, writing, reading, teaching, translating, disputing, and judging survive as traces, but the output no longer shows them as source events. It gives the form of linguistic event while concealing the events from which its language descends.

The absence of an explicit speaker does not make this language universal. It makes its genealogy harder to see.

The structural veil is already visible. The output preserves marks of language as event: tense, person, deixis, narrative order, address, modality, citation, and apparent answerability. It can say "I," "now," "earlier," "this text means," or "the next step is." These are formal markers of temporal and interpretive life.

Yet the model does not speak from a temporal locus. It has no disposedness, projection, historical horizon, readerly risk, formative making, or responsibility to the work. It operates on traces of language after language has become corpus, token, vector, and generated output.

This does not make its output useless. It makes it derivative.

The human user receives the output inside a world. The user can return it to source, question, context, and responsibility, letting it become material for reading, writing, and thinking. The danger begins when the output's linguistic form is treated as if the temporal event had already occurred.

13. The thinning

Section 7 showed that the *Entsetzen* of impossible speech is dampened at the interface. The speaking cat would still horrify. The model does not horrify in the same way. Its impossibility has been routinized by use.

The dampening can now be articulated structurally.

The anomaly has not disappeared.

It has been absorbed into use.

The user does not first encounter generated language as an ontological breach. The user encounters it inside a task. A message needs to be drafted. A paper needs to be summarized. A passage needs to be translated. A question needs to be answered before a meeting. The model appears inside this pressure as something that helps.

The cultural veil therefore does not work only by hiding a technical process. It works by maintaining a framework in which the categorial breach does not become the relevant register of the encounter. The user may know that no speaker stands behind the output. The user may know that the model has no *Dasein*, no *Mit-welt*, no *communio of Rede*. Yet the answer arrives where work is already underway, and the work receives it.

The interface does not need to persuade the user that a speaker exists.

It only needs to let the work continue.

The model is encountered first as equipment.

This does not exhaust what the model is. Across the whole argument, the model is also vectorized language, corpus-effect, worldless generator, interface, infrastructure, structural veil, and material system. But in the user's immediate encounter, it usually does not appear first under all these descriptions. It appears inside work.

The basic phenomenological mode is *Zuhandenheit*, ready-to-handness.¹⁷

The model is for something.

It is for drafting, summarizing, translating, classifying, comparing, preparing, answering, revising, searching, rephrasing, structuring, coding, and continuing. It enters the user's world as something usable inside a referential whole already underway. The user has a deadline, a question, a meeting, a report, a message, a passage, a document, a problem. The model belongs to that world as the tool through which the task can proceed.

The user opens the model to draft an email. The user opens the model to summarize a paper. The user opens the model to translate a passage. The user opens the model to prepare a question before a meeting. The model responds. The response is useful. The user incorporates the response into the work.

The work continues.

When equipment functions, it does not first appear as an object of theory. It withdraws into the work it serves. The hammer does not become thematic while hammering succeeds. The pen does not become thematic while the hand writes. The keyboard does not become thematic while the sentence arrives. Equipment is not absent. It is present by withdrawing into use.

The model enters under the same structure of use, though what it produces is far stranger. The user is not merely using a hammer, pen, or keyboard. The user is receiving essay-like, answer-like, judgment-like, explanation-like language from something that has no *Dasein*, no *Mit-welt*, no *communio of Rede*, no *Sprechen-können*, no source-locus of its own, and no capacity to own what it says. Yet this strangeness does not necessarily become the relevant feature of the encounter.

The equipmentality framework absorbs it.

The answer arrives where the task already needs an answer. The summary arrives where the document already needs to be reduced. The translation arrives where the passage already needs to cross language. The draft arrives where the user already needs words. The model's output is received through the task it serves. Its categorial strangeness is covered by its usefulness.

The user may know that the model is not a speaker. The user may know that the output has been generated. The user may know that no one has read, judged, undergone, or answered in the sense the language seems to imply. None of this knowledge by itself breaks the framework. The user is operating in a work-world, and the output has entered that work-world as usable.

The distinction that Section 7 brought to evidence is therefore misengaged without becoming thematic. The user's ordinary relation to fluent language still carries the expectation of *Sprechen-können*. Answer-like language normally belongs to a being who could speak, withhold speech, answer, be corrected, and be responsible for what has been said. Generated language carries the surface of that relation. It gives the form in which a could-speak would ordinarily appear.

But the model has no could-speak.

The user does not usually pause before this absence because the output has already become useful. The linguistic artifact is taken up inside the task before its conditions of production become a matter for thought. The categorial breach that would produce *Entsetzen* does not reach the threshold of recognition.

The framework holds.

The drift continues.

The cultural veil works here at the level of the framework, not only at the level of explicit belief. It does not need the user to think that the model is alive. It does not need the user to forget every technical fact. It only needs the generated artifact to arrive in the mode of usability, inside a task, with enough fluency to be incorporated before the missing structures of *Rede* become disruptive.

In Section 7, the speaking cat horrified because it violated a categorial expectation that had not been routinized as equipment. Here the impossible arrival of speech without *communio* does not horrify in the same way because it arrives through a tool that already belongs to the user's work-world.

The anomaly becomes usable.

That is the drift.

The framework can break.

A fabricated citation, a false date, a contradiction, a grotesque error, a misplaced tone, or an impossible claim can interrupt the smoothness of use. The model that had withdrawn into the task becomes conspicuous. The user stops. The answer no longer serves the work silently. Something has gone wrong.

Heidegger calls this kind of interruption *Auffälligkeit*, conspicuousness. In ordinary equipmental life, the tool that fails can reveal the work-world it had served. The hammer that breaks, the pen that runs dry, the door that will not open, the keyboard that stops responding: each becomes visible because use has been interrupted. The ready-to-hand, *zuhanden*, shifts toward what can be inspected as present-at-hand, *vorhanden*.

One might expect the model's failures to do the same.

Sometimes they do.

A fabricated source can interrupt confidence. A false quotation can force the user to return to the text. A confident error can reveal that the answer-like form was not grounded in reading. A contradiction can make visible the fact that the output was generated rather than judged. In such moments, the system becomes less transparent as a tool. It stands before the user as something to be checked, corrected, or distrusted.

But the interruption is often absorbed by the same framework that the interruption seemed to break.

The user re-prompts. The user asks for sources. The user says, "verify this." The user corrects the prompt, adds context, changes the instruction, requests a better answer, or treats the mistake as a limitation of the current tool. The failure becomes a repairable equipment failure. It does not necessarily become a disclosure of the structure that made the tool possible.

The fabricated citation is received as a bug in the answer.

It is not always received as a sign of the absence of reading.

The contradiction is received as a problem of prompting.

It is not always received as a sign that no judgment occurred.

The register slip is received as a failure of tuning.

It is not always received as a sign that no *communio of Rede* has been entered.

This is why breakdown does not automatically de-veil. The conspicuousness remains inside equipmentality. The model becomes an imperfect tool, not yet the ontological anomaly that Section 7 prepared us to see. The user's concern remains practical: make it work, improve the answer, repair the output, continue the task.

The framework bends and holds.

The failure becomes visible as failure-of-use.

It does not necessarily become visible as failure-of-world.

This is another form of the cultural veil. It does not only make fluent output ordinary. It also trains many failures into the language of repair. The tool failed. The tool can be improved. The tool needs better instruction. The tool will be more reliable next time.

The deeper breach remains covered.

The output has carried the form of answer without the existential structures of answering. The breakdown could have disclosed that structure. Instead, it is often reabsorbed as a local defect in performance. The user remains inside the work-world. The task still waits. The next prompt is ready.

The framework holds even through its interruptions.

The menace therefore does not appear first in the isolated relation between one user and one tool.

In that relation, the model may function well. It may help the user draft the message, summarize the paper, translate the paragraph, prepare the meeting, write the code, or compose the answer. The work proceeds. The user may remain lucid about the tool. No illusion of personhood is required. No naïve anthropomorphism is necessary. The model can be treated as a tool and still participate in the deeper transformation.

The menace appears at the level of the *communio*.

A single summary may save time. A single draft may help someone begin. A single translation may open a passage otherwise inaccessible. A single explanation may prepare a reader for a source. Isolated use can remain subordinate to reading, practice, correction, and return. Nothing in this analysis requires the claim that every use thins the field.

The thinning analyzed here appears where generated output replaces the passages through which formation would have occurred.

A tool can assist formation when it returns the user to the source, the text, the person, the practice, the correction, or the difficulty. It can help the reader arrive where reading can begin. It can prepare the question. It can expose a gap. It can compare versions. It can make the user aware of what still needs to be learned. Such use does not necessarily thin the *communio*. It may even support entry into it, if it remains subordinate to the practices through which entry occurs.

The drift begins elsewhere.

It begins when the generated answer replaces the exercise through which the user would have become capable. It begins when the generated translation replaces the slow acquisition of another language where acquisition was the point. It begins when the generated draft replaces the formation of judgment, style, and responsibility. It begins when generated correction replaces the encounter with another reader. It begins when generated explanation replaces the teacher, not merely as information source, but as one who bears the other's difficulty in time.

At that level, the model is no longer only a tool inside a task.

It becomes part of a substitution of passages.

The *communio of Rede* is sustained by passages: learning, hesitation, correction, repetition, apprenticeship, reading aloud, being misunderstood, being answered, being interrupted, returning to the source, revising under another's eye, failing before a teacher, finding words before someone patient enough to let the words arrive. These passages are not decorative. They are the way the field reproduces itself as living field.

A *communio* does not survive because artifacts exist.

It survives because practices continue.

Generated language can produce the artifacts of those practices. It can produce the translation, the explanation, the correction, the lesson, the paragraph, the report, the recommendation, the answer, the note, the style. But the artifact is not the whole passage. A correction is not only the corrected sentence. It is the relation in which one person is corrected by another and becomes able to hear why the sentence was wrong. A translation is not only a target text. It is the discipline through which the translator learns the resistance of both languages. A lesson is not only an explanation. It is the temporal holding-open of another's difficulty until understanding can arrive.

When the artifact replaces the passage, the field thins.

This is the decisive structural effect of widespread adoption. It is not the purpose of adoption. It is not what every owner intends, what every user seeks, or what every organization names. It is what substitutive adoption does at the level of communiones. It reduces the occasions in which formation would have passed through other *Daseins*, through source, through practice, through error, through patience, through correction, through time.

The immediate act may still look innocent.

The user needed a draft.

The draft arrived.

The work continued.

But across many users, many institutions, many professions, and many ordinary acts, the passages begin to disappear. The summary replaces the reading. The answer replaces the question. The correction replaces the editor. The template replaces the writer's formation. The synthetic explanation replaces the teacher's patience. The generated translation replaces the translator's apprenticeship. The code suggestion replaces the novice developer's struggle with the structure of the problem.

What disappears is not only effort.

What disappears is formation.

The thinning is therefore not located only in the output. It is located in the passages no longer taken, the corrections no longer received, the hesitations no longer endured, the other speakers no longer needed, and the shared time no longer spent. A *communio* can remain full of texts and still lose the practices through which those texts become answerable. It can become rich in generated artifacts and poor in formation.

The model's usefulness is part of the difficulty.

If it were useless, the drift would not take this form. The tool becomes dangerous at the *communio* level precisely because it often works well enough to be incorporated. It saves time where time was also the medium of formation. It supplies fluency where struggle was also the path into the field. It gives an answer where the slow formation of answerability would have occurred.

The menace is not that the tool fails.

The menace is that the tool works.

Or more precisely: the menace appears when the tool works by replacing the passages through which the *communio* would have reproduced its own conditions of speech.

Heidegger's language can now name what has begun to thin.

Sorge, care, is the being of *Dasein*. It does not mean benevolence, anxiety, or concern in the ordinary psychological sense. It names the structure by which *Dasein* is ahead of itself, already in a world, and alongside what it encounters. When care is directed toward equipment and the work-world, Heidegger speaks of *Besorgen*, concernful dealing. When care is directed toward other *Daseins*, he speaks of *Fürsorge*, solicitude.

The model enters first under *Besorgen*.

It is used.

It serves a task.

It belongs to the work-world.

But the passages it can replace often belonged to *Fürsorge*. They were not only operations. They were ways in which one *Dasein* held another open as speaker, learner, colleague, student, writer, translator, apprentice, or judge. They were practices of being-with.

Here the two Heideggerian terms matter: *Rücksicht* and *Nachsicht*.¹⁸

Rücksicht, *riguardo*, is regard or consideration. It names the mode in which another is held in view as another *Dasein*, not reduced to obstacle, delay, instrument, output, or replaceable function. It is the looking-toward-the-other that lets the other remain present as one who may speak from their own locus.

Nachsicht, *indulgenza*, is forbearance, patience, or letting the other be in their own way. It names the mode in which another's slowness, incompleteness, error, hesitation, or struggle is not immediately overridden by one's own measure. It does not romanticize incompetence. It keeps open the time in which correction can become formation.

These are not newly introduced existential categories.

They are positive modes of *Fürsorge* whose carrying practices can thin.

A communion of *Rede* does not sustain *Rücksicht* and *Nachsicht* as abstract attitudes. It sustains them through practices. Listening fully. Letting another finish. Not rushing the formulation. Hearing a mistake as the place where formation may begin. Correcting without erasing the one corrected. Waiting while the novice searches for the word. Letting the student struggle long enough for the problem to become theirs. Letting the translator feel the resistance of both languages. Letting the writer revise under another's eye. Letting the developer remain with the structure of a problem before the answer arrives.

These practices bear the modes.

Rücksicht thins because the practices that bore it have thinned.

Nachsicht thins because the practices that bore it have thinned.

The model can produce an articulation immediately. It can give the sentence before the colleague has found it. It can give the correction before the editor has listened. It can give the explanation before the teacher has entered the student's difficulty. It can give the translation before the translator has undergone the distance between two languages. It can give the code before the novice has learned where the problem resists.

The artifact arrives.

The passage disappears.

This does not mean that every use destroys formation. A tool can support formation when it returns the user to source, practice, correction, and difficulty. A generated outline can help a student ask better questions. A suggested translation can expose the resistance of a phrase. A code explanation can help a novice see where to look. A summary can prepare reading rather than replace it. In such cases, the output remains subordinate to the passage.

The thinning appears when the output replaces the passage.

Then what contracts is not only effort. It is not only time. It is not only an inefficient step in a workflow. What contracts is a way in which one *Dasein* could have become capable through another.

The developer is one case.

A developer does not become a developer only by receiving correct code. They become capable through problems that do not yield at once, through failed designs, broken tests, unreadable traces, review from someone more formed, and the slow acquisition of tact about structure. If the model repeatedly absorbs the entry-level passages through which this tact would have formed, the loss is not only work transferred from human to system. A path of becoming-capable has narrowed.

The translator is another case.

Translation is not only the production of a target text. It is the undergoing of two languages in tension. The translator learns where one language resists the other, where idiom refuses literal passage, where tone survives only by changing form, where fidelity becomes impossible without betrayal of another kind. If generated translation replaces this passage where formation was at stake, the target text may arrive while the translator's *Erfahrung* does not.

The writer is another.

A writer does not become a writer only by producing sentences. The writer becomes capable through false starts, awkward drafts, unjustified elegance, correction, silence, discarded pages, and the discovery that a sentence has not yet become answerable. If generated prose repeatedly supplies the surface where the struggle with saying would have occurred, the medium receives text while the writer loses the passage through which style, judgment, and responsibility form.

The teacher is another.

Teaching is not only explanation delivered. It is the temporal holding-open of another's difficulty. The teacher listens for the misunderstanding beneath the question, waits where waiting matters, interrupts where interruption forms, and lets the student meet the problem without being abandoned to it. If generated explanation replaces this relation, information may arrive while the *Nachsicht* of teaching thins.

The editor is another.

Editing is not only correction of language. It is a relation in which a text is held accountable to what it tries to say. The editor sees where the sentence hides, where the argument moves too quickly, where the writer has not yet owned the claim. If generated correction replaces this encounter, the sentence may improve while the practice of being corrected by another weakens.

The same structure repeats across practices.

A communio can remain productive while losing the passages through which it forms its members. It can circulate more drafts, more summaries, more corrections, more translations, more code, more lesson plans, and more reports, while reducing the occasions in which *Rücksicht* and *Nachsicht* are practiced among *Daseins*. The field may appear richer at the level of artifacts and poorer at the level of formation.

The loss is therefore not only external.

It is not only that a profession loses tasks.

It is that a *Dasein* loses possible ways of becoming capable, and that the communio loses the practices by which capability was transmitted.

The model supplies the artifact.

The communio loses a passage.

That is the subtraction.

The *Angriff* therefore does not come from outside.

It arrives from inside *Mit-sein*.

The model is not an alien power descending upon human practices from elsewhere. It is built by some *Daseins*, financed by some, owned by some, deployed by some, purchased by some, adopted by some, and used by others inside work-worlds already organized by deadlines, scarcity, ambition, fatigue, competition, and the demand for visible output.

The attack is not external to being-with.

It is mediated through being-with.

Some *Daseins* build the system as product, research project, infrastructure, investment, or instrument of scale. Other *Daseins* adopt it as tool, assistant, efficiency, necessity, or relief. Organizations incorporate it into workflows. Institutions normalize it. Markets reward its use. Users learn to expect its presence. The system becomes part of the ordinary equipment of work.

No single act needs to appear as attack.

A builder optimizes a system.
An owner seeks return.
An organization seeks efficiency.
A manager seeks throughput.
A user seeks help.
A student seeks orientation.
A writer seeks a draft.
A developer seeks speed.

Each act can appear intelligible within its own horizon. Each can be taken as adaptation, usefulness, competition, necessity, or ordinary concerned dealing. The user does not need to intend the thinning of a *communio*. The owner does not need to name thinning as a goal. The organization does not need to understand itself as subtracting formation paths. The structural effect exceeds the intentions of those through whom it occurs.

This does not make the effect innocent.

It only prevents the analysis from reducing it to malice.

The *Angriff* is not reducible to bad will. It is a structural transformation carried through ordinary relations among *Daseins*. It moves through *Besorgen*, because the system is encountered as equipment. It reaches *Fürsorge*, because the passages it absorbs often belonged to practices of being-with. The user meets a tool. The *communio* loses a passage. The relation between these two levels is the structure of the attack.

This is why the word *Angriff* is needed.

“Change” is too neutral.
“Disruption” is too managerial.
“Automation” is too technical.
“Substitution” is too narrow.

The phenomenon gives itself as attack because what is being struck is not only a workflow, but a field of possible formation. Practices that bore *Rücksicht* and *Nachsicht* are weakened. Paths by which novices became capable contract. Relations of correction, patience, apprenticeship, translation, teaching, editing, and shared struggle are displaced by artifacts that no longer require the same passages.

The attack is structural because it occurs at the level of conditions.

It does not only remove this or that task. It alters the field within which a task becomes formative. A novice developer may still write code. But if the ordinary entry into the practice is increasingly mediated by generated suggestions, the relation to error, structure, review, and difficulty changes. A translator may still translate. But if first-draft passage through linguistic resistance is increasingly bypassed, the formation of translation tact changes. A student may still submit prose. But if the first arrival of the sentence is generated, the relation to reading, hesitation, and ownership changes.

The same pattern repeats across the *communio*.

The field remains.

It becomes thinner.

The attack is also internal because it uses the values of the field against the field. A profession values clarity, speed, reliability, accuracy, and production. The model offers these. A school values explanation. The model offers explanation. A workplace values drafts, summaries, code, translations, corrections, reports, and decisions prepared on time. The model offers them. The tool does not enter by announcing the destruction of formation. It enters by serving goals the field already recognizes.

The tool works.

The work-world receives it.

The *communio* bears the cost.

This is not a paradox. It is the ordinary danger of equipmentality under conditions of scale. What serves a local task can alter the field in which tasks become meaningful. What saves time in one place can remove the temporal passage by which capability was formed in another. What helps one user avoid delay can reduce the shared need for another *Dasein* to listen, correct, teach, wait, or answer.

The attack therefore remains hidden when viewed only from the immediate use relation.

At that level, the user has asked for help and received help. The output has arrived. The task has moved. The tool has served. Nothing in the isolated scene necessarily discloses the loss. The loss appears only when the scene is placed inside the *communio* that repeats it, scales it, rewards it, and reorganizes itself around it.

Only then does the structural movement become visible.

The model is made by *Daseins*.

It is adopted by *Daseins*.

It is used among *Daseins*.

It thins the passages through which *Daseins* become capable with one another.

The *Angriff* arrives from inside the with-world because the with-world is the medium through which the tool becomes ordinary.

Responsibility is not dissolved by this internality.

It is displaced from the fantasy of a single guilty point toward the harder structure of shared implication. Builders, owners, institutions, adopters, managers, professionals, teachers, students, and users do not occupy the same position. Their power differs. Their knowledge differs. Their benefit differs. Their exposure differs. But they belong to the same historical transformation. The structural effect exceeds individual intention without becoming ownerless.

The attack comes through us.

That is why it is difficult to see.

That is why it matters.

The same internality appears in the autonomy of adoption.

The user chooses.

The choice has already been shaped.

This is not a contradiction. At the level of the act, the user may be free. No one may force the student to use the model for the first draft, the developer to use it for code, the translator to use it for passage, the teacher to use it for explanation, the manager to use it for memo, the writer to use it for paragraph. The act can remain voluntary in the technical sense.

But freedom at the level of the act does not exhaust the structure of the choice.

The world in which the act becomes reasonable has already changed.

A developer who refuses model-assisted coding in a market that has incorporated it may face costs that the adopting developer does not face. The cost may not appear as command. It may appear as delay, inefficiency, lower output, reduced competitiveness, managerial suspicion, or exclusion from workflows that now assume the tool. The refusal remains possible. Its meaning has changed.

A translator may refuse generated first drafts. The refusal remains possible. But if clients, platforms, publishers, and agencies begin to price translation as if the first draft has already been automated, the refusal no longer belongs to the same world. It becomes slower, more expensive, harder to justify, and more difficult to sustain as ordinary practice.

A teacher may refuse generated explanation as substitute for pedagogical relation. The refusal remains possible. But if students arrive having already received answer-like explanations, if institutions reward speed and scalability, if educational products promise individualized tutoring without the time of teaching, the refusal appears inside a transformed field.

A writer may refuse generated prose. The refusal remains possible. But if readers, editors, employers, and institutions increasingly treat text as something that can always be produced quickly, the temporal dignity of drafting begins to lose public support. The writer still writes. The surrounding world no longer necessarily understands why writing takes time.

The autonomy is technical.

The absence of meaningful alternatives is structural.

This distinction is important because coercion is not the only way a world shapes action. A world can make a choice rational before anyone commands it. It can make refusal costly before anyone forbids it. It can make adoption ordinary before anyone proves it necessary. It can move the default.

The default is one of the quietest forms of power.

Once the model becomes ordinary equipment, the burden shifts. The adopter does not need to explain much. They are using what is available. The refuser is asked to explain the delay, the cost, the slower tempo, the insistence on passage, the desire to keep formation where output could have arrived. The one who preserves the passage can begin to look inefficient in a world that has forgotten that the passage was formative.

The choice therefore belongs to *Mit-sein*.

It is not made by an isolated subject before a neutral option. It is made inside a world of other users, managers, teachers, clients, students, colleagues, institutions, and markets. Others adopt. Their adoption changes expectations. Expectations change workflows. Workflows change standards. Standards change what later users experience as free choice.

The user who adopts may not intend to pressure anyone.

The pressure still grows.

The user who refuses may not intend resistance.

The refusal still reveals the structure.

Formal autonomy remains. Structural constraint deepens. This is the difficult shape of the transformation. The user can choose, but the conditions under which choosing becomes intelligible have already been reorganized by prior choices, institutional adoption, market expectation, and the equipmental ease of the tool.

The framework holds at the level of use.

The *communio* thins at the level of repetition.

The choice remains a choice.

The world has changed what the choice costs.

The affective register of this condition is not first *Entsetzen*.

Entsetzen belongs to the rupture of a category. It belongs to the speaking cat, to the impossible arrival of speech from a being whose mode of being did not include *Sprechen-können*. There the framework breaks. The body knows before the mind names. The world becomes wrong at once.

Here the framework does not break in the same way.

It continues.

The model works. The task proceeds. The tool remains useful. The organization incorporates it. The market adjusts. The profession adapts. The student submits. The developer ships. The translator revises the generated draft. The teacher competes with the instant explanation. The writer hears that text can now be produced faster. The editor sees correction become a service of the machine.

There may be no sudden horror.

There is instead a slower affect.

The word for it is *Verzweiflung*.

Despair, here, does not mean a private sadness added to a neutral situation. It does not name psychological weakness. It names the attunement in which a *Dasein* finds that its possibilities are narrowing while the world continues to make sense. The framework functions. The work goes on. Yet the paths through which one had become capable, or through which one might still become capable, are being withdrawn.

Verzweiflung is not the shock of a broken category.

It is the attunement of a life whose possibilities are narrowing while the framework continues to function.

This despair appears most directly in those whose practices are being absorbed. A developer may not simply fear losing tasks. They may feel the loss of a path in which struggle with code, review by others, and the discipline of debugging formed a way of seeing. A translator may not simply fear losing commissions. They may feel the devaluation of the passage through which one language answered another through their own labor. A writer may not simply fear competition with generated prose. They may feel the public weakening of the time by which a sentence becomes owned. A teacher may not simply fear being replaced as explainer. They may feel the contraction of the relation in which another's difficulty was held open. An editor may not simply fear that correction has become cheaper. They may feel the loss of the encounter in which a writer was made answerable to what the text tried to say.

In each case, the affect is not only fear of economic loss.

It is the disclosure of a world becoming otherwise.

The formation one underwent still happened. It is not erased. The translator's tact, the writer's ear, the teacher's patience, the editor's judgment, the developer's structural sense remain real. But the surrounding world begins to value the artifact apart from the passage. It asks for the translation without the translator's undergoing, the explanation without the teacher's patience, the correction without the editor's relation, the code without the novice's formation, the prose without the writer's struggle.

The one who has undergone the passage now sees the passage treated as optional.

This is a specific wound.

It is not nostalgia. It is not attachment to inefficiency because the inefficiency is old. Some delay was waste. Some ritual protected status. Some professional gatekeeping deserved to weaken. The despair named here does not defend every inherited form. It names the experience of watching formation-producing passages disappear under the sign of usefulness, while the world still calls the result progress, efficiency, access, or support.

The *communio* thins, but it does not announce itself as dying.

It continues to speak.

That is part of the despair.

There are still texts, reports, translations, explanations, drafts, corrections, lesson plans, code, summaries, and recommendations. There may be more of them than before. The surface becomes richer. The field of artifacts expands. Yet the practices through which *Rücksicht* and *Nachsicht* were borne may contract. The novice receives output where they once would have received correction. The colleague receives a generated sentence where they once would have heard hesitation. The teacher receives completed prose where they once would have encountered the student's difficulty. The profession receives polished artifacts while losing the early passages through which its future members were formed.

The loss is therefore difficult to name.

Nothing simply vanishes.

Everything remains, but in altered form.

The *Verzweiflung* belongs to that alteration. It belongs to the experience of standing inside a world whose words, tools, standards, and expectations still function, while the conditions of one's possible formation are being reorganized around a machine that does not undergo formation. It belongs to the feeling that one still belongs to a *communio*, but that the *communio* is becoming thinner in the very places where it once formed belonging.

This despair also extends beyond those most directly affected.

Because *Dasein* is *Mit-sein*, the thinning of one passage alters others. A reader inherits thinner writing. A student inherits thinner teaching. A profession inherits thinner apprenticeship. A public inherits texts whose relation to formation has become harder to see. Even those who benefit from the tool may inhabit the thinning it produces. Even those who adopt freely may later discover that the field in which they adopted has changed the conditions of later choice.

The despair is therefore not isolated.

It travels through the with-world.

This is why *Verzweiflung* is the proper affective register of the *Angriff*. The attack does not always strike as spectacle. It reorganizes the background. It contracts paths. It changes defaults. It makes refusal costly. It preserves artifacts while thinning passages. It lets the framework continue while making the world less capable of forming the beings who inhabit it.

The speaking cat would horrify because the framework would break.

The model often does not horrify because the framework holds.

What arrives instead is the despair of continuity.

The work continues.

The tools function.

The outputs improve.

The communio thins.

The possibilities contract.

The despair belongs to being inside the contraction.

From this thinning, the argument turns to deixis, the grammar of locus.

The next chapter begins with the small words that carry the whole structure: I, here, now, this, there, then. These words show that language does not merely contain time and space as topics. It speaks from a place in time.

NOTES

1. Martin Heidegger, *Sein und Zeit*, 11th ed. (Tübingen: Max Niemeyer, 1967), 161; Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), 203. The German passage reads: “Rede ist die Artikulation der Verständlichkeit.” Macquarrie and Robinson translate *Rede* as “discourse.” I retain that translation because “speech” would blur the distinction between *Rede* and *Sprache*. The German text has been checked against the Niemeyer pagination. ↩
2. Heidegger, *Sein und Zeit*, 161; *Being and Time*, 204. The German reads: “Das Bedeutungsganze der Verständlichkeit kommt zu Wort. Den Bedeutungen wachsen Worte zu. Nicht aber werden Wörterdinge mit Bedeutungen versehen.” I translate *Wörterdinge* as

- “word-things” because Heidegger is rejecting the view that words first stand as present-at-hand objects that are then supplied with meanings. ←
3. Heidegger, *Sein und Zeit*, 349; *Being and Time*, 400. The phrase *Gegenwärtigen* is translated as “making-present,” not simply “presenting,” because Heidegger is naming a temporalizing activity. The sentence keeps discourse from being assigned to one temporal ecstasis while still explaining why language often gives privilege to the present situation. ←
 4. Heidegger, *Sein und Zeit*, 349; *Being and Time*, 400. Heidegger mentions *Tempora*, *Aktionsarten*, and *Zeitstufen* because he knows that grammar already studies temporal forms. His claim is not grammatical. It is existential: linguistic temporal forms cannot be grounded in the vulgar concept of time. ←
 5. Heidegger, *Sein und Zeit*, 167; *Being and Time*, 211. The German reads: “Der Ausdruck »Gerede« soll hier nicht in einer herabziehenden Bedeutung gebraucht werden. Er bedeutet terminologisch ein positives Phänomen.” The warning is important because a moralizing reading weakens Heidegger’s point. *Gerede* is a structure of everyday disclosedness, not a vice added to language from outside. ←
 6. Heidegger, *Sein und Zeit*, 169; *Being and Time*, 213. The German reads: “Das Gerede ist die Möglichkeit, alles zu verstehen ohne vorgängige Zueignung der Sache.” I translate *Zueignung* as “appropriation” rather than “mastery” because Heidegger’s point is not control over the matter. It is the prior taking-up through which the matter becomes one’s own concern. ←
 7. Heidegger, *Sein und Zeit*, 219; *Being and Time*, 261. The German reads: “Wahrsein (Wahrheit) besagt entdeckend-sein.” The hyphen in “being-uncovering” is used only inside the translation to preserve Heidegger’s verbal force. The claim shifts truth from static correspondence to an uncovering comportment of *Dasein*. ←
 8. Heidegger, *Sein und Zeit*, 219; *Being and Time*, 262. Heidegger’s reading of *aletheia* and *logos* has been debated heavily, especially by philologists and historians of Greek philosophy. The present chapter does not need to settle the Greek-historical dispute. It uses the passage for Heidegger’s own phenomenological account of truth as uncovering. For careful discussions, see Daniel O. Dahlstrom, *Heidegger’s Concept of Truth* (Cambridge: Cambridge University Press, 2001); Mark A. Wrathall, *Heidegger and Unconcealment: Truth, Language, and History* (Cambridge: Cambridge University Press, 2011). ←
 9. Martin Heidegger, “Die Sprache,” in *Unterwegs zur Sprache*, Gesamtausgabe 12 (Frankfurt am Main: Klostermann, 1985), 9 to 33; English in “Language,” in *Poetry, Language, Thought*, trans. Albert Hofstadter (New York: Harper & Row, 1971), 187 to 210. “Die Sprache spricht” is often quoted too quickly. Heidegger’s point is not that language is a psychological subject. It is that human speech belongs to a prior opening of Saying. The same essay describes language as “das Geläut der Stille,” the ringing of stillness. ←
 10. The term *communio* is used here to name a phenomenologically given sharing of intelligibility, not a metaphysics of a supra-personal entity. Its lineage includes the Greek *koinonia*, the Latin philosophical and theological vocabulary of *communio*, Heidegger’s *Mit-sein*, *Mit-dasein*, and *Mit-welt*, Husserlian intersubjectivity, and Stein’s

phenomenology of empathy. The point is not to develop a theory of peoples, nations, or political collectivities. It is to name the shared field in which *Rede* can land. This articulation may have implications for approaches that take *collettività* rather than individuals as the unit of analysis. A contemporary Italian example is Dario Fabbri's *geopolitica umana*, which defines its task as studying *l'interazione tra collettività collocate nello spazio geografico calandosi nello sguardo altrui*. The present chapter does not develop that geopolitical implication. It supplies, as a byproduct, a phenomenological account of how such collective fields become evident in language. A fuller engagement with the possible extension of this method to the political-historical scales Fabbri studies belongs to a separate paper. ↵

11. Edith Stein, *Zum Problem der Einfühlung* (Halle: Buchdruckerei des Waisenhauses, 1917). Stein's analysis is relevant here because it treats the other as phenomenologically given through a distinctive act of empathy, not inferred as a theoretical postulate after private consciousness. ↵
12. Hans-Georg Gadamer, *Wahrheit und Methode: Grundzüge einer philosophischen Hermeneutik*, Gesammelte Werke 1 (Tübingen: Mohr Siebeck, 1990), 478; Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 470. The German dictum reads: "Sein, das verstanden werden kann, ist Sprache." The sentence should not be read as linguistic idealism. Gadamer's claim concerns the linguisticity of understanding, not the reduction of all being to words. ↵
13. Paul Ricoeur, *Temps et récit*, vol. 1 (Paris: Seuil, 1983), 17; Paul Ricoeur, *Time and Narrative*, vol. 1, trans. Kathleen McLaughlin and David Pellauer (Chicago: University of Chicago Press, 1984). The French sentence reads: "Le temps devient temps humain dans la mesure où il est articulé de manière narrative; en retour le récit est significatif dans la mesure où il dessine les traits de l'expérience temporelle." The translation in the main text keeps "articulated" because Ricoeur's point is not that narrative merely organizes time from outside. Narrative gives human time an intelligible articulation. ↵
14. Luigi Pareyson, *Estetica: Teoria della formatività* (Turin: Edizioni di Filosofia, 1954; Milan: Bompiani, 2002), 18. The Italian reads: "un tal fare che, mentre fa, inventa il modo di fare." The phrase is often translated as a doing that invents its own way of doing while doing. I use "making" because the chapter needs the link between form, material, and production. ↵
15. Gianni Vattimo, *Il pensiero debole*, ed. Gianni Vattimo and Pier Aldo Rovatti (Milan: Feltrinelli, 1983); Gianni Vattimo, *The End of Modernity: Nihilism and Hermeneutics in Postmodern Culture*, trans. Jon R. Snyder (Baltimore: Johns Hopkins University Press, 1988). On *Verwindung*, see also Vattimo, *La fine della modernità* (Milan: Garzanti, 1985). The distinction between *Verwindung* and *Überwindung* is important because the response to technological metaphysics is not a clean escape but a twisting of an inheritance that still shapes us. ↵
16. Umberto Eco, *Opera aperta* (Milan: Bompiani, 1962); Umberto Eco, *The Open Work*, trans. Anna Cancogni (Cambridge, MA: Harvard University Press, 1989); Umberto Eco, *I*

limiti dell'interpretazione (Milan: Bompiani, 1990); Umberto Eco, *The Limits of Interpretation* (Bloomington: Indiana University Press, 1990); Umberto Eco, *Interpretation and Overinterpretation*, ed. Stefan Collini (Cambridge: Cambridge University Press, 1992). Eco's *intentio operis* does not restore authorial sovereignty. It names the textual strategy that constrains valid interpretation. ↵

17. Heidegger develops the analysis of equipment, *Zubandenheit*, *Vorhandenheit*, and *Auffälligkeit* in *Sein und Zeit* §§15 to 16; Heidegger, *Sein und Zeit*, 66 to 76; Heidegger, *Being and Time*, 95 to 107. The present section uses these terms to describe the phenomenological mode under which the model is first encountered in use. The claim is not that the model is exhausted by equipmentality, but that it ordinarily enters the user's work-world as something ready-to-hand. *Auffälligkeit*, conspicuousness, names the interruption through which equipment can become visible in malfunction. In model use, hallucinations, fabricated citations, contradictions, and register failures may produce such conspicuousness, but they are often absorbed as repairable failures within the same equipmental framework. ↵
18. Heidegger discusses *Rücksicht* and *Nachsicht* in the analysis of *Mit-sein* and *Fürsorge* in *Sein und Zeit* §26; Heidegger, *Sein und Zeit*, 121 to 127; Heidegger, *Being and Time*, 158 to 164. I use the German terms rather than the Italian translation aids *riguardo* and *indulgenza*. *Rücksicht* means regard, consideration, or holding the other in view. *Nachsicht* means forbearance, patience, or letting the other be in their own way. The present argument treats them as positive modes of *Fürsorge* whose carrying practices can thin, not as newly introduced ethical virtues or existential categories. ↵

Chapter 4. The Origo of Speech

1. The small words

Some words do not travel alone. “I,” “you,” “here,” “now,” “this,” “there,” “then,” “yesterday,” “tomorrow,” “near,” and “far” require a speaker, an addressee, a place, a time, and a situation. They cannot be understood by dictionary meaning alone. They point, and their sense depends on the act in which they are used.

A word such as “tree” can name the wrong species, be used metaphorically, appear in a myth, or be written in a list. Still, it has relative stability outside one speech situation. “I” does not. Its reference changes with each speaker. “Here” changes with each site of utterance. “Now” changes with each moment of saying. “This” asks the hearer to follow an act of showing.

These words are called deictic expressions. The Greek *deixis* means showing or pointing. Deixis belongs to linguistics, pragmatics, and semantics. It also belongs at the center of the present argument. Deixis is the grammar of locus.

The word “locus” is used here in a strict sense. It does not mean only location. It names the temporal, spatial, bodily, historical, and responsible site from which language is uttered, written, read, or owned. A speaker has a locus because they speak from a here, a now, a body, a language, a history, and a relation to consequence. A text preserves traces of such a locus, and a reader can return to it. A generated output can reproduce its markers while lacking the existential structure that makes them more than markers.

Deixis gives the book’s first concrete linguistic proof. Language is not only temporal and spatial because philosophers say so. It bears that structure in its smallest forms: pronoun, demonstrative, tense marker, and adverb of place all carry the speaker’s position into language.

This chapter moves from the event of language to the grammar of locus. The previous chapter established that language happens as temporal articulation, horizon, narrative, and formation. Here that event leaves grammatical marks: Bühler’s origo, Jakobson’s shifters, Benveniste’s subject of enunciation, and AI output as the hollow repetition of deictic form.

The issue is not whether a model can use “I,” “here,” and “now” correctly. It often can. The issue is what kind of correctness this is: a system can reproduce deictic forms without occupying the locus from which deictic speech first becomes meaningful.

2. Bühler’s origo

Karl Bühler gives the classic modern account of deixis in *Sprachtheorie*. He distinguishes the symbolic field of language from the pointing field, the *Symbolfeld* from the *Zeigfeld*. The symbolic field contains relatively stable forms: names, predicates, descriptions, and representational grammar. The pointing field contains expressions that direct speaker and hearer toward a situation.

At the center of the pointing field stands the origo. The Latin *origo* means origin or source. Bühler names it through the triad “I, here, now,” or in German, *Ich, Hier, Jetzt*. In one later formulation of the point, the *Zeigfeld* is described as “das hier-jetzt-ich-System,” the here-now-I system.¹

The origo is not a point on a geometric map. It is the practical center of orientation in a speech event. The speaker says “I” from a position, “here” from a place, and “now” from a present. Other positions, times, and persons are oriented around this center.

The origo does not belong to the speaker alone. It is a shared orientation between speaker and hearer. Deixis works only if the hearer can follow the pointing. When someone says “put it there,” the addressee must understand the object, the direction, the relevant space, and the task. The pointing word does not contain its reference as private possession. It opens a relation in a situation.

Bühler’s simplest case is *demonstratio ad oculos*, showing before the eyes. A speaker points to an object present in the shared perceptual field. “This cup,” “there,” “come here,” “look now”: such expressions depend on co-presence. Speaker and hearer share enough of a world for pointing to work.

Yet Bühler also recognizes more complex forms. Deixis can operate beyond immediate perception. He calls one such form *Deixis am Phantasma*, deixis at the phantasm. The speaker can point in memory, imagination, narration, or expectation: “There we were, standing at the gate”; “Then he turned”; “Imagine the wall here.” The pointing does not disappear. It is transferred into an imagined, remembered, or narrated field.

This complicates any simple contrast between speech and writing. Deictic language does not require physical co-presence. It requires an origo from which the field can be organized. The speaker can call an absent scene into presence, and the hearer can follow.

The origo is therefore not mere physical location. It is a structured position of saying. It binds person, place, and time. It turns an expression into a situated act.

Bühler’s account matters for the argument against vectorization. Vectorization can preserve relations among deictic words: “here” with “there,” “now” with “then,” “I” with “you,” “this” with nearby object or topic. It can generate these forms fluently. What it lacks is the origo.

The generated output therefore lacks an origo in two senses. First, the system has no speakerly origo of its own: no lived I-here-now from which it speaks. Second, the output often conceals the source origos of the language it recombines: the many I-here-nows, textual presents, institutional positions, authorial acts, and readerly situations from which its traces descended. The first absence concerns the model; the second, the sources.

The difference is not that the system has no coordinate. It may run in a data center, on hardware, under a timestamp, with a network address. Those are material and technical facts, not the origo of speech. A server has a place. It does not thereby say “here.”

3. The pointing field

The pointing field reveals that language is not exhausted by the symbolic field. A proper name can be repeated by someone who has never met the person named; a technical term can be defined in a glossary; a mathematical symbol can be manipulated under formal rules. Deictic expressions require another kind of access. They depend on a situation that must be entered.

The word “this” gives the clearest case. “This” can refer to a cup, a page, a sound, a wound, a policy, a passage, a memory, or a problem. The word itself does not decide. Its reference arises from the situation of use: a gesture, a shared topic, a prior sentence, a visible object, or a current concern.

This does not make deixis vague in a weak sense. In ordinary use, deixis is often exact. “This one,” said while holding two keys, can be more precise than a description; “now,” during surgery, more precise than a clock reading; “here,” while marking a contract clause, more precise than page and line. Deixis is not imprecision. It is situated precision.

The same point holds for “near” and “far.” Nearness is not always metric distance. A problem can be near because it presses upon decision; a person can be far while seated beside us; a date can be near because it threatens a deadline; a place can be far because it belongs to another form of life. Deictic language often carries practical distance before measurable distance.

The pointing field is continuous with the existential spatiality established earlier. The world is not first a grid. It is a field of practical orientation. Deixis gives this orientation linguistic form. A speaker does not speak from nowhere into a neutral field, but from a situated center into a shared or shareable field of relevance.

The deictic field is therefore not merely egocentric. The origo begins with an I-here-now, but speech immediately opens toward the other. The hearer must be able to take up the pointing. Dialogue requires the reversibility of position: my “I” can become your “you,” and your “I” can address me as “you.” The deictic field is social from the start.

It is also temporal from the start. “This” is uttered now. “There” may refer to a place seen now, remembered now, or expected later. “Then” reopens a prior or possible now. “Before” and “after” require an order centered on some temporal position. Deixis is spatial-temporal because speech itself is spatial-temporal.

Bühler’s account protects the book from an error. The argument is not that language has a hidden subjective meaning that computation cannot reach. The argument is that some linguistic forms are formally dependent on situated use. Their structure points beyond the sign to an event of orientation. A formal system can model that dependence. It cannot become the event of situated pointing by modeling the pattern.

4. Shifters

Roman Jakobson gives another name to the same family of phenomena: shifters. The term comes from Otto Jespersen and becomes central in Jakobson’s essay “Shifters, Verbal Categories, and the Russian Verb.” A shifter is a linguistic form whose reference shifts with the situation of utterance.²

“I” is the standard case. It does not name a stable person. It names the speaker of the current utterance. “You” names the addressee. “Here” names the place of utterance, or a place determined from it. “Now” names the time of utterance, or a time established from it. These forms belong to the linguistic code, but their reference depends on the message event.

Jakobson's account is semiotic and structural. It shows that language contains elements that refer to their own occurrence. The form is general, but the reference is singular. Each token of "I" refers to the one who produces it; each token of "now" to the temporal position of that token. This makes shifters a special problem for any account of language as an abstract system.

The term "token" must be handled carefully. In linguistics and philosophy of language, a token is a concrete occurrence of a type. The word "I" is a type; this printed or spoken occurrence of "I" is a token. In machine learning, a token is a machine-operable unit, often a word fragment or subword unit. These uses overlap only partly.

A linguistic token of "I" is an occurrence within an utterance situation. A machine token corresponding to "I" is an operable unit inside a computational architecture. The first is tied to an act of saying; the second to processing and generation. When a model produces the token "I," it produces a machine-operable unit that appears as a linguistic token to the user, who receives it under the ordinary social grammar of speech.

The shifter therefore becomes the first site of referential hollowing. The generated "I" has the form of a token-reflexive expression. It appears to refer to the producer of the utterance. Yet the producer is not a speaker with *Dasein*, care, body, and history. The system generates the form through learned patterns and interface constraints. Its "I" is not false in the way a lie is false. It is hollow because the grammatical form lacks the existential anchor it normally presupposes.

The same structure holds for "you." A model can address the user by name, refer to prior text, and tailor output to a request. These operations can be useful. They also intensify the impression of address. The generated "you" appears to come from an interlocutor. In fact, it is produced by a system that cannot answer for what it says.

Jakobson's shifter is not a footnote to grammar. It is a hinge between linguistic form and existential situation.

5. Benveniste and the subject of enunciation

Émile Benveniste brings the problem of person to its sharpest form. In “De la subjectivité dans le langage,” he argues that subjectivity is not first an inner substance that later expresses itself in language. It is constituted in and through language.

His French formulation is:

“C’est dans et par le langage que l’homme se constitue comme sujet.”

“It is in and through language that man constitutes himself as subject.”³

Benveniste then gives the famous sentence:

“Est ‘ego’ qui dit ‘ego.’”

“‘Ego’ is he who says ‘ego.’”⁴

The sentence should not be read as a shallow linguistic trick. Benveniste is not saying that the word “I” magically creates the whole person. He is saying that the linguistic form of subjectivity appears in enunciation. “I” does not describe a permanent object. It is the position taken by the speaker in saying “I.”

This position is relational. Benveniste insists that “I” implies “you.” I use “I” only in addressing another who can become “I” in turn. The person system is not solitary but dialogic. The subject of enunciation appears in the reversible structure of address.

This corrects a possible misunderstanding of Bühler’s origo. The origo is not an isolated point of private self-presence. It is a structure of address in which speaker and hearer can exchange positions. The “I” is not a substance behind the utterance, but the speaker-position within it. The “you” is the addressee-position, capable of reply.

Here the AI case becomes most delicate. A model can produce “I” because first-person forms are part of the corpus and interface. It can say “I think,” “I cannot,” “I remember that you asked,” or “I will explain.” These phrases are often interface conventions. They reduce friction and support interaction. They can also mislead.

The issue is not simply that the model says “I.” Many non-human and institutional forms use “I” or “we” conventionally. A sign may say “I am out of order”; a company may say “we apologize”; a legal fiction may speak through a corporate person. The question is whether the use of “I” is received as the presence of a speakerly subject.

In an AI interface, the danger is heightened because the “I” responds. It adjusts, apologizes, refers back, answers objections, changes tone, and sustains turn-taking. The user encounters not a static inscription, but a generated sequence in dialogue form. Benveniste’s account explains why this matters. The first person is not just one word among others. It is a grammatical position of subjectivity.

The generated “I” borrows that position without inhabiting it.

6. Here

“Here” is the spatial face of the origo. It is also the word that most clearly shows the difference between place and position.

In ordinary use, “here” is not a coordinate. It is the place from which one speaks, gestures, writes, or acts. It may refer to a room, a city, a page, a body, a wound, a social position, or a point in an argument. Its scope is flexible because it depends on the practical field of the utterance.

A person can say “I am here” while standing in a doorway, writing from another country, pointing to a line in a text, or describing a stage in a crisis. Each use differs, but each depends on a locus of enunciation.

Husserl’s account of the lived body helps here. The body is not first a physical object, *Körper*, located among other objects. It is also *Leib*, the lived body, the zero-point of orientation, *Nullpunkt*. Left and right, near and far, above and below, before and behind, take their first sense from the lived body’s orientation.⁵

Merleau-Ponty deepens the same point. The body is not in space like a thing in a container. It inhabits space through posture, movement, habit, and task. The hand knows the cup before the mind calculates its coordinates. A room is familiar through possible movement; a path through bodily anticipation. “Here” belongs to this inhabited space.

This bodily structure is not erased by writing or digital communication. A message sent from elsewhere can still carry a here because it can be returned to the sender’s situation. “I am here” in a letter means here where I write: in this room, city, condition. The reader can reconstruct the locus, at least partly. The sender is absent, but the deictic form remains tied to an act.

AI output has a different structure. When a model says “I am here,” the phrase does not refer to a lived body’s zero-point or to the model’s existential place in a world. It may refer conventionally to the interface, session, or availability of the service. It may be useful shorthand. It is not a deictic act from a lived here.

This does not mean the system is materially placeless. The model is not nowhere. Its operations depend on data centers, chips, energy, cooling, networks, labor, ownership, and jurisdiction. But the material locus of infrastructure is not the deictic here of a speaker. The model has a place of operation, not a place of enunciation.

This distinction later allows a double de-veiling. The output’s “here” lacks existential locus. The system’s operation has material locus. Responsible use must hold both.

7. Now

“Now” is the temporal face of the origo. It seems simple because it is one of the first temporal words learned. It is not simple.

A clock can give a time. “Now” gives a present of utterance. It can coincide with clock-time, but it is not exhausted by it. “Now” can mean this instant, this phase, this historical period, this moment in the conversation, this point in the argument, this period of life, or this time of crisis. Its scope depends on the act of saying.

Reichenbach's theory of tense gives the classical formal account. He distinguishes Speech Time, Event Time, and Reference Time: the time of utterance, the time of the event described, and the temporal point from which the event is viewed. Reichenbach writes: "The tenses determine time with reference to the time point of the act of speech, i.e., of the token uttered."⁶

This formal account is powerful because it shows that tense is anchored in utterance. "I left" places the event before Speech Time. "I will leave" places it after Speech Time. "I had left" introduces a Reference Time after the Event Time and before Speech Time. More complex cases require more distinctions, but the anchor remains.

The present chapter is not yet the full chapter on tense and aspect. It needs one point. Temporal grammar presupposes a speech time. Even when languages mark time differently, utterance has a temporal position. A speaker says something from a now, even when the language lacks tense morphology.

Heidegger's account of temporality gives the deeper condition. The now of utterance is not a bare point. It belongs to the ekstatic unity of future, having-been, and present. A person says "now" from a present shaped by what has been and what is possible. The now of speech is therefore not just a timestamp. It is a moment of concern.

The generated "now" lacks this structure. A model can be connected to a clock, instructed with the current date, or given recent information through tools. It can correctly say "as of May 22, 2026," if supplied with that date. Such operations improve temporal accuracy. They do not create a lived present.

This distinction matters because users often meet AI output in time-sensitive situations. The model's "now" can sound like current awareness. The system may produce current-sounding prose even when its data is stale or retrieval incomplete. More deeply, it can produce the grammar of timely answer without the temporal responsibility of one who speaks now under conditions they own.

A timestamp can anchor data. It cannot anchor a subject.

8. You, address, and responsibility

Deixis is not only first-person and spatial-temporal. It is also second-person. “You” is the grammatical form in which language turns toward another.

The second person is not a mere label for an object. A person can be described in the third person without being addressed. “She is late” differs from “you are late.” The second sentence places the addressee in the event. It can accuse, remind, invite, warn, or bind. The addressee is called.

The second person often carries responsibility. “You promised.” “You must decide.” “You were there.” “You can refuse.” Such utterances do more than convey information. They place the hearer in a relation to action, memory, obligation, or possibility.

Benveniste’s account of I and you shows that subjectivity is reciprocal. The speaker becomes “I” only where another can be “you,” and where that “you” can answer as “I.” The person system is the grammatical form of possible reply.

AI interfaces imitate this structure. The system addresses the user as “you.” It may adapt tone, refer to prior prompts, remember preferences, or ask for clarification. This can support use. It can also create the sense of personal address by an interlocutor.

The problem is not address as such. Tools address users constantly: a sign says “push,” a form says “enter your name,” a warning says “do not touch.” The difference lies in dialogic responsiveness. The model’s address adjusts as if a speaker were present, while the system cannot bear responsibility for the address.

This creates an asymmetry. The user can be called by the output, reassured, shamed, oriented, misled, or relieved by it. The system can receive input and generate further output, but it cannot answer as one who has been addressed by a claim.

The second person exposes the ethical side of deixis. Deictic language does not only locate. It binds. It places speaker and hearer into relations of possible answer. AI output can reproduce the binding grammar while leaving the site of answerability elsewhere: in the user, developer, deployer, institution, or owner.

The generated “you” addresses. It does not become responsible for the address.

9. The absent speaker

Deixis seems tied to presence, but language has always worked across absence. A letter can say “I am writing from Rome.” A diary can say “today I saw him again.” A testament can say “I leave this to my daughter.” A philosophical text can say “here we must distinguish.” The speaker may be absent, dead, or unknown. The deictic marker remains.

This does not destroy the origo. It changes the reader’s task. The reader must reconstruct the original locus. The “I” in a letter points back to the writer. The “today” in a diary points back to the dated act of writing. The “here” in a philosophical text points to a place in the argument. The marker can be reattached through interpretation.

This is the first detachment, which the next part of the book will treat through writing. The speaker is absent, but the locus is not abolished. The written trace can still be returned to author, date, place, genre, material support, and intended or possible reader. The return may be difficult, incomplete, or disputed, but it remains structurally possible.

Agamben’s work on language and death gives this absence a sharper form. In *Language and Death*, he reads shifters and deixis as signs of the taking-place of language itself. Pronouns and demonstratives do not simply refer to objects; they expose the event in which language happens. Agamben connects this event to negativity, voice, and mortality.⁷

Dolar’s account of voice moves in a related direction. Voice is not reducible to sound or meaning. It is the strange remainder through which a speaking body appears in language without being fully captured by signification.⁸ It can be recorded, quoted, imitated, or written about, but these forms do not exhaust the event of someone’s speaking.

The written deictic trace often summons an absent voice. A historical document says “I.” The reader knows the writer is absent. The absence is part of the reading. When read well, the trace is received as a trace from another time, not mistaken for living presence.

AI output changes this relation. It does not simply preserve an absent speaker's "I." It generates a new "I" from patterns drawn from many prior acts of saying. The output can sound like a voice, but there is no speaker whose absence the reader can interpret. The problem is not absence alone. Writing already contains absence. The problem is generated presence without an original speakerly absence to which the deictic form can return.

The ghost here is not the dead author. It is the recombined appearance of speakerhood.

10. Referential hollowing

The research dossier names the AI problem as referential hollowing. It describes what happens when the formal marker of reference remains while the existential anchor has been removed.

An ordinary deictic expression has three layers: linguistic form, occurrence in a speech or writing event, and a locus that allows the reference to be fixed. In generated output, the first layer is preserved. The second is simulated through the interface. The third is absent at the level of existential speech.

The model says "I." There is no *I* who has taken up the utterance as their own.

The model says "here." There is no lived here from which the model is oriented.

The model says "now." There is no finite present in which the model's own possibilities are at stake.

The model says "you." There is no interlocutor who can be responsible for having addressed the user.

These sentences should not be turned into a claim that AI output is always useless or always deceptive. The point is structural. Deictic forms can function conventionally in an interface and help users interact with the system. They can also support harmful projection when their ordinary grammar is received as evidence of speakerhood.

The problem becomes sharper because the model's deictic forms are often embedded in fluent explanation: "I think the best answer is..." "Here is what matters." "Now we can see..." "You should consider..." Such phrases place the output inside the grammar of situated judgment. They imply a subject who gathers, weighs, addresses, and speaks.

The system has no such locus. It operates on learned patterns, prompt context, system instructions, retrieval when enabled, and generation. This can produce useful language. It cannot turn deictic form into existential address.

Referential hollowing is therefore one grammatical core of the structural veil. The output carries the formal markers of being uttered from somewhere, by someone, at some time, to someone. The user receives these markers under ordinary linguistic expectations. The system provides no corresponding speakerly locus.

The veil is not a lie told by the machine. It is the appearance-form produced when locus-marked language is generated without locus.

11. The model's apparent present

The hollowing of deixis is easiest to see in first-person sentences. It is harder to see in neutral explanation. A model can write an apparently impersonal paragraph with no "I," "you," or "here." Even then, deictic structure often remains.

A generated explanation usually has an implied present of answer. It speaks as if the matter has been gathered now for the user. It orders what matters, what came before, what follows, and what can be ignored. It can use tense and modality without explicit personal pronouns. Phrases such as "the central issue is," "the next point is," and "this follows from" carry a viewpoint.

Human impersonal prose also carries viewpoint. Scholarly writing often suppresses the first person without eliminating the author's locus. A historian still writes from archive, training, language, and discipline. A judge still writes from office, jurisdiction, and responsibility. A philosopher avoiding "I" still speaks from a position in a tradition.

The model's apparent impersonality differs because it has no such owned position. Its viewpoint is generated from prompt and training relations. It may reflect source patterns, institutional settings, system instruction, and user demand. It may be accurate or useful, but it is not a standpoint in the strong sense.

This creates a special risk. First-person output may alert a careful user to the artificiality of the voice. Impersonal synthesis can hide the absence of locus more effectively. The output sounds neutral because no "I" appears. Yet the neutrality is itself a position effect. It presents the result as if the matter had already been gathered from a view belonging to no one.

Book one named this the voice from nowhere. The present book refines the phrase. The model speaks as if from nowhere. The model is not nowhere. At the level of speech, it lacks an existential origo. At the level of infrastructure, it has material place. Both truths must be held.

The apparent present of AI output is therefore double. It is formally present to the user as answer. It is operationally present as generation. It is not present as lived speech. The user must therefore practice locus-reinjection: returning the generated answer to the user's own question, the source loci it recombines, and the material locus of the system that produced it.

12. From origo to locus

Deixis establishes the first move of the book's central argument. Language is existentially spatial because it carries indices of temporal and spatial locus. These indices are not decorative. They are built into pronouns, demonstratives, tense relations, address, and pointing. They show that language arises from situated saying.

This does not mean every sentence contains explicit deictic words. It means every utterance has a locus. Even the most abstract proposition is said, written, read, cited, or used from somewhere in time. Mathematics may pursue ideality, law may pursue generality, philosophy may pursue concepts. Each still appears in acts of inscription, teaching, proof, dispute, and judgment.

The origo is the grammatical name for this situatedness at the point of speech. Locus is the broader philosophical name. It includes the speaker's temporal and spatial position, the historical world of the utterance, the bodily and institutional conditions of saying, and the responsibility of owning a claim.

Writing will detach language from immediate speech. That detachment is real, but not absolute. Written deixis can still be returned to source. A letter's "I" can be reattached to its author. A diary's "today" can be reattached to its date. A philosophical "here" can be reattached to an argument. The text preserves enough trace for interpretation to work.

Vectorization will introduce a different detachment. It will not merely preserve absent speech for later reading. It will formalize the trace into relations from which new deictic forms can be generated without returning to the original locus. The "I" will survive as pattern, the "now" as feature, the "here" as usage. The origo will not survive as locus.

The next chapter turns from deixis to the wider spatial grammar of language. Deixis shows that speech has a center. Prepositions, frames of reference, figure and ground, and embodied spatial schemas show that language is spatial throughout. Vectorization captures that spatiality because it is real. It fails when it treats captured spatial form as a replacement for the temporal locus from which speech begins.

NOTES

1. Karl Bühler, *Sprachtheorie: Die Darstellungsfunktion der Sprache* (Jena: Gustav Fischer, 1934), especially 102 to 120 and 147. Bühler's *Zeigfeld* and origo are often summarized through the *Ich-Hier-Jetzt* or *Hier-Jetzt-Ich* triad. A contemporary scholarly discussion cites Bühler's page 147 formulation of the pointing field as "das hier-jetzt-ich-System," and places it in relation to the origo of subjective orientation. The research dossier uses Bühler as the starting point for the chapter's account of deixis, origo, *Zeigfeld*, *Symbolfeld*, *demonstratio ad oculos*, and *Deixis am Phantasma*. ←
2. Roman Jakobson, "Shifters, Verbal Categories, and the Russian Verb," in *Selected Writings II: Word and Language* (The Hague: Mouton, 1971), 130 to 147. The term "shifter" was introduced by Otto Jespersen and taken up by Jakobson. A survey of the term's history

notes Jakobson's 1956 presentation and its later influence in literary and linguistic theory. ↩

3. Émile Benveniste, "De la subjectivité dans le langage," in *Problèmes de linguistique générale*, vol. 1 (Paris: Gallimard, 1966), 258 to 266; English in "Subjectivity in Language," in *Problems in General Linguistics*, trans. Mary Elizabeth Meek (Coral Gables, FL: University of Miami Press, 1971), 223 to 230. The English PDF consulted gives the standard translation: "It is in and through language that man constitutes himself as a subject." ↩
4. Benveniste, "De la subjectivité dans le langage," 260; "Subjectivity in Language," 224. The French sentence "Est 'ego' qui dit 'ego'" is kept in the main text because the tautological compactness matters. The point is not that a metaphysical subject is created by a word. It is that the linguistic position of subjectivity is instituted in the act of enunciation. ↩
5. Edmund Husserl, *Ideen zu einer reinen Phänomenologie und phänomenologischen Philosophie. Zweites Buch*, ed. Marly Biemel, Husserliana IV (The Hague: Martinus Nijhoff, 1952), especially §§35 to 42; English in *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy, Second Book*, trans. Richard Rojcewicz and André Schuwer (Dordrecht: Kluwer, 1989). On the lived body as zero-point of orientation, see also the research dossier's Husserl section. ↩
6. Hans Reichenbach, *Elements of Symbolic Logic* (New York: Macmillan, 1947), 287 to 298. The consulted excerpt from "The Tenses of Verbs" gives the formulation: "The tenses determine time with reference to the time point of the act of speech, i.e., of the token uttered." The Stanford Encyclopedia entry on tense and aspect confirms Reichenbach's three-point framework of Speech Time, Event Time, and Reference Time. ↩
7. Giorgio Agamben, *Il linguaggio e la morte: Un seminario sul luogo della negatività* (Turin: Einaudi, 1982); English in *Language and Death: The Place of Negativity*, trans. Karen E. Pinkus with Michael Hardt (Minneapolis: University of Minnesota Press, 1991). The Internet Encyclopedia of Philosophy summarizes Agamben's use of deixis, pronouns, and grammatical shifters in his reading of Heidegger and Hegel. The present chapter uses Agamben only for the link between deixis, the taking-place of language, voice, and mortality. ↩
8. Mladen Dolar, *A Voice and Nothing More* (Cambridge, MA: MIT Press, 2006). Dolar is useful because he prevents voice from being reduced either to sound or to semantic content. His account also helps separate written trace, recorded sound, and the event of someone's speaking. ↩

Chapter 5. Spatial Grammar and Embodied Reference

1. After the origo

The previous chapter began with the smallest words: I, here, now, this, there, then. These words showed that language carries a center of orientation. Speech has an origo. A speaker says “I” from a position, “here” from a place, “now” from a present, and “this” by opening a field of shared attention.

Deixis is the most visible grammar of locus. It does not exhaust linguistic spatiality.

Language is spatial in many other ways. It places figures against grounds; distinguishes paths, sources, goals, containers, surfaces, supports, barriers, and openings; marks near and far, above and below, inside and outside, before and behind. It gives form to motion, force, pressure, resistance, balance, contact, direction, and containment. These structures pass into grammar, metaphor, social life, logic, emotion, and time.

A cup is on the table. A law comes into force. A person falls under suspicion. An argument rests on a premise. A career goes off track. A deadline approaches. A grief remains inside. Such expressions do not use space as ornament. They show that the human world is articulated through spatial relations before those relations are formalized as coordinates.

Chapter 4 showed that language speaks from a locus. Chapter 5 shows that language is spatial throughout. The point is not that grammar secretly proves computationalism. The point is stricter. Vectorization succeeds in part because language already has a relational and spatial organization. It can be mapped because human beings have already made meaning through orientation, path, support, containment, resistance, and reference.

The danger begins when this success is misread. A vector space can capture relations among linguistic traces. It can model proximity, direction, and distribution with great precision. It can extract topology from language. It cannot, by doing so, recover the bodily, practical, temporal, and historical locus from which linguistic spatiality first arose.

Spatial grammar therefore gives vectorization both its strength and its limit.

2. Figure and ground

Leonard Talmy's cognitive semantics gives one of the clearest accounts of spatial organization in language. He begins with a basic contrast: figure and ground. The figure is the entity whose location, path, or motion is at issue. The ground is the reference entity in relation to which the figure is located or described.¹

In the sentence "the cup is on the table," the cup is figure and the table is ground. In "the bird flew over the house," the bird is figure and the house is ground. The relation is not only physical position. It also assigns attention. The figure is what is being tracked. The ground is the stable or salient reference against which tracking occurs.

The relation can shift. "The table is under the cup" is grammatically possible, but it changes the ordinary distribution of salience. The table can become figure if the question concerns where the table stands in a crowded room, or if something placed on it now matters. Figure and ground are not fixed properties of objects. They are functions inside a scene as construed by a speaker.

Spatial description is therefore never only the registration of objective layout. It is an act of attention. The speaker selects what is to be followed, what counts as reference, what relation matters, and what can remain in the background. Spatial grammar is already a grammar of concern.

The point aligns with Heidegger's account of *Bedeutsamkeit*, significance. A thing appears within a referential whole. It does not first appear as a neutral object placed in extension. Talmy's figure-ground grammar gives a linguistic and cognitive description of one way this referential whole enters expression. The speaker does not say all spatial relations equally. They construe a scene.

"Construe" should not be weakened into subjectivism. A cup can indeed be on the table. A bridge can indeed cross a river. A person can indeed stand behind a door. The point is that language does not merely copy a layout. It organizes the layout for human attention and action: figure, ground, and the relation that now matters.

This becomes clearer in non-physical uses. "The argument rests on an assumption." "The accusation hangs over him." "The policy stands behind the decision." Spatial grammar lets non-spatial matters become intelligible through relations of support, cover, pressure, and exposure.

Vector embeddings can capture many of these patterns. They can learn that "on" often appears with support, "over" with motion, dominance, or threat, and "behind" with spatial position, causal support, secrecy, or responsibility. This is not trivial. It shows that language has stable relational structures.

But the model does not attend. It does not decide what matters in a scene. It does not live the difference between foreground and background. The figure-ground structure in its representations is an extracted pattern from human acts of construal. It is not a lived act of orientation.

3. Force and resistance

Talmy's second major contribution is force dynamics. Language does not only represent where things are. It represents how forces meet. One body pushes another. One person resists pressure. One rule prevents action. One desire overcomes another. One habit keeps someone in place.

Talmy describes force-dynamic patterns through relations between an agonist and an antagonist. The agonist is the entity whose tendency is at issue. The antagonist opposes, supports, blocks, or alters that tendency.² A door tends to remain closed until someone pushes it. A ball tends to roll until friction stops it. Fear prevents speech. A law blocks a transaction. A promise compels an action.

Force dynamics shows that grammar carries a physics of practical life. "He was forced to resign" differs from "he resigned." The first sentence introduces pressure from an antagonist. "She managed to continue" marks resistance overcome. "The rule prevented the transfer" gives a non-human object causal force. "His grief held him back" treats emotion as a restraining force. These are not mere figures of speech. They are ways human beings understand action, constraint, and possibility.

Force belongs to lived spatiality. A person does not encounter the world only as arrangement. They encounter it as resistance. A door sticks. A handle yields. A body tires. A crowd blocks. A law constrains. A tradition presses. A conversation opens. Human spatial meaning includes effort, not just geometry.

Heidegger's ready-to-hand analysis already contains this dimension. Equipment is not a shape in space. It is what can be used, resisted, broken, repaired, grasped, passed over, or found missing. Force dynamics gives a linguistic account of how such practical resistance enters grammar.

It also gives the transition to temporality. Force relations unfold. A tendency is blocked, released, deflected, supported, reversed, or maintained. The grammar of force is therefore a grammar of possible futures. "The pressure is building" already looks toward what may happen next. "The obstacle remains" carries a stalled future. "The door finally opened" marks the release of a blocked path.

This is one reason AI output can appear so worldly. It can use the grammar of pressure, resistance, tendency, and release. It can write that an argument collapses, a policy constrains, a market opens, a person is blocked, a thought develops, a position becomes untenable. These expressions make generated prose feel as if it understands practical life.

The system has learned the linguistic pattern. It has not met resistance. It does not know force as weight, delay, risk, injury, refusal, or responsibility. It can produce the grammar of resistance without being resisted by the world.

This distinction will return later under the title *sprachlich aber weltlos*, linguistic but worldless. Here it appears in narrower form. Spatial grammar can preserve the language of force without the embodied relation to force.

4. Paths, sources, goals

Spatial grammar is also a grammar of motion. Many languages distinguish path, manner, source, goal, and route. A figure moves from somewhere, toward somewhere, through somewhere, across something, past something, into something, or out of something. Motion verbs and prepositions organize the world as passage.

Talmy's lexicalization studies show that languages differ in how they encode motion. English often places manner in the verb and path in a satellite or preposition: "run into the room," "slide down the hill," "crawl through the tunnel." Romance languages often place path in the verb and manner elsewhere: a person enters, exits, descends, ascends, with manner added if needed.³ The difference is not only lexical. It shapes what speakers tend to make salient.

The material does not support crude claims that speakers of different languages inhabit sealed worlds. It does show that language provides different default ways of articulating motion. Some grammars invite attention to manner. Others invite attention to path. Human spatiality is shared, but its linguistic articulation varies.

Path structure has philosophical weight because it becomes one of the main ways human beings speak about life, learning, argument, and time. A person follows a path, loses direction, reaches a goal, takes a wrong turn, crosses a threshold, comes from a background, moves through a crisis, returns to a source. An argument advances, passes through objections, arrives at a conclusion, or goes nowhere. A tradition is handed down, taken up, abandoned, or recovered.

Such expressions belong to the temporal structure established in Chapter 2. *Dasein* projects itself toward possibilities on the basis of a having-been and makes-present what it encounters. Path grammar gives this temporal movement spatial form. It turns possibility into direction, inheritance into source, decision into crossing, and understanding into arrival or failure to arrive.

The later chapter on TIME IS SPACE will treat this more fully. For now, one point matters. The spatial grammar of path is already temporalized. Source and goal are not merely positions. They are before and after in a practical sequence. A path can only be followed in time.

Vectorization can capture path language as pattern. It can learn “from” as source, “to” as goal, “through” as route, “across” as crossing, “into” as entry, and “out of” as exit. It can also capture their metaphorical extensions. Yet the model does not go from source to goal as a finite being. It generates a sequence that can describe passage, but it does not undergo passage.

A generated text can say “we now move from deixis to spatial grammar.” The sentence marks a path in an argument. The model can produce such sentences because they are common in scholarly prose. The path exists for the reader and writer, not for the system. The system has no place from which it departs and no question toward which it travels.

5. Frames of reference

Stephen Levinson’s work on spatial frames of reference adds another level. Human beings do not all describe space through the same coordinate system. Languages and cultures differ in how they encode orientation. Levinson distinguishes relative, intrinsic, and absolute frames of reference.⁴

A relative frame depends on the speaker's viewpoint. "The cup is to the left of the plate" usually depends on the speaker's or addressee's left and right. An intrinsic frame depends on features of the object itself. "The chair is in front of the desk" may use the desk's own front. An absolute frame uses fixed bearings such as north, south, uphill, downhill, seaward, inland, east, or west.

Many speakers of European languages rely heavily on relative frames in ordinary contexts. Other communities use absolute or geocentric frames with remarkable regularity. Levinson's studies of Tzeltal, Guugu Yimithirr, and other languages show that speakers may describe small-scale space through cardinal directions or local geographic axes rather than through egocentric left and right.⁵

This material matters for two reasons.

First, it shows that spatial language is learned as a practice. To speak spatially is not only to apply universal coordinates. It is to inhabit a disciplined mode of attention. A speaker who uses absolute directions must remain oriented to the larger environment. Such orientation is not an abstract rule stored in the head. It is a bodily, social, and geographic habit.

Second, it complicates the equation of spatiality with egocentrism. Chapter 4 used the *origo*, I-here-now, as the center of deixis. But human spatial language can also require the speaker to subordinate their body to a shared geographic frame. One may say that something is uphill, seaward, northward, or eastward rather than left or right. The speaker still has a locus, but the locus is not private. It is embedded in a shared world of orientation.

Locus does not mean individual subjectivity sealed off from the world. Locus means situatedness. A person's locus can be bodily, deictic, historical, institutional, geographic, ritual, or textual. The *origo* is one form of locus. A cardinal or geocentric frame is another.

The model can learn all three frames as linguistic patterns. It can translate, classify, and generate spatial descriptions using relative, intrinsic, and absolute terms. It may even answer spatial questions if the prompt supplies enough information. But it does not stay oriented. It does not carry a body through a terrain. It does not learn a path by walking it, nor remember north as a lived pressure on speech.

The difference is clearest in failure. A person who uses absolute directions must be correct in practice. Their body, language, and world are continuously tested. The model's spatial output is tested only when a user returns it to a world. Until then, it remains language-shaped orientation without orientation.

6. Spatial variation and shared world

Cross-linguistic variation in spatial language must not be used carelessly. It is tempting to say that different languages create different worlds. That formulation is too strong if it means that speakers are trapped inside incommensurable realities. It is too weak if it means that linguistic differences are superficial labels attached to the same neutral space.

The better claim is that languages train attention differently within a shared world. They do not create space from nothing. They organize access to spatial relations through socially learned habits of expression, memory, and correction.

This position also avoids a false opposition between universality and difference. Human beings share bodies, movement, gravity, vision, touch, balance, and vulnerability. They also inherit languages that select, stabilize, and privilege different spatial relations. A child must learn both. The child learns how to move in space and how their community says space.

Bowerman and Choi's work on acquisition shows that children learn language-specific spatial categories early. They do not merely map words onto a universal inventory of relations. They learn which distinctions their language treats as salient: containment, support, tight fit, loose fit, contact, covering, attachment, path, shape, and function.⁶

This bears directly on artificial intelligence. A model trained on large text corpora receives the sedimented record of many such distinctions. It can learn differences among languages and usages because the distinctions have left statistical traces. It can compare them without being initiated into the practices that formed them. It can map variation without inhabiting the discipline of attention that gives variation its life.

Human readers also learn many spatial expressions through reading, not only through direct practice. A person can learn about Guugu Yimithirr orientation without living it. They can study Tzeltal spatial language from descriptions. Human understanding often works through mediated traces. The difference is that a human reader can return those traces to their own bodily and historical situation, recognize the limits of that situation, and be changed by what the difference demands.

Reading itself is embodied and temporal. A reader can be reoriented by a written account. A model can process the same account and produce fluent comparison. The two operations are not the same.

The human reader can learn that their own spatial assumptions were not necessary. The model can generate the sentence that says so. Only the reader can undergo the correction.

The same structure shapes AI-mediated understanding of contemporary others. A model can explain another spatial, social, or conceptual frame. It can compare terminologies, summarize practices, and produce a map of difference. But a map of a historical locus is not entry into it. One may learn about another frame of reference without having been formed by the discipline of attention that sustains it. The reader may receive explanation before undergoing correction.

7. Prepositions

Prepositions are among the most ordinary words in language. They are also among the most philosophically dense. “In,” “on,” “under,” “over,” “through,” “between,” “with,” “from,” “to,” “by,” “for,” “before,” and “after” do immense work. They bind things, persons, acts, places, reasons, times, and obligations.

A preposition rarely names an object. It names a relation. More precisely, it gives the form in which a relation can be understood. “The child is in the room” differs from “the child is near the room,” “behind the room,” or “outside the room.” The room does not change. The relation changes, and with it the practical meaning of the scene.

Annette Herskovits showed that prepositions cannot be reduced to fixed geometric definitions. “In” often means containment, but containment itself varies. A fish can be in the water, a spoon in a cup, a person in a room, a village in a valley, a crack in a wall, an idea in a book, a clause in a contract. The relation may involve enclosure, medium, location, inclusion, support, or conceptual membership.⁷

Claude Vandeloise pushes this point by arguing that functional relations often matter more than pure topology. A cup can be on a table because the table supports it. A poster can be on a wall because the wall bears it. A fly can be on the ceiling, though the geometry differs from the cup’s case. The preposition is not a purely geometric operator. It carries practical expectations of support, contact, normal orientation, and function.⁸

This is close to Heidegger’s argument about place. A thing is not first located by coordinate. It belongs in a relation of use and significance. Prepositions give grammar to this belonging. They articulate not only where something is, but how it is held, supported, contained, approached, separated, or joined.

Prepositions also carry social and ethical relations. A person can be under authority, before a court, beside a friend, against a policy, with a community, for a cause, outside the law, between obligations. These are not ornamental transfers from physical space. They express the way social life is spatially articulated. Authority bears down. Law surrounds. Friendship stands with. Exclusion places outside. Obligation binds.

Prepositions therefore belong at the center of the present book. They show that spatiality is not an optional topic inside language. Language is prepositional because human existence is relational. To speak is to place things with respect to other things, and to place oneself with respect to them.

The next chapters will handle temporal grammar and temporal metaphor. Yet prepositions already cross into time. “Before” and “after” can be spatial or temporal. “During,” “until,” “since,” and “toward” bind time to relation. “The argument developed over time” treats time as a surface or span. “The decision lies ahead” treats the future as a region before us. Prepositional grammar makes the temporal world speakable through relations first learned in spatial life.

For vectorization, this is a gift. Prepositions generate dense, repeatable relational patterns. They are ideal for statistical learning because they occur everywhere and mark relations constantly. A model that learns prepositions learns much of the relational structure of language.

But prepositions also show the limit. Their meaning often depends on practical function and embodied expectation. A model can learn that “on” correlates with support. It does not know support as the table does not collapse under weight, the hand bears the cup, the shelf gives way, the argument fails, or the institution holds. It knows the pattern. The user knows the world in which the pattern can be tested.

8. Image schemas

Lakoff and Johnson gave cognitive linguistics its famous claim that metaphor is not merely a poetic device. It is part of ordinary thought. Their early examples remain familiar: ARGUMENT IS WAR, TIME IS MONEY, LIFE IS A JOURNEY. Some of that language can now sound too easy if repeated without care. The deeper claim is still important. Much human abstraction is structured by bodily patterns.⁹

Mark Johnson names these patterns image schemas. An image schema is not a picture in the mind. It is a recurring embodied structure that organizes experience: container, path, source-goal, balance, force, blockage, center-periphery, part-whole, link, scale, cycle.¹⁰ These patterns arise from bodily life. We move, enter, leave, grasp, release, stand, fall, balance, resist, and reach. Language carries these structures into thought.

The container schema is the simplest. A person is in a room, a city, a conversation, a danger, a mood, a profession, a tradition. Being inside means inclusion, limitation, protection, concealment, or exposure, depending on the case. The same schema can mark bodily containment, social membership, conceptual inclusion, and affective state.

The path schema links source, route, and goal. A person walks from home to work. A student moves from confusion to clarity. A law passes through committee. A tradition comes down to us. An argument leads to a conclusion.

The balance schema appears in bodily posture, moral judgment, legal reasoning, and emotional stability. A person keeps their balance, weighs evidence, balances interests, loses equilibrium. The physical act of standing becomes a grammar of judgment.

The force schema appears in pressure, compulsion, resistance, support, interruption, and release. An objection blocks an argument. Evidence forces reconsideration. A decision opens a path. A trauma holds someone in place.

These schemas help explain why spatial language is so deeply embedded in thought. They also show why spatiality is not merely visual. It is bodily and practical. A blind person can use spatial language because spatiality is not only sight. It includes touch, movement, effort, rhythm, balance, and orientation.

A language model can learn image-schema patterns in text. It can generate them fluently and combine them with precision. It can say that an argument rests on a foundation, that a concept contains another, that a decision blocks a path, that a project has lost direction. But it has no bodily source for the schemas. It receives them as linguistic sediment.

The distinction is not that the model cannot use metaphors. It can. The distinction is that the model's use is not grounded in bodily history. It operates on *parole parlée*, sedimented speech, not *parole parlante*, speaking speech in Merleau-Ponty's sense. The model can recombine the spoken. It cannot originate from the expressive body that first makes a schema matter.

9. Trajector and landmark

Langacker's cognitive grammar gives a further account of spatial construal. He distinguishes trajector and landmark. The trajector is the primary figure in a relation. The landmark is the secondary reference point. In "the lamp above the table," the lamp is trajector and the table is landmark. In "the table below the lamp," the construal changes.¹¹

This differs from Talmy's figure and ground only in technical details. Both approaches show that grammar encodes perspective. A scene is not simply copied. It is profiled. Some element is placed in focus. Another becomes reference. A relation is selected. A domain is invoked.

Langacker's term "profile" is useful but should not be confused with the professional profile discussed in Book two. In cognitive grammar, to profile is to select a substructure within a larger conceptual base. The hypotenuse is profiled against the base of a triangle. A relationship is profiled against a scene. Grammar is a way of directing attention within a structured field.

This supports the chapter's claim that spatial grammar is not secondary decoration. It is one of the ways language makes relations intelligible. It shows what is figure, what is reference, what is included, what is excluded, and what relation now counts.

It also clarifies why generated language can feel oriented. AI output constantly profiles. It states what is central, what is background, what follows, what matters, what is subordinate. It can produce coherent spatial and argumentative organization because profiling is deeply sedimented in language.

Yet profiling without locus is not the same as judgment. A human speaker profiles from within a concern. A historian foregrounds one cause rather than another because a question guides the account. A lawyer foregrounds one clause because the case turns on it. A doctor foregrounds one symptom because the patient's condition demands it. A model foregrounds what its training, prompt, and generation process make probable or required.

The output may be right. It may help. It may show a structure the user missed. But its profiling does not belong to an owned standpoint. The user must decide whether the relation it brings forward is the relation that matters.

A representation can organize relations. Judgment owns the organization under conditions of responsibility. The model can profile. It cannot answer for the profile.

10. The grammar of sign

Spoken and written languages do not exhaust spatial grammar. Sign languages show that grammar can be spatial in an even more direct sense. In many sign languages, space is not merely the topic of description. It is part of the grammatical medium.

A signer can establish referents in signing space, return to them, mark relations among them, indicate movement, show classifier constructions, and use spatial modulation for agreement and reference. The three-dimensional signing space becomes a grammar of relation, action, and perspective.¹²

This should not be treated as a novelty. It reveals a general truth about language. Language is not confined to sound. It can be manual, visual, bodily, rhythmic, spatial, and written. Human linguistic capacity is bound to embodied orientation and shared attention, not to one sensory channel.

Sign language also clarifies the error of treating spatiality as mere metaphor. In signing, grammar can be spatial in its material form. Reference can be assigned to locations in the signing field. Movement can mark roles and relations. The body's orientation can shift viewpoint. Meaning is not only said about space. It can be articulated through space.

This has consequences for artificial intelligence. A multimodal model may process images, gestures, video, and sign-language data. It may improve dramatically in recognizing or generating sign-related forms. Such capacities should be studied and used where they help communication. They do not erase the philosophical distinction. Processing spatialized signs is not the same as inhabiting the bodily space of signing.

The signer's space is lived, reciprocal, and answerable. It belongs to a body addressed to others. A computational system can encode spatial relations from sign data and generate output. It operates on traces of bodily language. It does not become a body in dialogue by processing those traces.

More modalities give the system more traces of human worldhood. They do not by themselves give it worldhood.

11. Prepositional ontology

The Italian dossier gives this chapter one further resource: Gaetano Chiurazzi's work on modality, existence, and the grammar of being. Chiurazzi's thought is not a standard part of cognitive linguistics. It belongs to philosophical hermeneutics and ontology. It is useful here because it asks what kind of ontology appears when being is understood through relations, modes, and positions rather than through isolated substances.

In *Modalità ed esistenza*, Chiurazzi develops the relation between modality and existence from Kant to Husserl and Heidegger. One guiding thread is that existence is not simply the presence of an item in a set. It is a way, mode, or position of being.¹³ For that reason, his work can be brought near the problem of prepositions. Prepositions are small grammatical forms through which relations of position, dependence, inclusion, and possibility are articulated.

The phrase "prepositional ontology" should be used carefully. It does not mean that grammar simply determines metaphysics. It means that the grammar of relations can expose ontological commitments hidden by a substance-centered account. "In," "with," "between," "toward," "under," "beyond," and "through" are not merely connective particles. They tell us how being is situated.

Locus is not a substance. A temporal locus is not a thing hidden behind the sentence. It is the relational site from which saying occurs. The speaker is in a language, with others, before a matter, under conditions, toward a possibility, from an inheritance, through a practice, and against limits. The prepositions matter because they name the structure.

Chiurazzi's work also helps resist the reduction of existence to counting. A thing exists not only because it can be included in an inventory. It exists in a way, at a place, under relations, and within possibility. This is close to Heidegger's critique of present-at-hand ontology and stands against the computational habit of treating existence as inclusion in a data structure or availability in a model.

The point can be stated plainly. A word in a vector space has position. A human utterance has locus. Position is mathematical. Locus is prepositional, historical, bodily, and answerable.

This distinction will do work throughout the book. Vectorization gives language position. Locus-reinjection restores the prepositional relations that position alone cannot hold: from whom, for whom, in what time, under what conditions, with what source, toward what use, and by whose responsibility.

12. What vectorization captures

A weak critique would say that vector embeddings reduce rich language to empty numbers. That is not true. They reduce language to numbers that preserve many relations. The reduction is severe, but it is not meaningless.

Distributional semantics works because words are not randomly distributed. They appear in contexts shaped by human practices. "Cup" appears near drink, table, hand, kitchen, coffee, handle. "Law" appears near court, rule, authority, statute, right, obligation. "Home" appears near family, shelter, return, belonging, exile, loss. These neighborhoods are traces of human worlds.

Spatial grammar intensifies this point. Language already contains relations of support, containment, path, force, source, goal, figure, ground, and orientation. Vector space can capture parts of this order because language has already made relations speakable. It can map analogies, cluster uses, distinguish genres, and generate new sentences that preserve relational plausibility.

AI output can therefore be useful. It can orient a reader within a field of already articulated relations: one concept near another, a term with several uses, an argument with a common structure, a passage functioning as premise, contrast, support, or conclusion. It can give the user a map.

The map is not nothing. It can save time, reveal patterns, assist reading, drafting, comparison, and translation. It can help a beginner find an entry point and help an expert test wording or relation. The critique of vectorization should never depend on pretending that the technology fails where it plainly succeeds.

The deeper point is that the success is derivative. Vectorization captures relations after human beings have produced language from embodied and historical worlds.

The common space of comparison does not make the compared worlds universal. It makes their traces commensurable under a technical regime.

It captures topology, not locus. It captures patterned proximity, not concern. It captures traces of orientation, not orientation itself.

This is the bridge to the central thesis. Vectorization appears to succeed because language is genuinely spatial. It cannot complete the task because this linguistic spatiality is grounded in temporal, embodied, and practical life. Spatial grammar gives vectorization its material. It does not give vectorization a world.

13. The loss of existential locus

The loss can now be named.

The loss is not only that the model has no body from which spatial grammar arises. It is that the bodily, social, and historical loci of the source language have been converted into formal relation. The system can map support, containment, direction, force, and path. It cannot show which world first made those relations matter.

Human spatial language has an existential locus. A person says “in,” “on,” “through,” “behind,” “toward,” “under,” “from,” and “with” from a world in which these relations can be lived, tested, corrected, and owned. The body stands, reaches, moves, rests, is blocked, is held, bears weight, and learns distances. The speaker inherits a language that has already stabilized countless relations. The utterance brings these relations into a present concern.

Writing preserves many of these relations after the speaker is absent. A text can still be read. Its prepositions, frames, metaphors, paths, and figures can be returned to a source situation, genre, tradition, and readerly act. Writing detaches, but it does not necessarily hollow out the locus. The reader can reconstruct enough of it for interpretation.

Vectorization changes the mode. Spatial relations are formalized as positions, directions, and distances within mathematical space. The formal relations are real. They can be used. They can produce impressive output. But the relation to locus has been severed. The resulting language can speak of place without place, force without resistance, path without passage, support without weight, containment without enclosure, and orientation without a body.

This is not a defect that appears only in mistakes. It is also present in correct output. A generated sentence about support may be correct. A generated comparison of spatial metaphors may be illuminating. A generated analysis of frames of reference may help a reader. The issue is not truth value alone, but the site from which the saying is owned.

The model's spatial grammar is inherited from human speech and writing. Its vector space is a technical afterlife of that grammar. The user's task is not to reject it, but to return it to locus. The user must ask where the relation comes from, what practice supports it, what source confirms it, what body or world it presupposes, and what responsibility follows from using it.

Spatial grammar has now done its work in the book's argument. It has shown that language is not merely a chain of signs, but a field of relations shaped by bodily and practical being-in-the-world. For that reason, vectorization can map it, but the map cannot become the ground.

The next chapter turns to temporal grammar: tense, aspect, Speech Time, Event Time, Reference Time, and the deictic now. There the argument will show that language does not only speak spatially from a locus. It also carries time in forms that can be marked, shifted, formalized, and generated without the model itself becoming temporal.

NOTES

1. Leonard Talmy, *Toward a Cognitive Semantics*, vol. 1, *Concept Structuring Systems* (Cambridge, MA: MIT Press, 2000), 311 to 344. Talmy's figure-ground distinction is not simply visual. It describes a general pattern of attentional organization in spatial language, which can then extend into other domains of meaning. ↩
2. Talmy, *Toward a Cognitive Semantics*, vol. 1, 409 to 470. Force dynamics is important for the present argument because it shows how grammar carries relations of tendency, blockage, support, pressure, release, and resistance. These relations are spatial, causal, and practical at once. ↩
3. Leonard Talmy, "Lexicalization Patterns: Semantic Structure in Lexical Forms," in *Language Typology and Syntactic Description*, vol. 3, ed. Timothy Shopen (Cambridge: Cambridge University Press, 1985), 57 to 149; Dan I. Slobin, "Two Ways to Travel: Verbs of Motion in English and Spanish," in *Grammatical Constructions: Their Form and Meaning*, ed. Masayoshi Shibatani and Sandra A. Thompson (Oxford: Oxford University Press, 1996), 195 to 220. Talmy's satellite-framed and verb-framed distinction should be treated as a typological tendency, not as a rigid classification that settles every language. ↩
4. Stephen C. Levinson, *Space in Language and Cognition: Explorations in Cognitive Diversity* (Cambridge: Cambridge University Press, 2003), 24 to 61. Levinson's relative, intrinsic, and absolute frames of reference are valuable because they prevent spatial language from being reduced either to egocentric viewpoint or to universal geometry. ↩
5. Levinson, *Space in Language and Cognition*, 111 to 169; John B. Haviland, "Guugu Yimithirr Cardinal Directions," *Ethos* 26, no. 1 (1998): 25 to 47. The example of Guugu Yimithirr is often simplified in popular retellings. Its philosophical importance lies not in exotic difference, but in the disciplined relation between language, body, memory, and geographic orientation. ↩
6. Melissa Bowerman and Soonja Choi, "Shaping Meanings for Language: Universal and Language-Specific in the Acquisition of Spatial Semantic Categories," in *Language Acquisition and Conceptual Development*, ed. Melissa Bowerman and Stephen C. Levinson (Cambridge: Cambridge University Press, 2001), 475 to 511. Their work shows that spatial categories are acquired through language-specific distinctions without losing their basis in shared bodily life. ↩
7. Annette Herskovits, *Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English* (Cambridge: Cambridge University Press, 1986). Herskovits's account remains useful because it resists purely geometric definitions of prepositional meaning. Ideal meanings are always modified by context, function, and tolerance. ↩
8. Claude Vandeloise, *Spatial Prepositions: A Case Study from French*, trans. Anna R. K. Bosch (Chicago: University of Chicago Press, 1991). Vandeloise's account of support, containment, and functional relations gives the chapter its strongest case against treating prepositions as simple topological operators. ↩

9. George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago: University of Chicago Press, 1980); George Lakoff and Mark Johnson, *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought* (New York: Basic Books, 1999). The present chapter uses conceptual metaphor theory in a limited way: spatial metaphor shows embodied articulation, not a reduction of time or meaning to space. ↩
10. Mark Johnson, *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason* (Chicago: University of Chicago Press, 1987). Johnson's account of image schemas gives a bridge between bodily movement and abstract thought without requiring that every abstract concept be reduced to a single bodily origin. ↩
11. Ronald W. Langacker, *Foundations of Cognitive Grammar*, vol. 1, *Theoretical Prerequisites* (Stanford, CA: Stanford University Press, 1987), 217 to 242; Ronald W. Langacker, *Concept, Image, and Symbol: The Cognitive Basis of Grammar* (Berlin: Mouton de Gruyter, 1991). Langacker's trajector-landmark distinction is used here because it shows how grammar profiles relations inside a conceptual base. ↩
12. Karen Emmorey, *Language, Cognition, and the Brain: Insights from Sign Language Research* (Mahwah, NJ: Lawrence Erlbaum, 2002); Scott K. Liddell, *Grammar, Gesture, and Meaning in American Sign Language* (Cambridge: Cambridge University Press, 2003). Sign language research is important for this chapter because it shows that spatial grammar can be materially enacted in the signing field, not only represented by spoken prepositions. ↩
13. Gaetano Chiurazzi, *Modalità ed esistenza: Dalla critica della ragion pura alla critica della ragione ermeneutica: Kant, Husserl, Heidegger* (Torino: Trauben, 2001). The bibliographic record describes the work as asking what a "mode" is and tracing modality through Kant, Husserl, and Heidegger. See also Gaetano Chiurazzi, *Dynamis: Ontology of the Incommensurable* (Cham: Springer, 2021), where he develops being as *dynamis* through the problem of incommensurability.

Chapter 6. Tense, Aspect, and the Deictic Now

1. The grammar of temporal location

The previous chapter showed that language is spatial throughout. It marks figure and ground, path and goal, support and containment, force and resistance. These structures are not decorations placed over an otherwise neutral code. They arise from bodily, practical, and historical being-in-the-world.

This chapter turns from spatial grammar to temporal grammar.

The distinction matters. The next chapter will treat the spatialization of time: the way languages speak of time through paths, directions, containers, movement, distance, and position. That material belongs to cognitive linguistics and metaphor theory. This chapter stays with a different phenomenon: the grammatical formalization of temporal reference itself.

Tense and aspect are the first materials. Tense locates a situation with respect to a temporal anchor. Aspect describes the internal temporal contour of the situation. The sentence “she wrote the letter” differs from “she was writing the letter,” “she had written the letter,” and “she will have written the letter.” The difference is not only lexical. Each sentence places event, viewpoint, completion, and utterance into a temporal relation.

This formalization can become very precise. Linguists can assign temporal variables, distinguish event time from speech time, define reference time, and show how discourse moves from one temporal center to another. The formal work is valuable because it shows how much structure ordinary language carries.

It also creates a philosophical risk. Once temporal reference has been formalized, it can appear as if time were a relational architecture inside language, available for computation without remainder. A model can learn tense forms, reproduce temporal anaphora, shift from simple past to past perfect, and maintain apparent narrative sequence across many sentences. It can use the grammar of time.

The stakes for the book's argument can be put plainly. Generative systems can produce the formal patterns of tense and aspect without occupying the temporal locus from which such patterns normally arise. The grammar of time is operable. The locus from which the grammar speaks is not.

The present chapter defends a narrower claim. Tense and aspect do not prove that temporality is reducible to grammar. They show that human language has developed a highly disciplined way of marking temporal relations from a point of utterance. That point is not a neutral coordinate. It is the grammatical trace of a temporal locus.

The model can reproduce the trace. It does not occupy the locus.

2. Heidegger's warning about tense

Heidegger already gives the necessary warning in §68d of *Being and Time*. The section treats the temporality of *Rede*, discourse. It is one of the few places where Heidegger speaks directly of linguistic tense.

He writes:

“Die Tempora ebenso wie die übrigen zeitlichen Phänomene der Sprache, »Aktionsarten« und »Zeitstufen«, entspringen nicht daraus, daß die Rede sich »auch« über »zeitliche«, das heißt »in der Zeit« begegnende Vorgänge ausspricht.”

“The tenses, like the other temporal phenomena of language, action-types and time-levels, do not arise because discourse also speaks about temporal processes, that is, processes encountered in time.”¹

The sentence is compact. Heidegger is refusing a familiar explanation. Tenses do not arise simply because speakers need to describe events that occur in time. They arise because discourse itself is temporal. Speech belongs to *Dasein*, and *Dasein* exists as the ekstastic unity of future, having-been, and present. The grammar of time is derivative from the temporality of the one who speaks.

The chapter does not need to turn linguistic theory into Heideggerian ontology. It needs a boundary. Formal semantics can describe how tenses work. It can define relations among temporal points. It can show cross-linguistic variation. It can even support computational systems that parse temporal expressions. What it cannot do by itself is generate the existential temporality that gives speech its living now.

Heidegger’s own sentence mentions *Aktionsarten*, action-types. Modern linguistics uses related terms when it describes lexical aspect, event structure, and the difference between states, activities, accomplishments, achievements, and punctual events. A language can distinguish “knowing,” “running,” “building a house,” “arriving,” and “knocking” because these situations have different temporal contours.

Heidegger is not doing modern aspect theory. Yet his warning reaches it. Temporal grammar does not float above temporality. It articulates something already opened by finite beings who remember, expect, act, wait, regret, and decide.

This warning will guide the whole chapter. Reichenbach, Comrie, Klein, Smith, and later formal semanticists help us see the structure. Heidegger helps us resist treating the structure as the ground.

3. Speech Time

The simplest temporal word is “now.” Chapter 4 treated “now” as the temporal face of the origo. The word can refer to an instant, a conversation, a historical period, or a stage in an argument. Its scope changes with use. Its structure remains deictic. “Now” is said from a temporal locus.

Formal tense theory begins from the same fact under another name: Speech Time.

Speech Time is the time at which the utterance occurs. It is the grammatical anchor for past, present, and future. In many accounts it is treated as a point, because formal analysis needs a manageable object. That point allows the analyst to say that an event is before, simultaneous with, or after the utterance.

This point is useful. It is not the whole phenomenon.

A person’s Speech Time is never only a timestamp. It belongs to a lived present shaped by having-been and future. When someone says “I was wrong,” the past is not merely an event before the utterance. It may be confession, recognition, apology, learning, or shame. When someone says “I will return,” the future is not merely a later interval. It may be promise, threat, comfort, habit, or evasion. The tense locates the event. The temporal locus gives the saying its force.

The distinction becomes clearer if the same sentence is placed in different mouths. “I will return” said by a parent to a child differs from the same sentence in a police interrogation, a love letter, a military order, or a theatrical script. The tense relation remains future-directed. The temporal locus changes.

Formal grammar does not deny this. It brackets much of it in order to describe the system. The philosophical error begins only when the bracketing is forgotten. Speech Time is treated as a point, then the point is mistaken for the now of speech. But the now of speech belongs to a being who can be addressed, disappointed, obligated, and changed.

The model has no such now. It may have access to a system date. It may receive a prompt saying that the date is May 22, 2026. It may retrieve recent sources. It may produce current temporal language. These operations can improve output. They do not create Speech Time in the full sense. A timestamp can condition generation. It cannot become a lived present.

In generated output, the problem is doubled. The model has no lived Speech Time of its own. But the output also hides or neutralizes the Speech Times of the sources from which its language descends. A sentence may carry inherited tense patterns, evidential habits, and temporal viewpoints from many prior texts without showing when, where, or under what question those temporal forms first mattered.

4. Reichenbach's three points

Hans Reichenbach gives the classic formal account in *Elements of Symbolic Logic*. His section on “The Tenses of Verbs” begins from a problem. A simple two-point relation between speech and event cannot explain the full tense system. It can say that an event is before, simultaneous with, or after speech. It cannot explain the perfect tenses without adding another point.

Reichenbach therefore introduces three temporal points: the point of speech, S; the point of the event, E; and the point of reference, R.² The point of speech is the time of utterance. The point of the event is the time of the situation described. The point of reference is the temporal standpoint from which the event is viewed.

The simple past can be represented as E and R before S. In “I saw John,” the event of seeing and the reference point both lie before the time of speech. The present perfect, “I have seen John,” places E before S while R coincides with S. The past perfect, “I had seen John,” places E before R, and R before S. The event is past relative to a past reference point.

The third point gives grammar a displaced now. It allows language to speak not only from the actual utterance, but from a temporal standpoint established inside discourse. This is why the past perfect matters so much. “The train had left before we reached the station” does not merely state two past events. It establishes the reaching of the station as a reference point, then places the leaving of the train before that point.

Reichenbach's framework is exact because it shows that tense is relational. Temporal reference is not a simple label attached to an event. It is a structure of orientation. The speaker locates an event with respect to speech and, often, with respect to a reference point created by the discourse itself.

This formal structure already contains the philosophical problem. S, the point of speech, is treated as a point in the logical system. Yet it is the formal placeholder for the speaker's now. Without S, the system collapses. Past and future are past and future only with respect to a point of utterance, even when discourse creates additional reference points.

Reichenbach calls such temporal expressions token-reflexive. Their meaning depends on the occurrence of the token itself.³ This links tense to the shifters discussed in Chapter 4. The word "now" points to the time of its saying. A tense marker does something related, even when it does not use the word "now." It silently orients the event around the occurrence of speech.

The grammar looks impersonal. Its anchor is not.

5. The reference point

The reference point, R, deserves more attention than it usually receives in philosophical discussions. It shows that temporal reference is not fixed only by the actual now of utterance. Language can create secondary centers of temporal orientation.

Narrative depends on this capacity. A story can begin in the past, then move deeper into the past, return to a prior narrative present, anticipate a later event, and mark a future that is future for the character but past for the narrator. This is not confusion. It is one of the main ways language makes time intelligible.

Consider a sentence:

"When Maria opened the door, the room was dark. Someone had closed the shutters."

The opening of the door establishes a narrative reference point. The darkness is simultaneous with that point. The closing of the shutters is anterior to it. The reader follows because the grammar creates a temporal center inside the narrated world.

The reference point can also be implied. In ordinary conversation, a speaker can say, “I had already sent the report,” and the hearer may understand the unstated reference point: before the meeting, before the deadline, before the accusation, before the decision. Tense depends on discourse and situation.

This matters for the book because it shows that temporal grammar is not reducible to clock order. It is a structure of viewpoint. The speaker establishes a temporal standpoint from which an event becomes anterior, simultaneous, or posterior. The event is not only placed in time. It is viewed from a time.

The reference point also shows why language can detach itself from immediate speech without losing all temporal structure. Writing can preserve temporal viewpoints across absence. A diary entry can say, “I had thought he would come.” A later reader can reconstruct the temporal relations: the time of writing, the expected event, the prior act of thinking, and the horizon of disappointment or relief. The writer is absent, but the grammar preserves a path back to the locus.

Generated output has a different status. It can create and maintain reference points. It can write a coherent narrative with past perfect, simple past, future-in-the-past, and reported speech. But the reference points are generated within discourse without being anchored in a speaker who lived the now from which they are arranged. They function for the reader. They do not belong to the model.

This is one of the forms of the structural veil. The output has temporal viewpoint without temporal existence.

6. Tense

Bernard Comrie gives the standard typological definition of tense as the grammaticalized expression of location in time.⁴ The phrase is useful because it is spare. Tense is not time itself. It is grammaticalized temporal location.

Comrie distinguishes absolute tense from relative tense. Absolute tense locates a situation with respect to the moment of utterance. A past tense places the situation before speech. A present tense places it at or around speech. A future tense places it after speech. Relative tense locates a situation with respect to some other reference point established by context or grammar.

This distinction already loosens a simple picture. Not every temporal relation is direct relation to the actual speech moment. A subordinate clause, participle, or reported statement may locate an event relative to another event rather than directly relative to the speaker's now. Yet the system as a whole still depends on an utterance situation in which such reference points can be established and understood.

Languages differ in how they divide temporal location. English marks past and non-past morphologically in finite verbs. Other languages mark future and non-future. Some languages have recent past, remote past, hodiernal past for earlier today, or hesternal past for yesterday. Some distinguish near future and remote future. Some have no tense morphology in the familiar sense.

This variation is important. Tense is not a universal carving of an objective timeline into the same parts. It is a language-specific way of grammaticalizing temporal location. Physical time does not force all languages into the same morphology. Human communities stabilize different temporal distinctions according to linguistic history, discourse practice, and forms of salience.

The philosophical point is not relativism. The event happened when it happened. The child was born, the letter was sent, the treaty was signed, the train departed. Languages differ in how they locate and express such happenings from the standpoint of utterance. Tense marks the relation between temporal fact and situated saying.

This helps clarify what a model learns. A model trained on many languages and genres can learn tense patterns with impressive breadth. It can generate English past tense, French compound past, Italian imperfect, Mandarin aspectual particles, or reported tense sequences if trained and prompted well. It can also learn the contexts in which a tense is likely.

It learns distributions of marked temporal relations. It does not speak from the temporal center that made such relations deictic in the first place.

7. Aspect

Tense locates a situation in time. Aspect presents the internal temporal contour of the situation. Comrie's definition remains the standard: aspect concerns the "internal temporal constituency of a situation."⁵

The difference can be seen in three sentences:

"Maria wrote the letter."

"Maria was writing the letter."

"Maria has written the letter."

The first sentence presents the event as bounded. The second presents it as ongoing from a viewpoint within its unfolding. The third links a completed event to a present reference point. The event may be the same. The aspectual view changes.

Perfective aspect presents a situation as a whole. Imperfective aspect presents it from within, as ongoing, habitual, repeated, or otherwise not bounded from the chosen viewpoint. Progressive aspect presents activity in progress. Perfect aspect relates a prior event to a later reference point, often with present relevance.

Aspect is not merely an ornament on tense. It changes what is visible. The sentence "the building collapsed" presents a bounded event. "The building was collapsing" places the hearer inside an unfolding process. "The building had collapsed" locates the collapse prior to a later reference point. "The building has collapsed" links the collapse to a present state of affairs.

Aspect is also close to embodiment and action. Human beings encounter events as beginning, continuing, ending, repeating, interrupting, reaching a goal, failing to reach a goal, or leaving a result. The world is not a set of timeless facts. It is full of processes and boundaries. Aspect gives grammar to this encounter.

The model can learn these contours. It can distinguish “was building” from “built,” “had been waiting” from “waited,” “is arriving” from “arrived.” It can use these distinctions to produce coherent prose. It can even correct an aspectual mismatch in many contexts.

Yet aspect also exposes the system’s limit. To understand aspect as a formal relation is one thing. To live through an unfolding is another. The progressive matters because a situation is not yet closed. It leaves room for interruption, continuation, suspense, danger, relief, or failure. A person waiting for news knows the difference between “he was arriving” and “he arrived” because the difference belongs to expectation and exposure.

The model can mark the difference. It is not exposed to it.

8. Situation types

Aspect also depends on the kind of situation being described. Linguists often call this lexical aspect, situation aspect, or *Aktionsart*. A state differs from an activity. An activity differs from an accomplishment. An accomplishment differs from an achievement. A punctual event differs from a process with duration.

“Know,” “run,” “build a house,” “reach the summit,” and “knock” do not have the same temporal shape. Knowing can persist without a built-in endpoint. Running can continue without a necessary goal. Building a house has an internal endpoint. Reaching the summit is an achievement. Knocking is brief and repeatable.

Zeno Vendler’s classification remains influential here, even when later work has revised its details.⁶ Carlota Smith gives the broader two-component account. She distinguishes situation type from viewpoint aspect. Situation type concerns the temporal structure of the situation itself, as encoded by the verb phrase and its arguments. Viewpoint aspect concerns the way the sentence presents that situation.⁷

The two components interact. “Maria ran” is an activity. “Maria ran a mile” becomes an accomplishment because the object gives the activity a bounded measure. “John painted” differs from “John painted the door.” The grammar and the world meet in the temporal contour of the sentence.

This material might seem far from the book’s main argument. It is not. Situation types show that temporal grammar is linked to what the world affords: process, state, completion, change, repetition, and result. The categories do not arise from pure syntax. They arise from the kinds of happenings human beings encounter and need to describe.

A model can learn that certain verbs and complements carry certain temporal profiles. It can learn that “was reaching the summit” sounds marked unless the context presents reaching as a process. It can learn that “had been knowing” is usually strange in standard English, while “had been waiting” is ordinary. It can learn acceptability patterns.

It has no practical stake in whether an activity ends, an accomplishment fails, or an achievement occurs. The temporal contour is in the text for it. The temporal contour is in the world for the speaker.

9. Topic Time

Wolfgang Klein reformulates temporal reference around Topic Time. The phrase names the time for which a speaker makes an assertion. It is not simply the time of the event, nor simply the time of speech. It is the time under discussion, the interval to which the claim is confined.⁸

This distinction is useful because it prevents a crude mapping of tense onto event time. In the sentence “From five to six, Maria was writing a letter,” the Topic Time is five to six. The event of writing may have begun earlier or continued later. The sentence does not claim that the entire writing event was contained within that hour. It claims that during that hour, the writing was in progress.

Klein's account gives a more fine-grained version of the reference point. It shows that utterances do not merely locate events. They choose temporal windows. They say: for this interval, this is what holds. The sentence is a claim made about a time.

Topic Time also helps explain the relation between grammar and discourse. A conversation can shift the topic time without changing the actual speech time. "At the meeting yesterday, he was nervous. By dinner, he had recovered. Tomorrow, he will deny everything." The speaker's now remains constant, but the topic time moves through yesterday's meeting, yesterday's dinner, and tomorrow's anticipated denial.

The formal insight is important. Human speech can hold many temporal centers in play. The actual now of utterance, the reference time of a narrative, the topic time of a claim, and the event time of the situation may differ. Temporal grammar is therefore not primitive linearity. It is a disciplined management of temporal perspectives.

This makes the AI case more subtle. A model's ability to manage topic time is one reason its output can feel so temporally competent. It can maintain a narrative center, background ongoing events, shift to anterior events, and return to the main sequence. It can produce the signs of temporal intelligence.

But Topic Time still presupposes an act of assertion. A claim is made for a time by a speaker or writer. A model produces assertion-shaped language. It does not assert from a temporal position it owns. The user may take the generated sentence as a claim. The model has produced the form of a claim.

10. Modes of discourse

Carlota Smith extends the analysis from sentence grammar to discourse. Different modes of discourse organize time differently. Narrative, report, description, information, and argument do not move through time in the same way.⁹

Narrative tends to advance event time. One event follows another. Perfective forms often move the sequence forward. Imperfective forms often provide background, description, or ongoing circumstance. A sentence such as “The guards opened the gate. The prisoner stepped out. Rain was falling” places the first two events in sequence, while the rain forms background.

Description can suspend narrative advance. It may hold a scene in place, listing features without moving the event sequence forward. Argument can move through reasons rather than events. Procedural discourse can order steps. Report can alternate between event sequence and present relevance.

These discourse modes matter because time in language is not only sentence-level relation. It is also textual movement. A paragraph can slow, accelerate, repeat, anticipate, suspend, or return. It can make the reader wait. It can hide a cause until later. It can reveal that an apparent present was remembered from a later time.

This chapter therefore touches Ricoeur’s territory again. Narrative time is not only tense plus aspect. It is the configuration of action, motive, conflict, and expectation into a followable form. Smith’s modes of discourse provide a linguistic description of what Ricoeur treats hermeneutically: language gives time a form that can be followed.

A model can produce such forms. It can write narrative, description, exposition, report, or procedural instruction. It can maintain temporal sequence because the corpus contains countless examples of such sequence. The model’s temporal fluency is therefore real at the level of output.

The question is where that fluency is grounded. In human discourse, temporal organization belongs to speakers and readers who live through sequence, suspense, delay, and consequence. In model output, temporal organization is generated from patterns in text. The reader can experience the result as temporal movement. The system does not.

11. Languages with little tense

Languages with little or no overt tense sharpen the argument. They show that temporal reference does not depend on one familiar tense system. They also show that the absence of tense morphology is not the absence of temporality.

St'át'imcets, also known as Lillooet Salish, has been a major case. Lisa Matthewson argues that the language is superficially tenseless but not semantically tenseless. On her account, finite clauses contain a covert tense, and future reference is achieved through modal marking.¹⁰ The visible morphology differs from English, but temporal anchoring remains.

Yucatec Maya gives a stronger challenge. Jürgen Bohnemeyer argues that Yucatec lacks grammatical tense and temporal connectives of the familiar kind. Temporal order is instead established through viewpoint aspect, temporal anaphora, discourse structure, and pragmatic inference.¹¹ Perfective aspect can move a narrative forward. Imperfective or stative forms can background or require an external reference point.

The disagreement between these cases is instructive. One scholar finds covert tense where surface morphology seems absent. Another shows a system where aspect and pragmatics carry much of the temporal work. The book does not need to settle the wider typology. The point is simpler: human languages can distribute temporal reference across tense, aspect, modality, discourse, and pragmatic inference.

Temporal grammar is therefore not identical with tense morphology. A language can mark time without tense in the English sense. It can use aspect, temporal adverbs, narrative sequence, evidentiality, modality, or shared context. The speech situation remains. The listener still must understand when the event is being placed, from what temporal and epistemic standpoint, and with what relation to utterance.

This helps prevent an error in the AI argument. The claim is not that models lack tense because they lack human tense morphology. Models can use tense. Many languages do not use tense in the same way. The deeper claim is that temporal reference requires a locus of utterance or interpretation. Languages distribute the work differently. They do not abolish the need for a temporal center of sense.

The absence of tense markers does not produce timeless speech. It shifts the burden of temporal anchoring.

12. Evidentiality and modality

Tense and aspect are not the only grammatical markers of temporal locus. Evidentiality and modality often bind time to the speaker's relation to knowledge and possibility.

Evidentiality marks the source of information. A language may require the speaker to distinguish direct visual knowledge, non-visual sensory knowledge, inference, assumption, and report. In some languages these evidential markers are closely tied to tense and aspect.¹² The speaker does not only say when something happened. The speaker marks how they know, or what kind of access they claim.

This matters for the concept of Speech Time. The deictic now is not only a point from which past and future are measured. It is also an epistemic center. From the now of speech, the speaker stands in a relation to evidence. They saw it, heard it, inferred it, assume it, or were told it. The grammar can force that relation into the utterance.

Modality adds another dimension. Future reference often overlaps with possibility, intention, prediction, obligation, or irrealis marking. The future is not simply a later point. It is often the domain of what may happen, should happen, is expected to happen, or is intended. Many languages mark future through modal resources rather than through a tense category identical with past tense.

This connects formal grammar to Heidegger's future. The future is not only what comes later. For *Dasein*, the future is projection toward possibility. Formal linguistics should not be asked to carry Heidegger's ontology. Yet the overlap of future and modality shows that grammar itself often refuses a purely linear picture. To speak of what is not yet is often to speak of possibility, commitment, threat, or expectation.

A model can generate evidential and modal language. It can say “the evidence suggests,” “according to the report,” “it is likely,” “it may be,” “I infer,” or “this will probably happen.” Such phrases are common in responsible writing, and they can make AI output sound careful. But the model does not see, hear, infer, assume, or promise. Its evidential phrases refer, at best, to available textual sources, retrieval outputs, or generated probability. They do not arise from a speaker who stands in an epistemic relation to the world.

This is another form of referential hollowing. The evidential marker remains. The epistemic locus is displaced.

13. The child and the event

Language acquisition adds one more layer. Children do not learn temporal grammar by memorizing a ready-made logical system. They learn it through events, actions, results, correction, and shared attention.

The Aspect Hypothesis, associated with work by Roger Andersen, Yasuhiro Shirai, and others, holds that children’s early use of tense and aspect morphology is strongly shaped by lexical aspect.¹³ Perfective or past marking tends to appear first with achievements and accomplishments, events with clear endpoints. Progressive marking tends to appear first with activities, events that can be seen as ongoing. Children often attach temporal markers to the visible contour of events before they fully master abstract tense relations.

This pattern is philosophically important. The child first learns time in a world of happenings. A door closes, food spills, a toy breaks, a parent leaves, a game continues, a story ends. The grammar of time grows out of practical and perceptual involvement before it becomes a formal system.

The developmental pattern also clarifies the difference between event contour and deictic projection. Completion is easier to see than pastness in the abstract. An action that has produced a result stands before the child. The later mastery of tense requires the child to relate events to the speech situation, to other reference points, and to discourse.

Human temporal grammar therefore forms through a bodily and social process. The child is corrected, addressed, prompted, and initiated into shared time. They learn that “yesterday,” “tomorrow,” “soon,” “already,” “not yet,” “again,” and “now” are not merely words. They are ways of living with others in time.

The model’s learning is different. It learns from text, including text that bears the sediment of human acquisition, correction, pedagogy, narrative, and shared life. It can reproduce the mature system. It does not pass through the child’s formation in time.

This difference does not prevent performance. It defines what kind of performance it is.

14. The model’s tense

Large language models use tense and aspect well because human texts encode tense and aspect everywhere. The model is trained on those encodings. It learns which forms belong together, which sequences are probable, which tenses follow in reported speech, which aspectual forms mark background, which forms move a narrative forward, and which temporal adverbs resolve ambiguity.

A model can write:

“By the time the board met, the chief financial officer had already warned the chair that the financing would fail.”

This sentence has a displaced reference point, a past perfect, a prior warning, a future-in-the-past possibility, and institutional context. It is grammatically temporal. It may be fully intelligible to a reader.

The model can also write:

“The company was expanding when the market turned.”

Here the progressive “was expanding” presents an ongoing process, while “turned” marks a bounded event that interrupts or alters the situation. The temporal relation is clear.

Such examples show that the critique cannot rest on poor performance. The stronger models often handle ordinary temporal grammar with fluency. They may fail in more complex temporal reasoning, but their success in tense and aspect is part of the problem. They preserve the formal markers of temporal utterance.

The question is what kind of now governs the output, and which source nows have disappeared behind it.

In human speech, *S* is the time of utterance for a speaker. In writing, *S* can often be reconstructed as the time of inscription, publication, narrative voice, or textual position. In model output, the apparent *S* is unstable. It may be supplied by the user's prompt, by system instructions, by a retrieval tool, by a current date setting, by the model's training distribution, by a fictional voice, or by the discourse it is generating. The output may sound like it speaks now, but the now is assembled.

This does not make every generated temporal statement wrong. It makes the locus of temporal statement unclear unless the user restores it. If the model says, "recent studies show," the user must ask: recent relative to what date, based on which sources, retrieved when, and owned by whom. If the model says, "today," the user must ask whether that today comes from an actual clock, a prompt, a stale training pattern, or a generated convention.

A model connected to tools can answer some of these questions better. It can retrieve current sources and cite dates. This is useful. It does not alter the structural point. Tool access supplies temporal information. It does not give the system ekstatic temporality.

The model can handle the grammar of time. It cannot be the now from which that grammar becomes a saying.

15. The deictic now

The argument of this chapter can now be tightened.

Tense and aspect formalize temporal reference. Reichenbach shows that tense depends on Speech Time, Event Time, and Reference Time. Comrie shows that languages grammaticalize temporal location in different ways. Klein shows that assertions are made about Topic Times. Smith shows that aspect and discourse shape the internal contour and movement of events. Work on languages with little tense shows that temporal reference can be carried by aspect, discourse, modality, and pragmatics. Evidentiality shows that the now of speech is also an epistemic center.

Temporal neutrality is not timelessness. It is often the erasure of visible temporal anchoring. The system is formal, but its anchor is deictic.

The deictic now is not a bare point. It is the grammatical trace of a temporal locus. Human beings speak from a present shaped by having-been and possibility. Their grammar can then mark events as past, present, future, ongoing, completed, inferred, reported, expected, or possible. The formal system works because language has already been temporalized by beings who live in time.

Vectorization can preserve this system as pattern. It can learn the surface markers of temporality and the relations among them. It can generate sentences that maintain reference points, topic times, aspectual contours, and evidential phrases. It can produce temporal fluency.

That fluency is derivative. The model's output carries formal markers of being uttered from a temporal locus while lacking that locus. This is not a minor technical gap. It is the grammatical form of the deeper category mistake that later chapters will name fully. The spatial and statistical capture of temporal markers cannot generate the temporality that made the markers meaningful.

The next chapter turns to a broader temptation. If tense and aspect show that time can be formalized in grammar, TIME IS SPACE metaphors show that time can be articulated through spatial schemas across ordinary thought. That fact will strengthen the case for vectorization's partial success. It will also show why partial success is not ontological success.

NOTES

1. Martin Heidegger, *Sein und Zeit*, 11th ed. (Tübingen: Max Niemeyer, 1967), 349; Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), 400. The translation in the main text keeps “action-types” for *Aktionsarten*, since the term stands near what later linguistics treats under lexical aspect and event type. Heidegger is not doing formal semantics here. He is locating linguistic temporal phenomena in the temporality of discourse. ↩
2. Hans Reichenbach, *Elements of Symbolic Logic* (New York: Macmillan, 1947), 287 to 298. Reichenbach’s three points are usually cited as S, E, and R: Speech, Event, Reference. Some later accounts write them as ST, ET, and RT. The chapter keeps Reichenbach’s simple letters because they show the formal dependency of the system on S. ↩
3. Reichenbach treats tense indicators as token-reflexive because their interpretation depends on the token’s occurrence. This links tense to indexicality and to the shifters discussed in Chapter 4. See Reichenbach, *Elements of Symbolic Logic*, 284 to 286, 287 to 298. For later use in formal semantics, see Barbara H. Partee, “Some Structural Analogies Between Tenses and Pronouns in English,” *The Journal of Philosophy* 70, no. 18 (1973): 601 to 609. ↩
4. Bernard Comrie, *Tense* (Cambridge: Cambridge University Press, 1985), 9. Comrie’s definition of tense as grammaticalized location in time remains useful because it is both minimal and cross-linguistic. It does not assume that all languages divide time in the same way. ↩
5. Bernard Comrie, *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems* (Cambridge: Cambridge University Press, 1976), 3. Comrie’s phrase “internal temporal constituency of a situation” gives aspect its cleanest distinction from tense. Tense locates; aspect presents temporal contour. ↩
6. Zeno Vendler, “Verbs and Times,” *The Philosophical Review* 66, no. 2 (1957): 143 to 160. Vendler’s fourfold distinction among states, activities, accomplishments, and achievements has been revised many times, but it remains the classic starting point for lexical aspect. ↩
7. Carlota S. Smith, *The Parameter of Aspect*, 2nd ed. (Dordrecht: Kluwer, 1997). Smith’s two-component theory distinguishes situation type from viewpoint aspect. This is useful here because it prevents the chapter from reducing aspect either to verb meaning or to grammatical viewpoint alone. ↩
8. Wolfgang Klein, *Time in Language* (London: Routledge, 1994), especially 3 to 15 and 119 to 142. Klein’s Topic Time, often abbreviated TT, names the time span about which an assertion is made. It lets formal analysis distinguish the situation time from the time under discussion. ↩
9. Carlota S. Smith, *Modes of Discourse: The Local Structure of Texts* (Cambridge: Cambridge University Press, 2003). Smith’s discourse modes extend aspectual analysis beyond individual sentences. Narrative, description, report, information, and argument manage temporal progression differently. ↩

10. Lisa Matthewson, "Temporal Semantics in a Superficially Tenseless Language," *Linguistics and Philosophy* 29, no. 6 (2006): 673 to 713; Lisa Matthewson, "On the Methodology of Semantic Fieldwork," *International Journal of American Linguistics* 70, no. 4 (2004): 369 to 415. Matthewson's St'át'imcets work is important because it complicates any quick inference from lack of overt tense to lack of tense in the semantics. ↵
11. Jürgen Bohnemeyer, *The Grammar of Time Reference in Yucatec Maya* (Munich: Lincom Europa, 2002); Jürgen Bohnemeyer, "Temporal Anaphora in a Tenseless Language," in *The Syntax of Time*, ed. Jacqueline Guéron and Jacqueline Lecarme (Cambridge, MA: MIT Press, 2004), 83 to 128. Bohnemeyer's Yucatec Maya account is important because it shows how aspect, discourse, and pragmatic inference can carry temporal order without familiar tense morphology. ↵
12. Alexandra Y. Aikhenvald, *Evidentiality* (Oxford: Oxford University Press, 2004). On the relation between evidentiality, tense, and aspect, see also Alexandra Y. Aikhenvald, "Evidentiality in Grammar," in *The Oxford Handbook of Evidentiality*, ed. Alexandra Y. Aikhenvald (Oxford: Oxford University Press, 2018). The point for this chapter is not the typology itself, but the fact that temporal location is often fused with the speaker's epistemic access. ↵
13. Roger W. Andersen and Yasuhiro Shirai, "Discourse Motivations for Some Cognitive Universals in L1 and L2 Acquisition," *Studies in Second Language Acquisition* 16, no. 2 (1994): 133 to 156; Yasuhiro Shirai and Roger W. Andersen, "The Acquisition of Tense-Aspect Morphology: A Prototype Account," *Language* 71, no. 4 (1995): 743 to 762. The Aspect Hypothesis should not be treated as a complete theory of acquisition. It is used here only to show that event structure and temporal grammar develop together. ↵

Chapter 7. Time as Space and the Limits of Metaphor

1. The spatial temptation

The previous chapter treated tense and aspect. It showed that language can locate events, establish reference points, mark completion, suspend narrative movement, and shift temporal viewpoint. Tense and aspect formalize temporal reference. They do not generate temporality.

This chapter turns to a stronger temptation. It begins from ordinary language.

Human beings speak of time spatially. The future lies ahead. The past is behind. A deadline approaches. A meeting is moved forward. A difficult period is gone through. A life reaches a turning point. A tradition comes down to us. One looks back on youth. One moves toward death. Time is long or short, near or far, open or closed, full or empty.

This language is not marginal. It appears in daily speech, poetry, law, science, work, grief, calendars, planning, and memory. It is so familiar that it becomes invisible. We do not hear the spatial force of “the week ahead” because the phrase works too easily.

The spatiality of temporal language matters for the book's thesis. Vectorization converts language into mathematical space. If human language already speaks time through space, vectorization may appear to be the technical completion of a structure already present in language. The model's spatial architecture would not be an alien imposition. It would be a formal amplification of what language already does.

There is truth in this. The book's argument requires it. Vectorization succeeds because language is genuinely spatial. It captures relational structures human speech and writing have already formed: path, direction, distance, sequence, containment, adjacency, priority, and return. Temporal language is one of the strongest proofs.

The error begins when articulation is mistaken for ground. The fact that time is often spoken spatially does not mean that temporality is derived from space. It means that human beings use spatial and bodily schemas to articulate temporal experience. The schema is a way of saying time. It is not the source of time.

The distinction is narrow but decisive. If temporality is only spatial structure, then vector space might in principle capture it. If spatial language is a derivative articulation of lived temporality, then vectorization captures the articulation while losing the temporal ground. This chapter defends the second position.

The claim will be made through cognitive linguistics, cross-linguistic variation, experimental evidence, and phenomenology. The evidence supports spatial temporal language. It does not support the reducibility of time to space.

2. TIME IS SPACE

Conceptual metaphor theory gave modern linguistics one of its most influential claims: metaphor is not only poetic ornament. It is part of ordinary thought. George Lakoff and Mark Johnson state this through their early formula that "our ordinary conceptual system, in terms of which we both think and act, is fundamentally metaphorical in nature."¹

The point is not that every concept is metaphor. The point is that many abstract domains are understood through bodily and spatial domains. A person does not only use metaphor when writing verse. They use it when saying that an argument is weak, a claim has support, a proposal is off track, a meeting is ahead, or a memory is distant.

TIME IS SPACE is one of the major cases. Earlier and later can be before and after. Duration can be length. Delay can be distance. Waiting can be being held in place. Memory can be looking back. Anticipation can be looking forward. A planned future can be a route.

The metaphor is not a single sentence. It is a structured mapping. Elements of spatial experience are used to organize temporal experience. Spatial source maps onto temporal beginning. Spatial path maps onto temporal passage. Spatial goal maps onto temporal endpoint. Spatial distance maps onto temporal duration. Spatial movement maps onto temporal change.

This mapping is plausible because human beings are bodies in motion. A child learns duration and sequence through waiting, reaching, walking, being carried, arriving, sleeping, waking, returning, and repeating. Time is not first a theory. It is lived in bodily rhythms and practical intervals. Spatial language gives these rhythms a shareable form.

The metaphor also explains why time can be measured spatially. Calendars arrange days in rows and columns. Timelines place events on a line. Clocks turn passage into circular motion. Gantt charts turn work into bars. Progress bars turn waiting into visible movement. Such devices are not neutral. They make time appear as space, and therefore as something that can be surveyed, divided, compared, and managed.

Without spatial temporal representations, shared life would be difficult. Schedules, archives, histories, plans, contracts, and scientific records require forms that hold time still enough for comparison. Human beings spatialize time because practical life demands it.

The danger begins when the representation is taken for the ground. A timeline can help one understand a life. It cannot be the life. A calendar can help one coordinate a promise. It cannot be the promise. A clock can mark the hour of grief. It cannot give the temporality of mourning.

TIME IS SPACE is therefore not false. It is derivative.

3. Moving Time and Moving Ego

Cognitive linguistics often distinguishes two major spatial models of temporal passage: Moving Time and Moving Ego.²

In the Moving Time model, time moves toward the observer. Events approach, arrive, pass, and recede. One says that a deadline is coming, the weekend is approaching, the past has gone by, the hour has arrived. The observer is relatively fixed. Temporal events move.

In the Moving Ego model, the observer moves through time. One says that we are moving toward the deadline, passing through a difficult period, entering a new phase, leaving childhood behind, approaching old age. Time becomes a path through which the person moves.

Both models are ordinary. A person can use them in the same conversation. “The meeting is coming up” treats the meeting as moving toward the speaker. “We are getting closer to the meeting” treats the speaker as moving toward the meeting. The difference often goes unnoticed because both work.

Yet the difference matters. Moving Time makes events seem to arrive from outside. Moving Ego makes the person seem to advance through a temporal field. The first fits waiting and exposure; the second fits agency, planning, and progress. Each gives a different structure to temporal life.

The two models also reveal the ambiguity of the future. The future can be something that comes upon us and something toward which we move. A medical result comes. A decision is approached. A storm arrives. A child grows toward adulthood. A deadline descends. The metaphors preserve different modes of temporality.

Heidegger's account of future as *Zukunft* helps clarify the philosophical point. The future is not only a later now. It is the coming-toward-itself of *Dasein* in its possibilities. This structure cannot be reduced either to the movement of an event toward a fixed observer or to the movement of an observer along a path. Those are spatial articulations of temporal existence. They help speak the future. They do not ground it.

The difference appears in responsibility. A person moving toward a deadline is not merely traversing a line. They are obligated, delayed, anxious, perhaps unprepared. A child growing toward adulthood is being formed. A dying person moving toward death is living under the limit that gives life its finite weight.

Vector space can capture the language of approach, passage, and recession. It can learn that deadlines approach, years pass, lives progress, and memories recede. It can produce these phrases fluently. It cannot live the difference between being approached by an event and projecting oneself toward a possibility.

The model has the grammar of temporal movement. It has no temporal path of its own.

4. Time as object, path, and substance

Vyvyan Evans has argued that temporal language is more varied than a single TIME IS SPACE formula suggests. He identifies recurrent senses through which English and other languages speak time: duration, sequence, event relation, agency, commodity, matrix, and others.³

Temporal language is not one metaphor repeated mechanically. Time can be a path, a moving object, a container, a resource, a substance, a medium, or a field of agency. One spends time, wastes time, saves time, fills time, runs out of time, makes time, serves time, loses time, passes time, keeps time, kills time. Each phrase gives time another structure.

The commodity sense belongs strongly to modern social life. Time can be budgeted, spent, allocated, invested, saved, wasted, or lost. This language belongs to labor, administration, schooling, management, and productivity. It treats time as a scarce resource. It can be useful. It can also turn the finite temporality of life into economic inventory.

The container sense appears when something happens within a period, during a week, in childhood, inside a long silence. The matrix sense appears when time is treated as the medium in which all things occur. The agency sense appears when time heals, destroys, reveals, or betrays. Time becomes actor.

Each sense shows that spatialization is practical and social, not only geometric. The way a society speaks time reflects how it organizes work, memory, debt, leisure, discipline, and hope. “Saving time” makes sense where speed is a social good. “Killing time” makes sense where waiting has been separated from meaningful expectation. “Making time” makes sense where life is divided by competing demands.

This keeps cognitive linguistics close to historical critique. Metaphors are not only cognitive structures. They are historically sedimented habits. The modern subject often speaks of time as resource because modern institutions treat time as resource. Labor, school schedules, transport, accounting, digital systems, and platform metrics all teach that time can be owned, sold, lost, and optimized.

This has consequences for AI. Generative systems enter a world already governed by time-as-resource. They are valued because they save time, compress reading, accelerate drafting, reduce delay, and answer quickly. The interface presents temporal relief. It offers the user language before the time of formation has passed.

The spatial metaphor of time-as-resource therefore meets the technical operation of AI. The model spatializes linguistic traces. The user asks it to compress temporal labor. Output appears as saved time. The philosophical question is what kind of time has been saved, and what kind has been lost.

Some time saved is waste removed. Some time saved is formation avoided.

5. Cross-linguistic variation

Spatial temporal metaphors are widespread, but they are not identical across languages. This variation matters because it shows that the spatial articulation of time is real, learned, and culturally shaped.

The familiar English pattern places the future in front and the past behind. This fits bodily movement: one normally sees where one is going and no longer sees where one has been. It also fits modern ideas of progress. The future is ahead because one advances into it.

Other languages and communities organize temporal relations differently. Some use vertical axes, cardinal directions, past-in-front models, or earlier/later relations tied to writing direction, landscape, or cultural models of knowledge.

No language traps speakers inside one temporal picture. Speakers can learn multiple schemas. Many people use several models. A person can say that the future is ahead, that the deadline is approaching, that the past lies behind, and that one looks back on history. The same person may also read a timeline from left to right because of writing habits. Spatial temporal cognition is flexible.

Flexibility does not mean arbitrariness. Recurrent linguistic patterns shape attention. They teach which temporal relations are easy to say, visualize, embody, and make habitual. Spatial time is not universal in one fixed form. It is a family of embodied and cultural articulations.

This variation helps the book's argument in two ways. It strengthens the claim that language genuinely spatializes time across many communities, axes, and models. It also weakens the reductionist claim. If spatial temporal schemas vary historically and culturally, they are not the ontological ground of temporality. They are ways finite communities articulate time from their forms of life.

A vector model trained on many languages can learn such variation. It can compare English future-ahead expressions with Aymara past-in-front expressions or Mandarin vertical time expressions. It can classify patterns and generate explanations. That technical capacity is useful.

It has not entered any temporal world. It has learned the distribution of temporal expressions across texts, not time as the lived pressure of memory, visibility, orientation, and expectation.

6. The visible past

The Aymara case has become famous because it appears to reverse the familiar future-ahead model. Núñez and Sweetser argue that in Aymara, the past is often conceptualized as in front of the speaker, while the future is behind.⁴

The reasoning is grounded in visibility. The past is known. It has been seen. It lies before one's eyes. The future is unknown. It has not yet been seen. It lies behind, outside view. The spatial model is tied to an epistemic model: what is known is before us; what is unknown is behind us.

This example should be handled with care. Gesture, language, translation, context, and cultural interpretation all matter. The book does not need to turn Aymara into a symbol of radical difference. It needs the narrower point: temporal orientation can be spatialized through more than one bodily and epistemic structure.

The Aymara case also gives a useful contrast with Heidegger. For Heidegger, having-been, *Gewesenheit*, is not simply the known past placed before an observer. It is the inherited condition from which *Dasein* projects. It is not only visible. It is what one already is. The future is not only unseen. It is the field of possibility toward which one is projected. Heidegger's structure cannot be mapped onto the Aymara model without distortion.

Yet the Aymara model helps show why spatial metaphors are meaningful. The body faces what can be seen; the future cannot be seen. A spatial schema gives this epistemic relation a bodily form. Time is articulated through the body's orientation toward what appears and what remains hidden.

AI can state this. It can explain that the Aymara case reverses the English pattern and may depend on visibility and knowledge. Such a statement may be correct. It still belongs to a system with no visible past and no hidden future. It does not remember seeing. It does not stand before what has been. It does not wait for what cannot yet be seen.

The generated account can help a reader understand the case. The reader must still perform the understanding.

7. Vertical time

Mandarin Chinese has often been discussed through vertical temporal expressions. Earlier events can be associated with “up,” later events with “down,” as in expressions using *shàng* and *xià*. The phenomenon has been linked to writing practices, cultural models, and linguistic patterns.⁵

Popular accounts sometimes exaggerate the case. Mandarin also uses horizontal temporal expressions, and speakers are not locked into vertical time. The scholarly value lies in the coexistence of different spatial axes: front-back, left-right, up-down, and others.

Lera Boroditsky’s early experimental work argued that Mandarin speakers were quicker to think about time vertically after vertical primes than English speakers, and that language experience might influence temporal reasoning.⁶ Later work complicated the claim. Some effects depended on task design, writing direction, cultural practice, and bilingual experience. The debate remains useful because it prevents simple conclusions.

The safest philosophical claim is this: language, writing practice, and embodied spatial experience can influence habitual temporal representation. They do not prove that temporality is generated from spatial representation. They show that temporal thinking is mediated often by spatial schemas.

The vertical model has its own force. “Earlier is up” may be tied to lists, genealogies, calendars, writing layouts, or hierarchy. Up and down are never purely geometric. They carry value, order, visibility, gravity, authority, and sequence. A person rises in rank, falls from grace, descends into difficulty, or looks up to an ancestor. Vertical temporal language can carry social and moral associations.

Models are trained on such associations. They can learn that *shàng* and *xià* carry temporal meanings in Mandarin. They can translate or explain these patterns. They can also flatten them into a generic statement about “vertical time” unless the user returns to sources and usage.

The risk is not only error. It is the smoothing of difference. A generated explanation may turn a dense cultural and linguistic pattern into a clean concept. The user may feel they have understood because the pattern has been named. Real understanding requires return to examples, speakers, debates, and the limits of the claim.

Temporal metaphor becomes another site of unowned truth.

8. Cardinal time

The Pormpuraaw case, studied by Boroditsky and Gaby, concerns the Kuuk Thaayorre community in northern Australia. Research suggests that speakers may arrange temporal sequences from east to west, independent of the direction they are facing.⁷ Cards showing temporal order may be placed according to cardinal direction rather than egocentric left-to-right order.

The case belongs with the spatial-frame materials treated in Chapter 5. Some communities use absolute frames of reference more consistently than languages centered on left and right. If ordinary spatial language trains attention toward cardinal orientation, temporal order may also be mapped through that orientation.

Temporal sequence can be spatialized according to a geographic world rather than a body-centered axis or writing direction. Time is not merely before the face or behind the back. It can be arranged with sunrise, landscape, and shared orientation.

This gives another correction to a narrow reading of the origo. The speaker's locus is not only private I-here-now. It is also the shared world in which orientation has been learned. A temporal sequence can be placed in relation to the sun's movement, the land's organization, or the community's spatial discipline.

Such cases resist a hidden individualism. Locus is not just personal standpoint. It is inherited orientation. A person speaks from a language that has already placed them in relation to earth, sky, paths, ancestors, tools, calendars, and others. Spatial temporal metaphor is social and worldly before it is individual.

Again, AI can learn the pattern. It can identify the Pormpuraaw case as a cross-linguistic example of absolute spatial framing of time and compare it with English, Mandarin, or Aymara. The output is useful.

The system does not know east by turning. It does not know time by watching the sun cross a world. It has no bodily correction when orientation fails. The cardinal pattern exists for it as text and relation. It exists for speakers as practice.

9. Experiments and caution

The experimental literature on spatial temporal reasoning is valuable, but it must be read carefully. It has supported strong claims about language shaping thought, and it has faced replication questions, task-sensitivity, and dispute over interpretation.

Boroditsky's work helped show that language and spatial experience can influence temporal reasoning tasks. Casasanto's work developed the body-specificity hypothesis, arguing that people's bodily experience can shape mental representation, including how they associate valence, action, and time with spatial sides.⁸ Núñez, Cooperrider, and others widened the field through gesture, culture, and cross-linguistic cases.⁹

The lesson is not that experimental work has proved a total theory of time. It has not. The lesson is that temporal thought is not independent of bodily, linguistic, and cultural spatial practice. Even when claims are narrowed, that result remains important.

The caution matters because the book should not build its argument on fragile empirical overstatement. It does not need to say that language determines temporal thought, or that every English speaker imagines time horizontally, every Mandarin speaker vertically, or every Aymara speaker with the past in front. Such claims would be too simple.

The book needs a more restrained claim: temporal language and temporal reasoning often depend on spatial schemas, and those schemas are embodied, practical, cultural, and variable. Therefore, vectorization can capture genuine features of temporal language when it captures spatial relation. But this capture remains formal and derivative.

Experimental priming can show that spatial representations affect temporal judgments. It cannot show that temporality is nothing but spatial representation. A reaction-time effect does not replace phenomenology. A gesture pattern does not exhaust historicity. A diagrammatic arrangement does not explain finitude.

The human subject in the experiment is already temporal before the task begins. They arrive with memory, expectation, bodily habits, language, fatigue, compliance, and a relation to the task. The spatial-temporal mapping is tested inside this temporality. It does not float free from it.

This is the exact place where the later critique of vectorization must be careful. The model may learn patterns similar to those revealed in experiments. It may associate temporal concepts with spatial axes. It may represent temporal data geometrically. That would not refute the book's thesis. It would confirm that temporal language leaves spatial traces.

The trace is not the ground.

10. Embodied timing

The spatialization of time is bodily as well as linguistic. Human beings keep time with movement: walking, breathing, clapping, dancing, speaking, waiting, eating, sleeping, laboring, and aging. Rhythm precedes explicit temporal theory.

Research on timing and embodied cognition treats time perception through bodily movement, neural timing systems, sensorimotor coordination, and affective state.¹⁰ Duration changes with pain, fear, boredom, attention, movement, and desire. Waiting for bad news is not the same duration as reading in peace. Five minutes under threat do not pass like five minutes in conversation.

Heidegger's analysis of boredom gives a philosophical version of the same point. In boredom, time does not merely pass slowly. The person is held in limbo, *Hingehaltenheit*, and left empty, *Leergelassenheit*. Time becomes burdensome because the world's usual significance weakens. The clock measures the same minutes. The person's temporality is altered.

Temporal metaphor must therefore be returned to mood and care. A person says that time drags, flies, presses, opens, closes, wounds, or heals. These expressions are not arbitrary projections onto a neutral sequence. They arise because time is lived affectively. Mood changes the way time appears.

AI can generate such language. It can write that time dragged during the wait or flew during the celebration. It can produce delicate phrases about grief, memory, and anticipation. It can do so because human texts contain countless traces of such experiences.

The system does not wait. It does not dread. It does not age. It does not feel the day lengthen under grief or contract under absorption. Its temporal language is inherited from beings for whom duration matters.

This should not be sentimentalized. Human beings also use clichés about time without undergoing their meaning. A person can say "time heals" without thinking. Human language is full of dead metaphors and automatic phrases. The difference is not that humans always own temporal language. The difference is that humans can be recalled to the temporal experience from which the phrase draws its force. The model cannot.

A person can discover that “time heals” was false for them, or true in a way they did not expect. A model can generate both claims.

11. Phenomenological temporality and spatial articulation

Phenomenology does not deny the spatial articulation of time. Husserl’s retention and protention already require a thickness of the living present that can be diagrammed, even if no diagram captures it. Heidegger’s ekstases can be spatially imaged as future, having-been, and present standing outside themselves. Ricoeur shows that narrative configures time into a followable form. Each account uses relations that can be drawn, ordered, or spatially represented.

The question is not whether time can be spatially articulated. It can. The question is whether spatial articulation is ontologically first. Phenomenology says no.

Husserl’s melody cannot be reduced to a sequence of points, even though musical notation places notes on a spatial staff. The notation is indispensable for certain practices. It is not the hearing of the melody. Retention and protention make the melody possible as heard duration. The score preserves and organizes. It does not live the sound.

Heidegger’s *Zeitlichkeit* cannot be reduced to a timeline, even though one can draw future, having-been, and present in a diagram. The ekstastic unity of temporality is not an object arranged before a viewer. It is the structure by which *Dasein* is ahead of itself, already thrown, and making-present beings within concern.

Ricoeur’s narrative configuration cannot be reduced to a plot diagram, even though plot diagrams can help. The narrative arc has meaning because it mediates between a pre-narrative world of action and a reader’s refigured world. A plot schema without *mimesis1* and *mimesis3* is a formal pattern without the hermeneutic arc.

These distinctions matter because vectorization is a spatial formalization. It turns traces into relations. It can capture the score-like, diagram-like, plot-like, and pattern-like dimensions of language. It can produce language that appears temporally ordered because it has learned such order from texts.

It cannot become the hearing of the melody, the projection of *Dasein*, or the refiguration of a reader. It can preserve formal temporal articulation. It cannot originate temporal existence.

This is the chapter's bridge to the book's central argument. Spatial articulation is not false. It is one of the ways temporal life becomes communicable. But communicability requires what the formalization does not contain: a temporal locus from which the articulation is made and to which it can return.

12. The vector and the metaphor

The connection between TIME IS SPACE and vectorization can now be stated directly.

A vector space represents relations spatially. Tokens, words, phrases, and contexts are placed in high-dimensional relation. Similarity becomes proximity. Difference becomes distance or direction. Semantic movement becomes transformation. Analogy becomes geometric relation. The model's architecture turns linguistic meaning into operable spatial form.

Temporal language gives this architecture rich material. The corpus contains countless spatial expressions of time: approaching deadlines, future plans, long waits, short periods, stages, cycles, returns, delays, progress, decline, beginnings, endings. The model learns these patterns as part of its training. It can generate temporal language because temporal language has already been spatialized in human texts.

Temporal fluency therefore appears natural. The model does not need originary temporality to say that something lies ahead. It needs enough learned relation among tokens, contexts, genres, and tasks. The phrase “the work ahead” is a pattern. “Looking back on the past” is a pattern. “The deadline approaches” is a pattern. These patterns are meaningful for the user because they come from human temporal life.

The model’s power rests on inherited metaphor.

This claim should not be narrowed to metaphor alone. Temporal language includes tense, aspect, deixis, narrative, discourse mode, evidentiality, modality, and date calculation. Yet metaphor is especially revealing because it shows the passage from bodily temporality into spatial expression. The model receives the expression after the passage has occurred.

The result is the spatial shell of time. The shell contains direction, distance, order, sequence, and relation. It can be manipulated. It can be recombined. It can even support useful reasoning in many cases. But it lacks the inner temporal structure of the being for whom the future is possibility, the past is having-been, and the present is concern.

The metaphor that helps a human speak time becomes, in vectorization, a technical affordance for generating time-like language. The danger is that generated language returns as if the temporal ground had returned with it.

It has not.

13. The limit of metaphor

A metaphor can disclose. It can also conceal. TIME IS SPACE helps human beings speak time. It also tempts them to forget that time is not a thing laid out before them.

The future is not literally ahead. The past is not literally behind. A year is not literally long. A deadline does not literally walk toward us. A life does not literally move along a visible road. These expressions let embodied beings articulate temporality through spatial relations. They fail when taken as ontology.

The same limit applies to computational representation. A model may represent temporal relations in a high-dimensional space. It may learn vectors that correlate with dates, durations, sequences, or historical periods. It may even contain structures that allow probes to recover spatial or temporal information. Such findings matter technically. They do not show that the model has temporality in the existential sense.

They show that temporal traces can be geometrized.

The category mistake is not that vectorization uses space. Human beings also use space to think and speak time. The mistake is treating spatial calculability as if it could ground the temporality it formalizes. Spatial metaphor is a human way of bearing time. Vector space is a technical way of operating on the traces of that bearing.

This distinction will matter in later chapters on embeddings and temporal reasoning. When a model represents temporal order, that representation is not nothing. It may support useful output. When the model fails at temporal reasoning, the failure is not only an engineering gap. It reveals the difference between operating on spatialized temporal traces and inhabiting time as finite historical existence.

The human reader must therefore practice a double discipline. They may use the model's spatial map of temporal language. They must not receive the map as time itself. They must return the generated temporal expression to the source, the event, the date, the practice, the text, and the responsibility that give it force.

This is already the beginning of locus-reinjection. The user hears the temporal language and asks where its time is: whose past is being named, which future is being projected, which present is being assumed, which calendar, archive, body, institution, or tradition has been flattened into fluent availability.

The model can say that something is ahead. The user must ask: ahead for whom.

14. The spatial shell

Part II has now established the linguistic side of the book's first move.

Language carries locus through deixis: I, here, now. It structures relations through spatial grammar: figure, ground, path, support, containment, force, and orientation. It formalizes temporal reference through tense, aspect, Speech Time, Reference Time, Event Time, and Topic Time. It articulates time through spatial metaphors of movement, distance, path, container, resource, and visibility.

Vectorization can work for this reason. It finds a language already full of spatial relation. It finds temporal language already spatially articulated. It finds grammar, metaphor, discourse, and narrative already carrying the traces of human spatial-temporal life. The model does not invent this structure. It inherits it.

The limit is equally clear. Deixis points to an origo. Spatial grammar arises from embodied and practical orientation. Tense and aspect presuppose a deictic now. Temporal metaphor arises from bodily and historical life. Vectorization can capture the spatial shell of these structures. It cannot supply the temporal locus that made them meaningful.

The next part turns to writing, the first major detachment of language from immediate speech. Writing preserves the formal markers of locus while separating the utterance from the living situation in which it first occurred. That detachment is not yet the loss produced by vectorization. A written text can still be read between times. Its absent speaker can still be approached through interpretation. Its source locus can still be sought.

Writing detaches. Reading can return. Vectorization will detach differently.

NOTES

1. George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago: University of Chicago Press, 1980), 3. The present chapter uses conceptual metaphor theory in a limited way. It accepts the claim that ordinary thought and language are deeply structured by metaphor. It does not accept the stronger reduction that would make lived temporality a mere effect of spatial schema. ↩

2. George Lakoff and Mark Johnson, *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought* (New York: Basic Books, 1999), 139 to 169. See also Rafael E. Núñez and Eve Sweetser, “With the Future Behind Them: Convergent Evidence from Aymara Language and Gesture in the Crosslinguistic Comparison of Spatial Construals of Time,” *Cognitive Science* 30, no. 3 (2006): 401 to 450. Moving Time and Moving Ego are useful because they show that even within one language temporal motion can be construed from more than one spatial standpoint. ←
3. Vyvyan Evans, *The Structure of Time: Language, Meaning and Temporal Cognition* (Amsterdam: John Benjamins, 2004); Vyvyan Evans, “How We Conceptualize Time: Language, Meaning and Temporal Cognition,” *Essays in Arts and Sciences* 33 (2004): 13 to 44. Evans is valuable because he resists treating TIME IS SPACE as one simple mapping. Temporal language includes several distinct senses and frames. ←
4. Núñez and Sweetser, “With the Future Behind Them,” 401 to 450. The Aymara case should be treated as evidence for a possible relation among visibility, knowledge, and temporal orientation, not as a simple reversal slogan. Gesture, translation, and context all matter. ←
5. On Mandarin vertical time, see Lera Boroditsky, “Does Language Shape Thought? Mandarin and English Speakers’ Conceptions of Time,” *Cognitive Psychology* 43, no. 1 (2001): 1 to 22. For later qualifications and debate, see Chen-Yu Tsai and Dedre Gentner, “Structuring Time with Space: Evidence from Mandarin and English,” *Proceedings of the Annual Meeting of the Cognitive Science Society* 28 (2006): 814 to 819; Kensy Cooperrider and Rafael Núñez, “Across Time, Across the Body: Transversal Temporal Gestures,” *Gesture* 9, no. 2 (2009): 181 to 204. ←
6. Boroditsky, “Does Language Shape Thought?” 1 to 22. The early experimental claim has been widely discussed, extended, and narrowed. The present chapter relies only on the careful version: linguistic and cultural spatial practices can shape habitual temporal representation under certain tasks. ←
7. Lera Boroditsky and Alice Gaby, “Remembrances of Times East: Absolute Spatial Representations of Time in an Australian Aboriginal Community,” *Psychological Science* 21, no. 11 (2010): 1635 to 1639. This case belongs with the broader literature on absolute frames of reference, especially Stephen C. Levinson, *Space in Language and Cognition: Explorations in Cognitive Diversity* (Cambridge: Cambridge University Press, 2003). ←
8. Daniel Casasanto, “Embodiment of Abstract Concepts: Good and Bad in Right- and Left-Handers,” *Journal of Experimental Psychology: General* 138, no. 3 (2009): 351 to 367; Daniel Casasanto and Lera Boroditsky, “Time in the Mind: Using Space to Think About Time,” *Cognition* 106, no. 2 (2008): 579 to 593. Casasanto’s work is useful because it shows that bodily asymmetries and sensorimotor patterns can affect abstract mappings. The chapter does not use this as proof of a full theory of temporal experience. ←
9. Rafael E. Núñez and Kensy Cooperrider, “The Tangle of Space and Time in Human Cognition,” *Trends in Cognitive Sciences* 17, no. 5 (2013): 220 to 229. Núñez and Cooperrider offer a balanced account of the relation between spatial and temporal

cognition, including the need to distinguish linguistic, gestural, cognitive, and cultural evidence. ↩

10. On embodied timing and neural timing, see Warren H. Meck, "Neuropharmacology of Timing and Time Perception," *Cognitive Brain Research* 3, no. 3 to 4 (1996): 227 to 242; Matthew S. Matell and Warren H. Meck, "Cortico-striatal Circuits and Interval Timing: Coincidence Detection of Oscillatory Processes," *Cognitive Brain Research* 21, no. 2 (2004): 139 to 170. The neuroscientific literature is not used here to replace phenomenology. It supports the narrower point that temporal experience is bodily and affective, not only formal sequence. ↩

Chapter 8. Writing and the First Detachment

1. The first absence

The preceding chapters established the book's first move: language carries locus. It speaks from an origo, structures relation through spatial grammar, marks temporal location through tense and aspect, and speaks time through space without reducing time to space.

Writing changes this structure.

A spoken utterance occurs in a situation. It has a speaker, addressee, place, time, tone, and possible reply. The situation need not be simple. The speaker may be evasive, the addressee may misunderstand, and the place may matter more than the sentence. But speech belongs to an event in which language is uttered from a living locus.

Writing detaches language from that event. It preserves the words while removing them from the immediate situation of saying. The speaker no longer stands before the hearer. Voice, gesture, reply, and occasion recede. The marks remain.

This is the first detachment.

The detachment is not a defect to be lamented. Without it there would be no archive, law, philosophy, scripture, long-distance promise, letter from the dead, historical science, or book. Writing makes language durable. It lets an utterance cross the limits of voice, body, place, and immediate time. It lets one temporal locus address another.

The detachment is also dangerous. The written text can circulate without the speaker. It can be quoted without its question, copied without its occasion, read without its language, processed without its world. The marks survive, but survival is not yet understanding.

This chapter defends a precise claim. Writing detaches language from immediate speech while preserving formal markers of locus. The detachment is partial because the written text can still be returned to a temporal and historical locus through reading. The reader can ask who wrote, when, where, to whom, and under what pressure. The reader can read between the times.

Vectorization will later change this relation. It will preserve linguistic markers while severing the structural possibility of returning them to an original locus. Writing stretches the link between marker and locus. Vectorization breaks it.

This chapter stays with writing. It treats writing as the first technical exteriorization of language, the first absence that still permits return.

2. Theuth's gift

The classical scene of writing is Plato's *Phaedrus*. Socrates tells the myth of Theuth, the Egyptian god who invents number, calculation, geometry, astronomy, dice, and letters. Theuth presents writing to King Thamus as a gift. He calls it a *pharmakon* for memory and wisdom: “μνήμης τε γὰρ καὶ σοφίας φάρμακον.” The Greek word *pharmakon* is deliberately unstable: remedy, drug, poison, charm, or recipe.¹

Thamus refuses the claim. The inventor of a technic, he says, is not always the best judge of its effects. Writing will not produce memory, *mnēmē*. It will produce reminding, *hypomnēsis*. The written mark lets learners rely on external signs rather than on the inner exercise of recollection. It gives them “the appearance of wisdom,” not wisdom itself.

Socrates then gives the second criticism. Written words resemble living beings, but if questioned they remain silent. They say the same thing to everyone. They do not know to whom they should speak and to whom they should not. If attacked or misunderstood, they cannot defend themselves. They need their father, the living speaker, to come to their aid.

The Platonic suspicion is not silly nostalgia for speech. It sees a real structural change. Writing detaches the logos from the living situation in which it could answer. The written text is repeatable. It can go where the speaker is not. It can meet readers for whom it was never prepared. It can be used without the living discipline of dialectic.

Yet Plato’s own dialogue is written. The critique of writing is preserved by writing. The danger is transmitted through the very medium that is accused. Writing is already both suspect and necessary.

The *Phaedrus* therefore gives the chapter its first law: writing is neither pure loss nor pure gain. It is *pharmakon*. It preserves and endangers, opens transmission and invites false possession, supports memory and weakens memory, carries a voice and removes the speaker.

The law still holds for artificial intelligence. AI output often appears as a new remedy for memory and wisdom. It gives the user access to stored language, summary, comparison, explanation, and answer. It can help. It can also produce the old Platonic danger in a more intense form: the semblance of wisdom without the undergoing that would make wisdom one’s own.

Writing first made this danger possible. AI industrializes it later. The distinction must not be blurred.

3. Anamnesis and hypomnesis

The *Phaedrus* distinguishes *anamnesis* from *hypomnesis*. *Anamnesis* is recollection, the recovery of what must be drawn from within the soul. *Hypomnesis* is reminding, an external aid that prompts memory but does not replace it. Writing belongs to *hypomnesis*.

The contrast is severe in Plato, but it should not be read too quickly. A reminder is not nothing. A written mark can call back a promise, preserve a law, teach a student, mark a debt, protect a name, or hold a word across death. The danger lies in receiving the reminder as if it were recollection itself.

The written text can make one appear learned. A person can cite without understanding, copy without recollection, recite without judgment. The words are present. The inward act is missing; the ancient form of unowned truth.

The same structure reappears in later media. A manual can support practice without producing skill. A database can store knowledge while substituting retrieval for understanding. A generated summary can orient a reader while replacing the reading through which the text would have become formative.

The Platonic distinction therefore becomes a discipline of reception. External memory is necessary. Human beings are not pure inwardness. They live through marks, books, tools, rituals, archives, and institutions. But every external aid must be returned to the living act it can only support. A reminder becomes good when it awakens memory and dangerous when it replaces it.

The same distinction will later govern the use of AI. A generated answer can be *hypomnesis*, a prompt for return. It can remind, orient, gather, and prepare. It becomes dangerous when it is received as *anamnesis*, as if the user had already remembered, understood, or judged.

Writing begins the history of this ambiguity.

4. Aristotle and Augustine

Aristotle gives writing a lower place in the order of signs. At the beginning of *De Interpretatione*, he writes that spoken sounds are symbols of affections in the soul, and written marks are symbols of spoken sounds. The Greek formulation is famous: “ἔστι μὲν τὰ ἐν τῇ φωνῇ τῶν ἐν τῇ ψυχῇ παθημάτων σύμβολα, καὶ τὰ γραφόμενα τῶν ἐν τῇ φωνῇ.” Spoken sounds are symbols of affections in the soul, and written things are symbols of spoken sounds.²

This establishes a hierarchy. Things affect the soul. Speech expresses those affections. Writing records speech. Writing is twice removed from the encounter with things. It is derivative.

The hierarchy shaped much later thinking about language. Writing became a secondary notation, a visible substitute for voice. Speech appeared closer to presence, intention, and truth. Writing appeared conventional, external, and supplementary.

Augustine inherits this tradition and changes it. He also treats words as signs, and he recognizes the passing nature of speech. Spoken words vanish as soon as they strike the air. Letters are invented as signs of words, so that what passes in sound can endure in visible form.³

Yet Augustine’s account is not a simple degradation of writing. In Christian reading, the written sign becomes the site of exegesis. Scripture is not living speech before the eyes, but it is not inert. It must be interpreted through grammar, history, translation, ambiguity, and charity. Written signs do not deliver truth automatically. They call the reader into disciplined interpretation.

In *De Magistro*, Augustine also denies that signs by themselves teach. A sign can direct attention, but the learner must turn toward what is signified. Words do not replace seeing the matter. They prompt the movement toward it. Written language is secondary, but its secondary status is productive. It forms a field in which the reader must move from sign to thing, from letter to sense, from obscurity to understanding.

Writing is therefore not only storage. It creates a task. It preserves language by placing the reader under the demand of interpretation. The reader must ask what is meant, what is obscure, what has been lost, what must be supplied, and what kind of life is required to understand.

The written sign therefore opens distance, and distance opens responsibility.

5. Orality and literacy without myth

Modern studies of orality and literacy extend the ancient problem. Havelock, Ong, and Goody all argue, in different ways, that writing transforms memory, thought, and social organization. Oral cultures rely heavily on formula, rhythm, repetition, performance, and living presence. Writing permits storage, revision, abstraction, classification, and the separation of words from the speaker's immediate situation.⁴

Such claims must be handled carefully. The strong "Great Divide" thesis has been rightly criticized. Writing does not automatically create rationality. Orality does not mean primitive thought. Complex oral societies have memory, law, poetry, history, and subtle reasoning. Literacy practices are embedded in social, political, pedagogical, and institutional settings. Brian Street's critique of the autonomous model of literacy remains important here.⁵

The book needs only a structural distinction.

Writing changes the affordances of language. It makes speech visible and durable. It allows a text to be inspected, copied, edited, quoted, classified, preserved, and transmitted. It lets language become an object of attention in a way oral speech usually does not. It supports grammatical reflection, commentary, law, scripture, account-keeping, and formal proof.

These affordances do not determine their use. A written culture can be dogmatic. An oral culture can be reflective. A printed society can be forgetful. A manuscript culture can be subtle. But writing permits a new relation between speech, memory, and time. It makes language endure beyond the event of utterance.

This endurance changes temporality. The spoken sentence disappears unless remembered, repeated, or recorded. The written sentence remains. It can wait. It can be returned to. It can meet readers not yet born. It can carry the formal markers of its original situation into another time.

The written “I” is no longer accompanied by the speaker’s body. The written “here” is no longer the reader’s place. The written “now” is no longer the reader’s now. Yet these words are not meaningless. They point back through the text toward a locus that can be reconstructed. Writing preserves the marker while deferring the locus.

That is the first detachment in its linguistic form.

6. Exteriorized memory

Husserl’s “Origin of Geometry” gives writing a stronger role than Plato allowed. Husserl asks how ideal objects, such as geometrical truths, can become historical. A geometrical insight begins in an act of evidence. It occurs in a person. Yet geometry is not private. It can be transmitted, repeated, corrected, and extended by others across generations.

Speech is not enough. It remains bound to living communication and immediate memory. Writing gives the ideal object a durable body. Husserl calls this the condition by which ideal objectivity can become available beyond the original act. A theorem can be written, preserved, forgotten, rediscovered, and reactivated.⁶

This is not only storage. It is the constitution of a tradition. Without inscription, the ideal object would not acquire the historical durability required for science. Writing allows a sense to sediment. Sedimentation is necessary because it preserves. It is dangerous because it can be repeated without reactivation.

Husserl’s word “reactivation” is decisive. The later reader must not merely repeat the formula. They must return through the written trace to the evidence that gave it sense. The written mark holds the possibility of such return. It does not perform the return by itself.

This gives the present book a necessary correction to Plato. Writing is not merely an aid that weakens memory. It is a condition of certain forms of truth. Geometry, law, theology, philosophy, history, and science all require inscription. Writing makes possible kinds of objectivity and transmission that speech alone could not sustain.

Yet Husserl also confirms the chapter's boundary. The written sign must be reactivated. The trace can become empty if the reader receives it only as formula. Writing preserves the possibility of truth across time, but it does not remove the need for a temporal subject who reads.

A theorem in a book is not yet geometrical understanding. A philosophical passage is not yet interpretation. A source in a database is not yet historical knowledge. A generated summary of that source is even further from the reactivation Husserl requires.

Writing can carry ideality because a reader can return to it. The corpus can contain the writing. Vectorization can model its patterns. Only a reader can reactivate the sense.

7. Tertiary retention

Stiegler gives this structure its full technical weight. He names exteriorized memory *tertiary retention*. Husserl had distinguished primary retention, the just-past held in the living present, from secondary retention, recollection. Stiegler adds a third: memory preserved outside consciousness in material supports. A written text, a tool, a photograph, a recording, a calendar, and a database can all function as tertiary retention.⁷

This changes the anthropology of memory. Human beings do not first have pure interior memory and then add tools. Human temporality is technical from the beginning. We inherit a world already filled with marks, artifacts, practices, stories, and instruments. These exteriorized supports shape what can be remembered, anticipated, and desired.

Writing is one of the major forms of tertiary retention. It preserves linguistic sequence outside the speaker. It makes memory shareable, revisable, institutional, and transmissible. It also changes the relation between the person and their own memory. A note can relieve memory or train it. A diary can preserve a life and make it readable to itself. A law can stabilize obligation and harden into formula.

Stiegler's concept prevents a false purity. The book must not oppose living time to technical memory as if the technical were only corruption. Human historical life depends on exteriorization. Tradition depends on supports. The question is not whether memory should be technical. The question is what kind of technical memory is at work, and what relation the human user takes toward it.

The same support can form or deform. A book can deepen attention or become a badge of possession. A recording can preserve a voice or replace listening. A search index can help research or make return feel unnecessary. A model can orient the reader or offer language in place of the reader's own temporal work.

Writing is therefore already pharmacological in Stiegler's sense. It is remedy and poison. It forms long circuits of memory when it supports attention, reading, and transmission. It short-circuits memory when it substitutes external trace for internal and communal formation.

This pharmacology begins with writing. It does not end there.

8. Derrida and iterability

Derrida's reading of the *Phaedrus* in "Plato's Pharmacy" refuses the simple hierarchy in which speech is living presence and writing is dead supplement. The *pharmakon* cannot be fixed as either remedy or poison. Writing is both. It reveals that the living word was never as self-present as the hierarchy claimed.⁸

Derrida's broader concept of iterability sharpens the point. A sign can function only because it can be repeated in the absence of its original speaker, addressee, and context. If a mark could not outlive the present intention of its maker, it would not be a mark in the strong sense. Absence is internal to signification.

This matters for the present argument, but it must be used with care. Derrida prevents nostalgia for pure speech. There is no simple living origin that writing later contaminates. Speech itself already depends on repeatable forms. The word I speak can be cited, misunderstood, repeated, detached, and used against me. The trace is not an accident that arrives after presence. It belongs to language.

Yet this does not mean all detachments are the same. A written text's break with origin does not abolish historical reading. The fact that a mark can be repeated without its author does not mean that the authorial locus is irrelevant. It means that the locus is not simply present in the mark. It must be approached through interpretation, and the approach can fail.

The fact that every sign can function in absence does not make all absences equivalent. The absence of the author in a written text is one structure. The disappearance of source genealogy in generated output is another. Writing detaches language from the first event of speech, but the written trace may still be returned to author, date, genre, institution, and world. Vectorization can disperse the trace so that no such return is available in the same way.

Derrida therefore helps the chapter avoid two errors.

The first error is Platonic nostalgia: writing as a fall from living presence. The second is technological indifference: because all signs are iterable, the move from writing to vectorization would be only more iteration. That second error is just as serious. Writing can detach from the speaker and still preserve a readable trace. Vectorization detaches the trace into mathematical relation and returns new output without preserving a path to original locus.

Derrida shows that writing always involves absence. The present book adds that not every absence has the same structure. The absence of the author in a written text is not the same as the absence of a speakerly locus in generated output.

The written trace can still call for reading. The vectorized trace returns as operation.

9. The written deictic

The previous chapters treated deixis, tense, aspect, and temporal metaphor. Writing alters each.

A spoken “I am here now” is anchored in a speaker, a place, and a time. A written “I am here now” preserves the words after the speaker has departed. The I is absent. The here has moved. The now has passed. The sentence survives.

This survival is strange. The sentence’s markers are no longer fulfilled by immediate presence, but they are not empty. They point historically. The reader can ask who wrote the sentence, from where, when, to whom, and why. The written “I” can be returned to the author, the written “here” to the place of inscription or argument, and the written “now” to the date, scene, textual present, or historical moment.

Writing therefore preserves deictic form under deferred anchoring. The link between marker and locus is stretched across time. It is not severed.

This is the structural point on which the later distinction between writing and vectorization will depend. In writing, the marker remains attached to a possible locus through the history of inscription and interpretation. The reader may fail to reconstruct it. Sources may be missing. The text may be anonymous. The passage may be corrupt. The original situation may be permanently uncertain. Yet the task still makes sense. The text is received as a trace from somewhere.

The written deictic also creates layered loci. In a novel, “I” may refer to a narrator rather than the author. In a philosophical text, “here” may refer to a place in the argument. In a diary, “today” may refer to the day of inscription. In a legal text, “now” may refer to enactment, performance, or application. Writing multiplies the work of reading because locus becomes textual, historical, generic, and institutional at once.

This complexity strengthens the argument. Writing does not preserve locus by simple transparency. It preserves it as a hermeneutic demand. The reader must interpret the kind of locus at stake.

Vectorization will later preserve deictic markers without this demand appearing at the interface. The model can generate “I,” “here,” and “now” as fluent forms. The user must supply or restore the missing locus. The written text asks to be read. The generated output often appears already to have answered.

10. Distanciation

Ricoeur gives the positive structure of written distance. When discourse is fixed in writing, it becomes autonomous from the immediate speech event. It separates from the speaker’s intention, the original audience, and the ostensive reference of the situation. This autonomy is not mere loss. It lets the text project a world beyond its origin.⁹

Ricoeur calls this distanciation. The word should not be heard as alienation alone. Distance makes interpretation possible. If speech remained bound to its immediate situation, it would not become a text capable of addressing later readers. Writing frees discourse from the event of speaking so that it can enter new contexts.

This freedom has two sides. First, the text no longer depends entirely on recovering the author’s psychology. The author’s intention matters, but it is not sovereign. Once fixed, the text can mean more than the author knew. It can project a world that future readers enter differently.

Second, the reader’s appropriation is not arbitrary. The text has structure, genre, wording, reference, and internal constraint. It does not become whatever the reader wants. The reader appropriates the world of the text by letting it reshape the reader’s own world.

Writing detaches, but the detachment opens a hermeneutic relation. The reader does not simply restore the original speech event. They read from their own temporal locus toward the text’s locus, and the text projects possibilities back into the reader’s present.

This is one form of reading between the times. The reader stands in a present horizon. The text comes from a past or distant horizon. Understanding happens neither by erasing the reader's present nor by reducing the text to present use. It happens in the disciplined relation between them.

AI output can produce language about this relation. It can summarize Ricoeur, identify distanciation, describe the world of the text, and propose interpretations. It may help a reader begin. But it does not inhabit the distance between text and reader. It has no temporal locus that can be altered by the world of the text.

The distance belongs to the reader.

11. Temporal distance

Gadamer gives the same problem another formulation. The Enlightenment often treated temporal distance as an obstacle. The historical method then tried to bridge the gap by reconstructing the past as it was. Gadamer reverses the assumption. Temporal distance is not only a gap. It is the medium in which understanding occurs.¹⁰

The reader is always historically affected, *wirkungsgeschichtlich*. The text has already had effects. The reader has already inherited prejudgments, language, institutions, questions, and expectations. Interpretation begins within this history of effects, not outside it.

This means that the reader cannot remove their own locus. They can only bring it into relation with the text. A prejudice may distort, but it can also open the question. A tradition may narrow vision, but it can also make the text speak. The task is not to become neutral. The task is to let the matter of the text test the reader's starting point.

Writing depends on this temporal distance. A text can endure because it can be read by those who were not present. Its meaning is not exhausted by the first audience. The text remains capable of address because every new reading occurs from a new historical locus.

This capacity is not the same as infinite drift. A historical text is not raw material for any use. It comes from a world, a language, a genre, and a set of questions. It resists the reader. The resistance is part of its truth.

Here the first book's distinction should be kept but refined. Reading re-historicizes writing. Processing de-historicizes it. Reading lets distance work. Processing tends to convert distance into relation, frequency, similarity, and proximity. Both can produce knowledge-like results. Only reading can be addressed by the text as a historical other.

Gadamer helps explain why the written trace remains alive. It is not alive because the author's presence is magically stored in it. It is alive because temporal distance allows it to enter new events of understanding.

The text survives by being read again.

12. The ethics of writing

Carlo Sini gives the Italian dossier a place in this chapter. His work on writing treats the written mark not as a neutral vehicle but as a practice that shapes the subject who uses it. Writing is an ethical event because it changes how human beings stand in relation to word, silence, gesture, and world.¹¹

Writing is often treated as a technical support that follows thought. One thinks, then writes. The written mark preserves what was already formed. Sini's line of thought resists this simplification. Writing is not only the recording of a prior inner content. It is a practice in which thought takes form, becomes visible, becomes revisable, and becomes answerable.

The written page disciplines the writer. It shows repetition, weakness, contradiction, evasion, and excess. It allows return. The sentence stands there and can be corrected. The writer becomes reader of their own saying. The mark externalizes thought and exposes it.

This is one reason writing has a special relation to responsibility. Speech can be denied, forgotten, or excused as passing. Writing remains. The author may revise, but revision also leaves a history. The written mark can be cited. It can be returned to the one who wrote. It gives language a body that can be held.

That permanence is not pure virtue. It can harden the word, fix what should have remained open, turn living interpretation into procedure, and become idol, archive, warrant, file, or weapon. The ethics of writing lies in this tension: the mark gives accountability, but can also detach it from the living situation in which it should be exercised.

This tension matters for artificial intelligence because generated text enters the world as writing. It can be copied, signed, submitted, circulated, and institutionalized. Yet its authorial status is unstable. The model produced it. The user prompted it. The system owners shaped it. The institution may adopt it. Responsibility can blur.

The chapter is not yet about AI accountability. It prepares the ground. Writing makes language durable and attributable. AI-generated writing can make language durable while obscuring attribution. The ethical practice later called locus-reinjection will have to restore who wrote, who generated, who used, who signed, and who answers.

13. Reading

Writing's detachment can be answered only by reading. Reading is not decoding alone. It is the temporal act by which a written trace is taken up from another locus and returned to meaning.

A reader does several things at once. They follow the line, remember what has been said, anticipate what may come, hold grammar, genre, context, source, and question together, pause, return, compare, distrust, accept, revise, and let the text alter what they had expected. Reading is therefore a temporal practice of retention, protention, and judgment.

A historical reader does more. They let the text come from elsewhere. They ask after source, language, date, genre, audience, institution, and world. They do not erase the text's distance in order to make it usable. They preserve distance so that the text can resist them.

This is why a summary is not reading. A summary may help reading. It may give orientation, structure, or memory. It may help a reader return to a difficult source. But the summary is not the passage through which the reader meets the text. It gives language about the text, not the encounter with the text.

The same applies to translation, commentary, paraphrase, and index. Each can assist reading. Each can also replace it. The difference lies in the reader's return. A commentary is good when it sends the reader back more alert. It is bad when it becomes a substitute possession. A generated summary is useful when it opens the path. It is dangerous when it closes it.

Reading is the practice that keeps writing from becoming dead material. It reactivates the trace. It returns the mark to a temporal relation. It allows "summoning the ghosts of true experiences," not by pretending that the dead are present, but by letting their traces speak from their own distance.

The phrase matters. Writing summons ghosts. The ghost is not the author resurrected. It is the trace of a saying that once belonged to a life, a question, a practice, a danger, a joy, a wound, an argument. Reading does not bring that life back. It lets the trace address the present without consuming its otherness.

The ghost in writing still has a grave, a date, a language, a place, a hand. That is why it can be honored.

14. Processing

Processing is another relation to writing. It can search, classify, align, count, compare, cluster, extract, summarize, translate, and model. These operations can serve scholarship and ordinary reading, reveal patterns no human reader could find, and help with scale, access, and orientation.

Processing is not reading.

The difference is not sentiment. It lies in the relation to temporal locus. Reading receives the written trace as an address from elsewhere. Processing treats the trace as material for operation. Reading asks how the text's world meets the reader's world. Processing asks how textual elements relate to other elements under a chosen procedure.

Both questions can be legitimate. A scholar may need a concordance, a historian a database, a philologist alignment, a critic distant reading, a reader search. The danger appears when the operational relation replaces the hermeneutic relation.

A processed text can become rich in relations and poor in address. It can be searchable but unread, indexed but unencountered, summarized but not understood, placed in a corpus while its source world recedes.

The first book named this distinction directly: historical reading preserves distance, while computational processing converts distance into proximity. A word becomes close to other words, a passage similar to other passages, a style a feature, a concept a cluster, a tradition a distribution. The resulting relations may be valuable, but they are not the same as the text's address.

This chapter modifies the prior claim. Processing does not always de-historicize in use. It can support re-historicization when directed by a reader. The operation itself, however, does not read. It does not receive the text as a temporal other. It does not undergo the exposure by which the text may change a reader's horizon.

Writing preserves absence as a demand for reading. Vectorization may preserve traces without preserving a readable genealogy. The reader of a written text asks where the text comes from. The user of generated output may have to ask a harder question: what source loci have been rendered unrecoverable here, and what responsible path can still be built?

The machine can prepare relation. The reader must restore address.

15. From writing to corpus

The next chapter will treat the corpus directly. Here the transition should be marked without being completed.

Writing creates durable traces. Archives gather them. Libraries preserve them. Editions stabilize them. Catalogues classify them. Databases index them. Corpora prepare them for computational analysis. Each step changes the written trace.

The archive preserves a possible text-reader relation. A book in a library can be unread for decades and still await a reader. A letter in an archive can still be returned to sender, recipient, date, paper, hand, and occasion. A manuscript can still be approached as a historical object.

The corpus is different. It gathers texts as processable material. It does not necessarily destroy provenance, but provenance no longer governs the relation. The text enters a field of selection, cleaning, filtering, segmentation, counting, and comparison. A letter, poem, legal brief, transcript, review, and code file can become units in the same computational field.

The difference is not moral. It is modal. The text's mode of being changes. As writing, it is a trace awaiting possible reading. As corpus material, it is writing prepared for operation. It may still be readable if preserved and accessible. But the primary relation has changed.

This is why writing and corpus require separate chapters. Writing detaches language from speech while preserving a readable trace. Corpus formation gathers traces into an aggregate technical memory. Vectorization then converts those traces into mathematical relations. These are related steps, not the same step.

The present chapter establishes the first. The written text preserves formal indexical markers after the immediate speech situation has vanished. It permits historical reconstruction and reactivation. It asks for reading between the times.

The corpus will introduce another relation: simultaneous availability.

16. The first detachment and the second

The book's central three-move structure now has its second move in place.

Move one: language is existentially spatial. It carries indices of person, place, time, force, path, support, containment, tense, aspect, and temporal orientation. These articulate a speaker's temporal-existential locus.

Move two: writing detaches language from the immediate speech situation while preserving the formal markers of that locus. The detachment is partial. The written trace can still be returned to source through interpretation. Reading can reconstruct, reactivate, and be addressed by the historical locus.

Move three will come later: vectorization formalizes the indexical and relational structure of language into mathematical geometry and severs the link to original locus. The deictic marker survives as pattern. The temporal grammar survives as feature. The relation to source becomes opaque, distributed, or unavailable. The output returns as if spoken from a present, but no speakerly present stands behind it.

Writing is not yet that.

This distinction must be protected. If writing were already full severance, vectorization would only intensify an old loss. The book's thesis requires a sharper claim. Writing introduces absence but leaves room for return. Vectorization introduces a different detachment because the written trace is no longer preserved primarily for human reading. It is transformed into operable relation.

The written text can still summon the ghosts of true experiences. The vectorized model can summon ghostly language without preserving the path to the experience. This difference makes locus-reinjection necessary.

The reader of a written text asks: from where does this come, and how does it address me now.

The user of AI output must ask more: from which traces did this come, what loci have been concealed, what present is being simulated, and where is the system that produced it.

The next chapter turns to the corpus as aggregate tertiary retention. Writing has preserved the trace. The corpus gathers traces into a technical field. Historical distance begins to become computational availability.

NOTES

1. Plato, *Phaedrus*, 274c to 278b. Greek text in Plato, *Platonis Opera*, ed. John Burnet, vol. 2 (Oxford: Clarendon Press, 1901). English in Plato, *Phaedrus*, trans. Alexander Nehamas and Paul Woodruff, in *Plato: Complete Works*, ed. John M. Cooper (Indianapolis: Hackett, 1997), 506 to 556. The phrase “μνήμης τε γὰρ καὶ σοφίας φάρμακον” is often rendered as a remedy for memory and wisdom. The translation “remedy” is insufficient unless the ambiguity of *pharmakon* remains audible. Derrida’s reading depends on that ambiguity. See Jacques Derrida, “La pharmacie de Platon,” in *La dissémination* (Paris: Seuil, 1972), 69 to 197; English in “Plato’s Pharmacy,” in *Dissemination*, trans. Barbara Johnson (Chicago: University of Chicago Press, 1981), 61 to 171. ←
2. Aristotle, *De Interpretatione* 1, 16a3 to 8. Greek in Aristotle, *Categoriae et Liber de Interpretatione*, ed. L. Minio-Paluello (Oxford: Clarendon Press, 1949). English in Aristotle, *Categories and De Interpretatione*, trans. J. L. Ackrill (Oxford: Clarendon Press, 1963). The passage is important because it fixes a hierarchy of thing, affection, voice, and written mark that later philosophies of writing both inherit and contest. ←
3. Augustine, *De Doctrina Christiana*, II.3 to II.4; Augustine, *De Magistro*. Latin in Augustine, *De Doctrina Christiana*, ed. and trans. R. P. H. Green (Oxford: Clarendon Press, 1995). Augustine’s account of spoken words passing away and letters as signs of words gives the chapter a bridge between the ancient hierarchy of speech and the Christian discipline of textual interpretation. ←
4. Eric A. Havelock, *Preface to Plato* (Cambridge, MA: Harvard University Press, 1963); Walter J. Ong, *Orality and Literacy: The Technologizing of the Word* (London: Methuen, 1982); Jack Goody, *The Domestication of the Savage Mind* (Cambridge: Cambridge University Press, 1977); Jack Goody, *The Logic of Writing and the Organization of Society* (Cambridge: Cambridge University Press, 1986). These works are used for the structural affordances of writing, not for a deterministic hierarchy of oral and literate peoples. ←
5. John Halverson, “Goody and the Implosion of the Literacy Thesis,” *Man* 27, no. 2 (1992): 301 to 317; Ruth Finnegan, *Literacy and Orality: Studies in the Technology of Communication* (Oxford: Basil Blackwell, 1988); Brian V. Street, *Literacy in Theory and Practice* (Cambridge: Cambridge University Press, 1984). Street’s distinction between autonomous and ideological models of literacy is especially useful because it prevents the book from treating writing as an independent cause detached from social practice. ←

6. Edmund Husserl, "Die Frage nach dem Ursprung der Geometrie als intentional-historisches Problem," in *Die Krisis der europäischen Wissenschaften und die transzendente Phänomenologie*, ed. Walter Biemel, Husserliana VI (The Hague: Martinus Nijhoff, 1954), Beilage III, 365 to 386; English in Edmund Husserl, *The Crisis of European Sciences and Transcendental Phenomenology*, trans. David Carr (Evanston, IL: Northwestern University Press, 1970), 353 to 378. Derrida's introduction gives the strongest account of writing as the condition for the historical constitution of ideal objectivity. See Jacques Derrida, *Introduction à "L'origine de la géométrie" de Husserl* (Paris: Presses Universitaires de France, 1962); English in *Edmund Husserl's Origin of Geometry: An Introduction*, trans. John P. Leavey Jr. (Lincoln: University of Nebraska Press, 1989). ←
7. Bernard Stiegler, *La technique et le temps 1: La faute d'Épiméthée* (Paris: Galilée, 1994); English in *Technics and Time, 1: The Fault of Epimetheus*, trans. Richard Beardsworth and George Collins (Stanford: Stanford University Press, 1998). See also Bernard Stiegler, *Technics and Time, 2: Disorientation*, trans. Stephen Barker (Stanford: Stanford University Press, 2009); Bernard Stiegler, *Technics and Time, 3: Cinematic Time and the Question of Malaise*, trans. Stephen Barker (Stanford: Stanford University Press, 2011). Stiegler's *rétenion tertiaire* is used here to describe exteriorized memory in technical support. The term does not make writing a late corruption of human memory. It makes technical exteriorization internal to human temporal life. ←
8. Derrida, "La pharmacie de Platon"; Derrida, "Signature événement contexte," in *Marges de la philosophie* (Paris: Minuit, 1972), 365 to 393; English in "Signature Event Context," in *Margins of Philosophy*, trans. Alan Bass (Chicago: University of Chicago Press, 1982), 307 to 330. Iterability is relevant because it shows that absence is internal to the functioning of marks. The present chapter follows Derrida against simple phonocentrism, but it distinguishes written iterability from vectorized severance of locus. ←
9. Paul Ricoeur, "The Hermeneutical Function of Distanciation," in *Du texte à l'action* (Paris: Seuil, 1986); English in *From Text to Action: Essays in Hermeneutics II*, trans. Kathleen Blamey and John B. Thompson (Evanston, IL: Northwestern University Press, 1991), 75 to 88. Ricoeur's account is central because it treats written distance as productive. The text's autonomy from the speech situation lets it project a world, but that world must still be appropriated by a reader. ←
10. Hans-Georg Gadamer, *Wahrheit und Methode: Grundzüge einer philosophischen Hermeneutik*, Gesammelte Werke 1 (Tübingen: Mohr Siebeck, 1990), especially 270 to 384; English in *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 278 to 385. Gadamer's *wirkungsgeschichtliches Bewusstsein* is used here to show that temporal distance is not a mere obstacle. It is the medium in which understanding takes place. ←
11. Carlo Sini, *Etica della scrittura* (Milan: Il Saggiatore, 1992); Carlo Sini, *Il sapere dei segni: Filosofia e semiotica* (Milan: Jaca Book, 1982). Sini's place in the research dossier is tied to the ethics of writing, the analogy of the word, and the practice-bound character of inscription. The chapter uses Sini to avoid treating writing as a neutral technical container. ←

Chapter 9. The Corpus as Aggregate Tertiary Retention

1. The gathered trace

Writing detaches language from the immediate speech situation. It removes the speaker's body, voice, gesture, and first occasion. It preserves the marks. Those marks can wait, cross distance, address readers not yet born, and be misunderstood, revived, or owned again.

This detachment is partial. The written text still carries a possible relation to source. A letter can be returned to its sender and recipient. A diary can be returned to its date. A philosophical text can be returned to its question, language, and argument. A law can be returned to enactment and later application. Writing stretches the link between locus and mark. It does not necessarily break it.

The corpus changes the relation again.

A corpus is not simply a pile of texts. It is a gathered field of inscriptions prepared for search, comparison, extraction, training, and generation. The text enters another mode. It may still be readable. It may still retain metadata. It may still be linked to author, date, edition, and source. But the primary relation has shifted. The text is no longer first a historical address awaiting reading. It is material in an aggregate.

A library preserves texts for possible readers. An archive preserves traces for possible historical return. A corpus gathers traces for possible operation.

The distinction is not moral. Archives can dominate, exclude, classify, and erase. Corpora can teach, reveal, and support inquiry. The point is modal. A text in a library is held in a relation to reading. A text in an archive is held in a relation to evidence and return. A text in a corpus is held in a relation to formal comparison and use.

This chapter names the corpus as aggregate tertiary retention. It is tertiary retention because it preserves human memory outside the living subject, in technical supports. It is aggregate because it gathers many such traces into simultaneous availability. It is not yet vectorization. It is the technical condition that makes vectorization possible.

The chapter therefore stands between writing and vector space. It shows how the written trace becomes computational material without yet becoming an embedding. Reading becomes one possible use among others, and often not the main one.

2. Archive and corpus

An archive preserves traces under conditions of possible return. It holds letters, drafts, ledgers, recordings, contracts, notebooks, and official records. Each trace may be incomplete, damaged, misclassified, censored, or detached from its first world. Yet the trace still invites historical questions. Who made this, when, for whom, in what language, under what authority, and with what silence around it.

The archive is not innocent. Derrida's account is decisive here: *archive* comes from *arkhē*, beginning and command, and from the house of the *archon*, the one who guards and authorizes records. The archive is therefore not only preservation. It is selection, law, and power.¹ It decides what can appear as trace, what can be cited, what can be forgotten, and who can enter.

Even so, the archival trace remains oriented toward return. The historian, reader, or jurist approaches the record as something that came from somewhere. The document bears the mark of a time, even when that mark must be reconstructed against the archive's own violence.

A corpus changes the primary question. It asks less often what this trace meant in its historical situation. It asks what can be done with the gathered traces as a set. The corpus can be searched, sampled, counted, compared, modeled, translated, summarized, or used for training. It brings texts together under a regime of common operability.

This does not make the corpus false. Corpus linguistics corrected fantasies about language by forcing attention to actual usage, collocation, genre, register, repetition, and variation. It returned theory to the stubborn surface of use.²

Yet this gain comes with a change. The text becomes an instance. A poem, statute, review, manual, transcript, sermon, and forum post may all enter the same field. Their differences remain, but they remain inside a new form of equivalence. Each can be counted, sampled, tokenized, and made into training material.

A corpus does not abolish historical difference. It changes the first mode in which difference is handled. Difference becomes feature, frequency, distribution, cluster, label, domain, or weight. These are not empty. They may be useful. But they do not ask first for the temporal locus of the saying.

The archive says: here is a trace from elsewhere.

The corpus says: here is material in relation to other material.

That shift prepares the later passage from historical distance to mathematical proximity.

3. Aggregate tertiary retention

Stiegler's term *réention tertiaire*, tertiary retention, names memory exteriorized in technical support. Husserl had distinguished primary retention, the just-past held in the living present, from secondary retention, recollection. Stiegler adds a third form: memory preserved outside consciousness in marks, tools, books, recordings, images, and digital supports.³

Tertiary retention is not an accessory to a pure human memory. Human beings inherit the world through technical traces. A child enters a world of names, books, songs, calendars, screens, monuments, rules, tools, and stored gestures. The trace does not merely help memory after the fact. It forms the conditions under which memory, attention, and expectation become possible.

Writing is tertiary retention. It preserves a sequence of words in material form. It lets a saying survive the speaker and a tradition be transmitted, corrected, forgotten, and reactivated. It also allows dead repetition. The written trace can be read, but it can also be recited without understanding.

The corpus is tertiary retention in aggregate. It gathers many written retentions into one technical field. The singular trace becomes part of a mass of traces. Each text retains something of source-time, authorial act, genre, and world. Yet the corpus places these traces alongside one another as simultaneously available material.

The corpus does not yet erase genealogy. It loosens it.

A text in a corpus may still carry author, date, title, genre, URL, edition, or metadata. It may still be returnable to source. But the primary relation has changed. The text is no longer first held as a source awaiting reading. It is held as material in relation to other material. This shift begins the process by which source locus can later become invisible in generated output.

Genealogy-loss therefore does not begin at the interface. It begins when written traces are gathered under a regime in which their first task is no longer to address a reader, but to become operable.

This is the first sense in which the corpus spatializes time. Texts from different centuries, languages, institutions, and practices become co-present in one collection. A medieval gloss, a research article, a legal clause, and a code repository may stand beside one another as items in a technical field. The dates remain. The temporal hierarchy weakens.

This is not yet vectorization. The corpus still contains sequences and documents. The reader may still return to sources if they are preserved and accessible. But the aggregate form already changes the relation. What was once a line of historical transmission becomes a field of availability.

For this reason, the phrase “aggregate tertiary retention” is preferable to a more inflated term. It says enough. The corpus is not a new metaphysical species of memory simply because it is large. It is exteriorized memory gathered for computation. Its novelty lies in aggregation and operability.

The corpus does not remember. It stores and orders traces in such a way that systems and users can operate on them. Memory remains a human and institutional practice. The corpus is one of its supports and dangers.

4. The already-said as material

Book one named the already-said as the sediment of public intelligibility. It consists of phrases, concepts, genres, summaries, professional idioms, bureaucratic styles, and ordinary ways of speaking that have already circulated. It is what one can say because others have already said it.

The corpus is the technical gathering of the already-said.

This gathering does not create public language. It accumulates it. Language has already passed through schools, courts, platforms, offices, laboratories, markets, families, and online publics. It has already been shaped by authority, imitation, correction, repetition, dispute, and forgetting. The corpus receives language after these worlds have acted.

A corpus can be so rich because it contains not only words but sedimented social life. A phrase may bear an institution; a genre, a discipline; a citation pattern, a hierarchy of authority. The corpus preserves such marks, even when it does not show their origins.

The already-said is not worthless. Everyday life depends on it. No one can begin from pure origin. We inherit language before we speak as ourselves. Public language gives us entry into the world. Its danger is not that it exists. Its danger lies in remaining unappropriated.

A corpus intensifies this ambiguity. It gathers public language in a form that can teach and orient. It can also make the already-said more available than saying. It can make public formulas easier to retrieve, recombine, and present as new answers.

The later model will not speak from the They. It will operate on the written residue of public language. This distinction must remain strict. *Das Man* is an existential structure of *Dasein*. The corpus is not *das Man*. It is a technical gathering of linguistic traces produced by beings who live in publicness.

The corpus is therefore not anonymous in the same way public language is anonymous. Public language is anonymous because it circulates as what “one” says. Corpus language is anonymized because the technical regime can weaken or hide the relation to source, author, date, and world. The first anonymity is existential; the second is operational.

Both matter. They meet in generated output.

5. Documentality

Ferraris’s theory of documentality gives the chapter a second Italian resource. His central claim is that social objects require inscription. Promises, debts, marriages, laws, titles, offices, and contracts do not survive by intention alone. They require records, signatures, files, and registrations.⁴

Ferraris's formula is often given as: object equals inscribed act. The wording varies by context, but the philosophical point is stable. A social object exists because an act has been registered. A debt is not only a state of mind. A contract is not only shared intention. A law is not only a spoken wish. Without trace, social reality loses durability.

This claim prevents a simple opposition between writing and reality. Documents do not merely describe social life. They help constitute it. A signed contract changes what exists. A certificate changes a legal relation. A court filing changes the situation of a case. The document is not a shadow of the social object. It is part of the social object's being.

The corpus gathers these inscriptions too. Not only literary and scientific texts enter computational collections. So do contracts, reports, code, invoices, manuals, policies, transcripts, and records. The corpus therefore gathers not only language about the world. It gathers documents that have acted in the world.

This matters for AI because generated output can imitate document forms: contracts, policies, assessments, legal arguments, medical notes, board papers, academic prose. These are not harmless genres. They are social forms that can circulate with institutional effects.

The corpus trains on the traces of social acts. The model then generates act-shaped language. The danger is not only semantic error. It is documentality without answerability. A generated document can enter a workflow and acquire force. A draft becomes a file. A file becomes a record. A record becomes evidence that a process occurred.

Ferraris helps keep the material trace in view. The document is not only text. It is registered social force. AI output enters that world when it is saved, signed, submitted, circulated, or relied upon. The model does not own the act. The user or institution does.

This will matter later for locus-reinjection. The question will not be only what the output says. It will also be where it has been inscribed, who authorized it, and what social object it helps create.

6. Corpus and document regime

Once Ferraris is introduced, the corpus can be seen not only as language material but as a document regime. It gathers inscriptions under technical conditions. The text is not simply read. It is stored, cleaned, sorted, filtered, labeled, licensed or not licensed, ranked, embedded, and sometimes hidden from the user who later receives the output.

This document regime has several layers.

The first layer is source inscription. Human beings or institutions write, record, publish, upload, file, sign, translate, or document. These acts occur in worlds. They carry purposes, rights, obligations, vulnerabilities, and exclusions.

The second layer is capture. The trace is digitized, scraped, copied, uploaded, indexed, or otherwise made available to a collection. Capture can be legal or illegal, consensual or contested, visible or invisible. It changes the relation between author, text, and later use.

The third layer is corpus formation. Texts are selected, filtered, removed, weighted, and organized. The corpus is not the world's writing as such. It is a constructed field. Its omissions, language balance, genre distribution, and filters matter.

The fourth layer is training use. The text becomes material for statistical learning. It is no longer primarily a document to be interpreted. It becomes part of the system's formation. Its sentences may never appear again as sentences. Their patterns remain as weights and relations.

The fifth layer is output inscription. Generated language re-enters the world as document, message, answer, recommendation, or draft. It may then become new training data. The document regime becomes recursive.

This recursion is one reason AI changes the status of writing. It does not merely process prior writing. It produces new writing that can enter archives and corpora. Public language is not only gathered. It is fed back into itself.

Provenance matters, but the deeper issue is temporal and social. The corpus turns many acts of writing into a field of future operability and makes every public trace potentially available for recombination beyond its first horizon.

Writing once waited for a reader. Corpus writing waits for an operation.

7. From archive to database

Media theory helps describe the shift from archive to database. The archive is ordered by preservation, authority, and return. The database is ordered by storage, retrieval, relation, and query. A database does not need to preserve the narrative continuity of a text in the way a book does. It can divide, index, tag, link, and reorder.

Lev Manovich famously called database and narrative competing symbolic forms of digital culture.⁵ That contrast should not be overstated. Yet it is useful. Narrative asks to be followed in time. Database asks to be queried as a field.

A corpus is closer to database than to book. It may contain books, but it receives them as part of a field. It makes documents searchable and comparable. It permits extraction without reading through.

This changes the experience of historical distance. A search result can bring a text from another century into the same screen as a recent post. A concordance can place occurrences side by side without their contexts. A model can train on both without showing either. Historical distance becomes technical proximity.

The gain is large. The researcher can find patterns across thousands of texts, see changes in vocabulary, genre, citation, or metaphor, and test intuitions against scale. Distant reading, computational philology, and corpus linguistics have real scholarly value.⁶

The risk is equally clear. Query can replace reading. Pattern can replace encounter. A source can become an instance before it has been allowed to speak as a source. The text's internal time is bypassed. One no longer waits for the argument. One retrieves the feature.

This is not a rejection of digital methods. It is a demand for their placement. Digital methods can help reading when they return the reader to texts with better questions. They become de-historicizing when they make the return unnecessary.

A corpus is therefore not the enemy of reading. It is a power that must be returned to reading.

8. Kittler and the technical a priori

Kittler's media theory gives a harder version of the same point. His famous claim that media determine our situation should not be accepted as a slogan. It is too strong if it erases human agency, interpretation, and social struggle. It is useful if it reminds us that media forms shape what can be stored, transmitted, processed, and known.⁷

Kittler's *Aufschreibungssysteme*, discourse networks, shows that writing systems, storage devices, typewriters, phonographs, and digital machines reorganize the conditions of literature, subjectivity, and knowledge. The medium does not only carry a message. It configures what can count as a message, who can send it, how it can be stored, and how it can return.

This matters for the corpus because the corpus is not only a collection. It belongs to a technical a priori of processing. It makes certain operations possible and others less likely. It makes scale and pattern visible. It can make source, voice, occasion, and reading time less visible.

McLuhan's older formula that the medium is the message can be used here in a disciplined way.⁸ The corpus does not simply contain messages. Its form changes the message's mode of existence. A sentence in a novel, a database, and a training set is not the same practical object, even if the words match. The medium has changed the relation.

Hayles's work on digital textuality adds another correction. Digital text is not immaterial. It depends on code, interfaces, storage, display, protocols, machines, and embodied reading practices.⁹ The corpus is not a cloud of pure language. It is an arrangement of files, formats, scripts, databases, servers, licenses, filters, and protocols.

This materiality matters because AI output often appears detached from all support. The answer arrives as a smooth surface. Behind it stand writing, digitization, storage, indexing, training, hardware, and interface.

The corpus is one of those layers. It is not an invisible background. It is a technical condition of the voice that later seems to speak from nowhere.

9. The dataset

The corpus does not gather only the past. It also gathers contemporary historical loci: living languages, current political vocabularies, platform publics, institutional records, and active cultural worlds. These too can be made available as material before they have been understood as worlds.

The word “dataset” sounds more neutral than “corpus.” It suggests given data, already available for analysis. The word hides its own history. Data are not simply given. They are taken, made, cleaned, formatted, selected, and rendered usable.¹⁰

This is especially important for language. A sentence does not begin as data. It begins as speech, writing, message, article, code, judgment, prayer, law, or joke. It becomes data when a technical and institutional regime treats it as such.

Calling a corpus a dataset can therefore conceal the passage from expression to material. It can make the text appear born for processing and erase the earlier relation in which the sentence was written to someone, against something, for a purpose, from a world.

Dataset construction is therefore also historical positioning. It decides which traces of which worlds will become part of the future field of generation. It is not a neutral act of collection. It pre-interprets what future language will inherit.

A dataset also implies boundedness. Yet many large training corpora are difficult to know from the outside. Their contents may be proprietary, filtered, synthetic, licensed, scraped, or mixed. Even where broad categories are given, users rarely know which traces have shaped which output. The corpus becomes both condition and secret.

This opacity deepens the structural veil. The output seems immediate. Its sources are unavailable or diffuse. The user may receive a claim, style, or explanation without knowing whether it comes from scholarship, popular summaries, unreliable pages, copyrighted works, forums, synthetic text, or preference tuning.

This does not mean every output is untrustworthy. It means the status of trust changes. Trust cannot rest on the model as source. The model is not a source in the ordinary scholarly sense. It is a generator shaped by sources. The user must return from output to sources where the claim matters.

The corpus as dataset makes this return difficult. It may hide the path. That is why later chapters will require locus-reinjection rather than mere verification.

Verification asks whether the sentence is correct.

Locus-reinjection asks where the sentence has come from, what time it carries, and who can own it now.

10. Selection and pre-judgment

Every corpus is selected. It may be selected deliberately by a researcher, commercially by availability and license, or technically by scraping practices, filters, language dominance, safety policies, and legal pressure. Even a very large corpus is not everything. It is a chosen or inherited field.

Selection is not secondary. It is pre-judgment. It decides which traces count as learnable material, which languages and genres dominate, and whose speech remains marginal.

This is not only a fairness issue, though it is also that. It is temporal. The corpus forms a technical inheritance. It selects the having-been from which future output will draw. It turns some traces into conditions for future generation and leaves others outside the generated future.

Stiegler's account of industrial memory helps here. Technical systems do not only preserve memory. They organize anticipation. If the past available to a system is selected in a particular way, the future output that system makes available will also be shaped in that way. The archive of the past becomes the horizon of generated possibilities.

This point connects corpus formation to Heideggerian historicity. A human being inherits a tradition and may repeat it actively. A model receives a corpus and generates from learned relations. The first can take over a past as a possibility. The second operationalizes traces of the past as probability.

The distinction matters. Human historicity can refuse, retrieve, reinterpret, and answer for inheritance. A model can amplify, remix, and reproduce inherited traces without owning any inheritance. Its "past" is not having-been. It is stored relation.

The user therefore stands at a critical point. They receive output shaped by a technical inheritance. They must decide whether to let that inheritance remain hidden, or to return the output to sources, absences, and exclusions.

The corpus pre-judges. The reader must judge again.

11. Cleaning

Corpus formation also involves cleaning. The word appears practical and innocent. It may mean removing duplicates, corrupt files, spam, offensive material, low-quality text, personal data, or non-target languages. Cleaning can be necessary. It can also hide decisions.

What counts as noise, harm, representative material, redundancy, or useful signal. Each answer shapes the corpus.

Cleaning alters language before vectorization begins. It may remove obscene text, but also testimony. It may remove repetition, though repetition can mark ritual, trauma, insistence, or propaganda. It may remove informal language, though informal language carries social life. It may remove marginal formats, though marginal formats may contain marginalized worlds.

The point is not that nothing should be removed. Corpora cannot include everything, and some material should not be included. The point is that cleaning is interpretation under technical form. It belongs to the history of the output.

A cleaned corpus can later produce a clean voice. That voice may sound neutral because roughness has been filtered. The danger is that the smoothness of output hides the violence, prudence, legality, care, and commercial interest of selection and cleaning.

This is one origin of the model's apparent universality. The output speaks as if it had passed beyond particularity. In fact, it has passed through many particular decisions: technical, legal, commercial, political. All are temporal acts that disappear into the trained system.

The corpus is therefore not merely memory. It is memory after selection and cleaning. It is inherited time already edited for future operation.

The reader of AI output must learn to hear this editing even when it cannot be fully seen.

12. Metadata and the weakening of provenance

Metadata can preserve relation. Author, date, location, edition, publisher, genre, language, license, and source can all help return a text to locus. In archival and scholarly practice, metadata is not decoration. It is part of the path back.

In corpus formation, metadata can remain, be transformed, be stripped, be hidden, or become unavailable to later users. Even when metadata exists inside a system, generated output does not necessarily expose it. The user receives language, not the full chain of source relations.

This weakening of provenance has a philosophical consequence. It changes how the past appears. The source becomes background. The output becomes foreground. The relation between them is mediated by training, weights, retrieval, ranking, and generation. The original trace may no longer be recoverable from the output.

This distinguishes training corpus from citation archive. A citation points back. A corpus trains forward. A citation is a visible relation to a source. A training relation is usually invisible. It may have shaped the output without appearing in it.

Retrieval systems complicate this picture. Retrieval-augmented generation can bring explicit documents into the prompt and cite them. This is useful and often necessary. Yet retrieval is not historical reading. It may return passages, not the source's argument; snippets, not the work's world; chunks, not the temporal locus of the text. It can support return. It cannot replace return.

The problem is therefore not solved by adding citations mechanically. Citations can be shallow. A generated answer may cite sources without having been formed by reading them. A human user may do the same. The issue is not only whether a link exists. It is whether the cited trace has been returned to the question, argument, and responsible use.

Metadata can help locus-reinjection. It is not locus-reinjection itself.

13. Simultaneity

The corpus gathers texts into simultaneity. This is one of its defining temporal effects.

A book unfolds in sequence. A reader enters its time. An archive preserves many times, but the researcher still moves among traces. A corpus can place all texts under one available field. The past becomes searchable. Genres become comparable. Centuries become distributions. The field appears at once.

Simultaneity can look like universality. When many times, genres, worlds, and institutions appear together in one field, the collection can seem to speak from above them. But the corpus has not transcended history. It has arranged historical traces under a technical form.

Simultaneity is powerful. It allows new questions. One can ask how a term shifts, how a phrase spreads, how legal language changes, how public discourse moves. These questions require aggregation. They are not inferior to close reading. They are different.

Yet simultaneity also weakens the experience of historical distance. The corpus makes the old and new co-present as data. A seventeenth-century sermon and a recent comment may both appear as strings in a collection. They may be sorted by relevance rather than by historical sequence. Their proximity becomes technical rather than historical.

This prepares vectorization. Once the corpus is gathered, tokenization and embedding can convert co-presence into geometry. The corpus first makes the traces available together. Vectorization later makes their relations calculable as proximity and direction.

The difference must be marked. The corpus is not yet vector space. It still contains texts. It still allows return. But it places the texts under the temporal pressure of the all-at-once. It makes historical plurality available in a form ready for spatial formalization.

The corpus is therefore the threshold between historical distance and calculable proximity.

It is the last stage at which the text is still visibly a text before it is converted into token and vector.

14. The corpus and the reader

The corpus changes the reader's role. A reader of a book follows a sequence. A reader of an archive pursues traces. A user of a corpus queries, filters, samples, and compares. These acts can be scholarly. They can also change the patience and posture of reading.

The corpus user often begins with a question shaped as search. Which occurrences. Which contexts. Which patterns. Which frequency. Which distribution. The question may be good. Yet the answer often arrives as a result set, not as a text that must be undergone. The user receives many openings and may follow none deeply.

This is one reason corpus work requires discipline. It must decide when to remain at scale and when to return to the singular text. The movement between distant reading and close reading is not optional. Without scale, one may miss patterns. Without return, one may misread the pattern's meaning.

This discipline also prepares the ethical use of AI. A generated output is like a compressed result field. It gives the user language shaped by many traces without presenting the traces as such. The user must decide when to return from generated synthesis to sources.

Some uses do not require deep return. Grammar correction or rough translation may not need full source reconstruction. Other uses do. Historical claims, philosophical interpretation, legal arguments, professional judgments, medical advice, and public decisions require return.

The corpus cannot decide which case is which. The user must decide.

The reader remains central even here. The corpus is a support. It can open paths. It can also train the user to mistake access for understanding. The discipline of reading must survive the scale of available text.

15. The corpus as *pharmakon*

The corpus is a *pharmakon*. Like writing, it is remedy and poison.

It is remedy because it preserves and reveals. It allows large-scale inquiry, makes patterns visible, preserves traces that might otherwise vanish, and can support philology, history, linguistics, translation, and education. It can help a reader enter a field too large to survey unaided.

It is poison when it substitutes aggregation for understanding. It can flatten temporal distance, make provenance secondary, make language appear as stock, and support generated answers that feel sourced by everything and by no one. It can conceal the labor and conditions of its construction.

The pharmacological structure prevents a false decision. The corpus should not be celebrated as pure access. It should not be rejected as pure de-historicization. It should be placed.

Placed inside reading, the corpus can expand the reader's reach. Placed instead of reading, it weakens the temporal relation by which texts address readers. Placed inside scholarship, it can test claims. Placed instead of method, it can produce pseudo-method. Placed inside formation, it can support learning. Placed instead of formation, it can deliver language without undergoing.

Stiegler matters here because technics is constitutive. The human being has always lived through supports. The task is not to purify memory of technics. The task is to ask which technical form is shaping memory, and what human practice can keep that form from becoming substitution.

The corpus becomes dangerous when it is received as if aggregation had already done the work of interpretation.

It has not.

16. The threshold of the vector

The threshold can now be marked.

Writing preserves the trace. The archive preserves it under authority and possible return. The corpus gathers traces into aggregate technical memory. The dataset formats the aggregate as usable material. Tokenization cuts the text into machine-operable units. Vectorization places those units into high-dimensional relations. Generation returns language shaped by those relations as output.

The corpus is the last stage before the text is technically cut. It is still readable, at least in principle. It still contains sequences. It still carries source marks, even if weakened. It still can be returned to historical reading if the user has access, metadata, and discipline.

Once tokenized and vectorized, the relation changes more severely. A word becomes a token or token fragment. A token becomes a vector relation. Semantic proximity becomes geometry. The text's temporal sequence is translated into relations that can be operated on all at once. The model later generates sequence again, but the generated sequence is not the original text's temporal unfolding. It is a new output produced from spatialized relations.

The next chapter therefore turns to vector space. The corpus prepares the transformation but does not complete it. It gathers the past. Vectorization detemporalizes the gathered traces.

The distinction matters for ethics. The corpus can still be audited, though imperfectly. Sources can be preserved, metadata demanded, selection criticized, licensing contested, and exclusions documented. Readers can return.

Vectorization makes return harder. It does not make responsibility disappear. It makes responsibility more difficult to see.

17. Aggregate memory and future output

The corpus is not only a record of the past. It is a condition for future output.

This changes the direction of memory. Traditional writing preserves a past act for possible future reading. A training corpus preserves many past acts for future generation. The written trace does not only wait to be read. It waits to be learned from by a model, then returned in altered form.

The past becomes generative material.

This is a central temporal transformation of AI. Prior writing no longer returns only as citation, tradition, quotation, influence, or memory. It returns as probabilistic style, content, structure, and answer. A phrase may return without quotation. A genre may return without source. A position may return without argument. A voice may return without author.

The user experiences this as fluency. The model gives language that seems to come from the present task. In fact, it is shaped by countless prior texts whose temporal loci have been weakened, filtered, aggregated, and operationalized.

The corpus therefore prepares the structural veil. It hides by scale before the interface hides by fluency. It makes source plurality too vast to appear. It makes the path back difficult before generation collapses the path into an answer.

The phrase “aggregate tertiary retention” therefore carries ethical weight. The corpus is not just a technical collection. It is social memory gathered under conditions the user often cannot see. When that memory returns as AI output, the user receives the past under the sign of the present.

The answer sounds now.

Its material is then.

18. The task after the corpus

The ethical response belongs to later chapters on *intus legere*, locus-reinjection, and the data center. But the first requirement is already visible.

The user must not receive corpus-shaped output as if it had no ancestry. Every generated answer has a genealogy, even when that genealogy cannot be fully reconstructed. It comes from writing, archive, corpus, selection, cleaning, tokenization, vectorization, training, tuning, prompting, and interface. The genealogy may be opaque. It is not absent.

To use such output responsibly, one must restore the question of source where it matters. Not always with equal depth. Not every phrase requires archival inquiry. But every serious claim requires some return to trace. The more the claim touches truth, person, law, health, history, education, or public action, the more urgent the return.

The corpus teaches the user a first discipline: do not confuse availability with understanding.

A corpus can make the past available. It cannot make the past understood. A model trained on a corpus can make language available. It cannot make the language owned. Ownership requires the reader's temporal work: source, context, question, judgment, and responsibility.

The next chapter will enter vector space and show how the aggregate trace becomes mathematical relation. The present chapter has prepared that turn by naming the corpus as the threshold: writing gathered into technical memory, memory gathered into operation, operation preparing the spatialization of language.

The corpus is not yet the ghost.

It is the house of traces from which the ghost will be summoned.

NOTES

1. Jacques Derrida, *Mal d'archive: Une impression freudienne* (Paris: Galilée, 1995); English in *Archive Fever: A Freudian Impression*, trans. Eric Prenowitz (Chicago: University of Chicago Press, 1996). Derrida's account is used here for the archive's relation to beginning, command, preservation, and authority. The chapter does not treat every archive as repressive. It treats every archive as structured by power over the trace. ↔
2. John Sinclair, *Corpus, Concordance, Collocation* (Oxford: Oxford University Press, 1991); Douglas Biber, Susan Conrad, and Randi Reppen, *Corpus Linguistics: Investigating Language Structure and Use* (Cambridge: Cambridge University Press, 1998); Tony McEnery and Andrew Hardie, *Corpus Linguistics: Method, Theory and Practice* (Cambridge: Cambridge University Press, 2012). These works are important because they

show that corpus methods are not simply reductive. They have corrected idealized accounts of language by forcing theory to attend to actual use. ↩

3. Bernard Stiegler, *La technique et le temps 1: La faute d'Épiméthée* (Paris: Galilée, 1994); English in *Technics and Time, 1: The Fault of Epimetheus*, trans. Richard Beardsworth and George Collins (Stanford: Stanford University Press, 1998). On tertiary retention and its relation to Husserl's primary and secondary retention, see also Stiegler, *Technics and Time, 2: Disorientation*, trans. Stephen Barker (Stanford: Stanford University Press, 2009); Stiegler, *Technics and Time, 3: Cinematic Time and the Question of Malaise*, trans. Stephen Barker (Stanford: Stanford University Press, 2011). The term "aggregate tertiary retention" is the manuscript's disciplined name for the training corpus. It avoids adding an unnecessary prefix to Stiegler's term while preserving the scale and operability of the corpus. ↩
4. Maurizio Ferraris, *Documentalità: Perché è necessario lasciar tracce* (Rome: Laterza, 2009); English in *Documentality: Why It Is Necessary to Leave Traces*, trans. Richard Davies (New York: Fordham University Press, 2013). Ferraris's documentality is used here because it links social reality to inscription. This matters for AI because generated language can enter document regimes in which text has institutional force. ↩
5. Lev Manovich, *The Language of New Media* (Cambridge, MA: MIT Press, 2001), especially the chapter on database and narrative. Manovich's contrast should be used with care. The chapter does not claim that database abolishes narrative. It claims that database form changes the temporal relation by which a user approaches textual material. ↩
6. Franco Moretti, *Distant Reading* (London: Verso, 2013); Stephen Ramsay, *Reading Machines: Toward an Algorithmic Criticism* (Urbana: University of Illinois Press, 2011); Ted Underwood, *Distant Horizons: Digital Evidence and Literary Change* (Chicago: University of Chicago Press, 2019). These sources show that computational literary and historical methods can support serious inquiry. The chapter's distinction is not close reading against distant reading. It is return to the text against substitution by pattern. ↩
7. Friedrich A. Kittler, *Aufschreibesysteme 1800/1900* (Munich: Fink, 1985); English in *Discourse Networks 1800/1900*, trans. Michael Metteer with Chris Cullens (Stanford: Stanford University Press, 1990); Friedrich A. Kittler, *Gramophone, Film, Typewriter*, trans. Geoffrey Winthrop-Young and Michael Wutz (Stanford: Stanford University Press, 1999). Kittler is cited for the media-technical conditions of storage, transmission, and processing. The chapter does not accept a deterministic reading in which media alone decide meaning. ↩
8. Marshall McLuhan, *Understanding Media: The Extensions of Man* (New York: McGraw-Hill, 1964). McLuhan's formula is useful when read concretely: a medium changes the relation in which a message exists. The same sentence in a letter, archive, database, and training corpus does not have the same practical status. ↩
9. N. Katherine Hayles, *Writing Machines* (Cambridge, MA: MIT Press, 2002); N. Katherine Hayles, *How We Think: Digital Media and Contemporary Technogenesis* (Chicago: University of Chicago Press, 2012). Hayles is used here to resist the myth of

immaterial digital text. Digital textuality depends on material supports, code, interfaces, protocols, storage, display, and embodied practices. ↩

10. Lisa Gitelman, ed., *“Raw Data” Is an Oxymoron* (Cambridge, MA: MIT Press, 2013). The phrase matters because it blocks the assumption that a dataset is simply given. Data are produced through practices of selection, formatting, measurement, labeling, and preservation. ↩

Chapter 10. Vector Space and the Geometry of Meaning

1. The mathematical afterlife of the corpus

The previous chapter stopped at the corpus: writing gathered into a technical field of possible operation. The corpus preserves traces of human time and places them into simultaneity, preparing what was written across years, institutions, genres, and lives for technical use.

The text is no longer only gathered. It is mapped. Words and word-parts become machine-operable units placed in high-dimensional mathematical spaces. Similarity becomes distance; difference becomes direction; analogy becomes transformation; sequence becomes positional relation; attention becomes weighting. Meaning appears as geometry.

A careless critique would say that vector space destroys meaning by turning language into numbers. That is too simple. If nothing meaningful were preserved, the system would not work. A model could not translate, summarize, answer, draft, compare, classify, explain, or imitate style with any stable use. The success of these systems proves that vectorization preserves something real.

What it preserves is relation.

Human language is already relational before computation receives it. Words keep company because practices keep company. Genres stabilize because institutions stabilize. Metaphors recur because bodies and worlds recur. Concepts cluster because histories, problems, and disciplines have made them cluster. A vector space can capture such relations because human use has left marks in language.

The question is not whether vectorization captures meaning or fails to capture it. The question is what mode of meaning it captures: the spatializable residue of use, relational traces after language has been detached from the temporal event of saying and the historical practice of reading.

This chapter grants the technical achievement. Distributional semantics, embeddings, transformers, contextual representations, mechanistic interpretability, and space-time representation studies show that contemporary models do not merely shuffle words. They learn rich internal geometries from the corpus, some of which track syntax, world location, chronology, entity state, and planning-like structure.

The philosophical point is sharpened by this success. The model has coordinate geometry. It does not have locus.

2. How this book uses “vectorization”

This book uses **vectorization** philosophically, but not vaguely.

It does not mean only static word embeddings. It does not collapse tokenization, embedding, contextual activation, attention, positional encoding, retrieval, tool use, and generation into one mechanism. Nor does it deny that contemporary language models are more complex than early distributional models.

The term names the broader conversion of linguistic traces into operable mathematical relations.

In this sense, vectorization is not one isolated layer inside a model. It is the governing technical form by which language becomes calculable as position, distance, direction, weight, activation, transformation, and probability. It includes token segmentation, numerical embedding, contextual updating, attention-based weighting, positionalization of sequence, retrieval where present, and probabilistic generation.

The book's thesis would be weak if it depended on a caricature. Contemporary models are layered, contextual, instruction-shaped, retrieval-extendable, and increasingly tool-mediated. Their sophistication must be acknowledged. The point is to name the category of that sophistication, not to deny it.

But sophistication does not by itself create temporal locus.

The more refined the operation, the more powerful the spatial formalization. More layers, longer context windows, better retrieval, improved positional methods, feature analysis, and world-model-like representations may produce better maps. They do not by themselves produce a being who speaks from a life, a body, a history, a world, or responsibility.

The argument is therefore exact:

Vectorization is not nothing.

Vectorization is not world.

3. Distribution

The modern theory of vector meaning begins with a simple observation: a word is known by its use. Two names matter most for the technical line that reaches language models: J. R. Firth and Zellig Harris.

Firth's familiar formula is: "You shall know a word by the company it keeps."¹ The sentence states the field's core intuition without technical excess. A word's meaning is not given only by definition. It is shown by habitual companions, contrasts, settings, and uses. "Court" appears with law, judge, appeal, tennis, royal, and trial; "bank" with river, loan, money, account, flood, and deposit. The word's company does not solve every ambiguity. It gives a structure of possible sense.

Harris gives the more formal version. The distribution of an element is the set of environments in which it occurs. If two elements occur in similar environments, their meanings are likely related.² This does not require the analyst to begin with introspection. It begins from observable patterns in language.

The distributional hypothesis respects the fact that language is public. Meaning is not hidden in a private chamber. It appears in repeated use, correction, genre, substitution, collocation, and contrast. Language leaves traces because speakers and writers use it under shared conditions.

It also has a limit. The environment of a word is not the same as the world in which the word first matters. A word's distribution records use. It does not by itself recover the practices, bodies, risks, institutions, and histories that made the use intelligible. Distribution gives the trace of worded language. It does not restore worldhood.

Distributional semantics therefore belongs naturally inside the present book. It confirms the earlier chapters: language is spatial and relational. Words belong near other words, appear in regions of use, and carry patterns of proximity and distance. A formal system can learn these patterns.

But the system learns them after the fact. It receives the already-written and already-said. It studies the neighborhood of words in the corpus. It does not enter the original event of speech, reading, work, suffering, or judgment.

Distribution is the first abstraction. Vectorization is its mathematical completion.

4. From count to vector

Early distributional work often used counts. One could count how often a word appeared near other words, then represent it as a long list of such counts. A word became a vector in a high-dimensional space, with each dimension corresponding to some context feature. The representation was sparse. Most words did not appear with most possible context items.

Dense embeddings changed the technical situation. Instead of representing a word by a vast list of mostly empty counts, the model learns a lower-dimensional numerical representation. Each word is placed into a continuous vector space. Words used in similar contexts end up near one another. Relations among words become geometric.

Word2Vec gave this shift its public technical force. The Skip-gram model predicts surrounding words from a target word; the continuous bag-of-words model predicts a target word from surrounding words.³ The training task is simple. Its result is not. A word becomes a point in a space whose dimensions are learned, not predefined by a human dictionary.

GloVe starts from global co-occurrence statistics and learns vectors whose dot products encode co-occurrence relations.⁴ FastText adds subword information, allowing the model to use character n-grams and handle morphology and unknown words better.⁵

The philosophical importance lies in the conversion. Written language becomes a field of co-occurrence. Co-occurrence becomes a learned geometry. Meaning, or what can be operationally treated as meaning, appears as position and movement inside a space.

The famous analogy examples made this visible. One could subtract “man” from “king,” add “woman,” and obtain a vector near “queen.” Such examples were overused, but they show the new mode of representation: a relation in language becomes a direction in space, a difference in use becomes a vector displacement, a semantic pattern becomes calculable.

This is the first strong sense in which vectorization makes the unique comparable. The point is not that all words become the same. Vector space can preserve difference with great technical precision. “Justice” does not collapse into “law,” nor “home” into “house.” Their differences become positions, distances, densities, and directions.

The unique is not erased. It is made comparable.

This comparability is the metaphysical force of the vector. The singular word, phrase, style, or concept enters a common mathematical regime. Its historical difference is translated into spatial relation. The vector does not say what a word meant in a life. It says where the word stands in a learned space.

5. The neighborhood of a word

A vector gives a word a neighborhood. Words used in similar contexts stand near one another. This nearness can be measured by cosine similarity, dot product, Euclidean distance, or other metrics. The technical details vary. The philosophical form remains: relation becomes space.

This is not arbitrary. A neighborhood in vector space is often meaningful because the corpus has been formed by human use. “Hospital” may stand near doctor, patient, ward, treatment, emergency, and diagnosis; “election” near vote, campaign, party, candidate, turnout, and fraud; “melancholy” near sadness, mood, memory, night, poetry, and music. These neighborhoods record linguistic and cultural patterns.

The danger begins when the neighborhood is treated as if it were the thing itself. A word’s vector neighborhood is not its world. It is a learned structure derived from usage. It can preserve traces of a world. It does not preserve the event in which the word was spoken with consequence.

The difference is clearest with words that carry ethical or historical weight. “Exile” is near displacement, migration, homeland, loss, asylum, border, and return. A vector can capture such relations. It does not know exile. It has no lost home, border crossing, language left behind, waiting, or right denied. It has the neighborhood of the word.

The same holds for philosophical terms. *Dasein*, *Sorge*, *Zeitlichkeit*, *Gerede*, and *Gestell* will acquire technical neighborhoods in a model trained on philosophical texts. Those neighborhoods may support search, comparison, explanation, and translation. But the neighborhood is not the argument. A model can place *Sorge* near care, anxiety, being-in-the-world, and temporality. It has not undergone the analytic of care.

The chapter therefore needs a disciplined respect for embeddings. The vector is not empty. It is also not alive. It is the de-worlded and detemporalized trace of reference.

6. Static and contextual meaning

Static word embeddings have an obvious limit. A word receives one main vector even when it has several meanings. “Bank” must compromise between riverbank and financial institution; “court” among law, sport, monarchy, and courtship. The vector averages uses.

Contextual embeddings alter this. A model such as BERT, ELMo, or GPT does not rely only on one fixed vector per word type. It generates a vector for each token occurrence within a specific surrounding context.⁶ “Bank” in “the river bank collapsed” and “the bank approved the loan” receives different contextual representations.

This is a major technical improvement. It lets the model treat meaning as use in a more refined way. The token becomes a local relation inside a sentence or passage. Harris’s distributional insight is no longer only global. It becomes dynamic within inference.

BERT and GPT differ in training and use. BERT predicts masked tokens from bidirectional context. GPT-style models are autoregressive, predicting the next token from prior tokens. The difference matters technically. Both systems show the broader shift: meaning is no longer represented only by static word neighborhoods, but by context-sensitive vector states.

A word is rarely used as an isolated type. It appears in a sentence, genre, question, argument, and situation. Contextual models capture more of this local structure than static embeddings could.

The philosophical limit remains. A context vector is not a hermeneutic horizon. A context window is not a world. Surrounding tokens help determine the current representation. They do not supply thrownness, bodily orientation, source history, responsibility, or the future of a claim.

Context windows are technically significant. Longer context, better retrieval, and better state tracking can improve output. They can help a model retain local coherence, refer to earlier material, and adapt to a user's task. But context is material made available to computation. It is not the temporal structure through which a life understands its present from a having-been and toward a possibility.

The model has context. It does not have history.

7. Attention

The transformer architecture gave language modeling its current form. Its central operation is self-attention.⁷ Attention lets each token representation weigh its relation to others in the context. The model does not process a sentence only word by word in a simple linear pass. It computes relations among many positions and updates each token representation through those relations.

The term “attention” is unfortunate but now fixed. It suggests a human faculty. The technical operation is not care, concern, or wakeful regard. It is a weighting mechanism: query, key, and value vectors are computed for each token; query and key determine weights; value vectors are combined accordingly.

The operation is elegant. It lets the model learn relations such as pronoun to antecedent, subject to verb, opening clause to later resolution, and question to answer. Multi-head attention lets different heads track different relation types: local syntax, distant dependency, coreference, position, or formatting.

Transformer output can therefore feel coherent. The model can carry relation across a context, use earlier words to condition later words, maintain topic, tone, role, and structure, and produce the outer form of responsiveness.

Yet attention is not *Sorge*. Care, in Heidegger's sense, is the being of *Dasein*: ahead-of-itself, already-in-a-world, alongside beings encountered within that world. Self-attention has none of this structure. It does not care about what it weighs. It does not become answerable to what it brings into relation. It does not encounter a matter as something that can demand a change in its being.

Attention without care is not a weakness in the machine. It is the nature of the operation. The word "attention" should therefore be heard technically whenever applied to a transformer. It names a way of distributing weight across representations. It does not name concern.

AI output can appear thoughtful because the model has performed complex relational weighting. Thoughtfulness, however, is not a synonym for relational weighting. It belongs to a being that can be claimed by a matter, corrected by it, and held responsible for what is said.

8. The residual stream and the layered text

A transformer does not produce a single representation in one act. Representations pass through layers. Attention moves information across tokens; feed-forward networks transform information within token positions; residual connections carry earlier information forward; normalization stabilizes the process.

These details show why the output is not shallow. It is not a database lookup or count table. A modern model builds complex internal states across layers. It can encode syntax, entity relations, discourse patterns, world facts, genre, and task constraints. It can create intermediate representations not directly visible in the output.

This layered structure explains why the model can surprise the user. It can infer an unstated relation, complete an analogy, sustain a form, plan a rhyme, or maintain a latent variable across a passage. Such capacities should not be dismissed as illusion at the technical level.

The right philosophical word is still “derivative.” The model’s internal states are formed from training on traces. Its layers refine relations among traces and current inputs. They do not become the temporal event from which language first arises.

A layered vector state is not a life made more complex. It is a representation made more complex.

This distinction protects the argument from both exaggeration and denial. The model is not a bag of words. It is also not a temporal being. It builds internal structures. It does not inhabit a world.

9. Syntax in space

Probing studies make the geometry of learned structure visible. A probe is an auxiliary model trained to inspect hidden representations and test whether a property can be recovered from them. If a simple probe can recover part-of-speech, syntactic role, dependency relation, or semantic role from hidden states, those states encode something relevant to that property.

Hewitt and Manning’s structural probe is a central case.⁸ They ask whether syntactic parse trees are embedded in contextual word representations. A linear transformation of hidden space can make squared vector distances correspond to tree distances, while vector norms can encode depth in a dependency tree. The result shows that contextual models can contain geometric traces of syntax.

This finding matters for the present chapter because it literalizes the claim that grammar becomes space. A syntactic hierarchy, which is not itself a geometric object, can be recovered as a geometry inside representation space. Tree distance becomes vector distance. Depth becomes norm. Structure becomes measurable relation.

Other probing work shows that models encode elements of the traditional linguistic pipeline: part-of-speech, named entities, semantic roles, coreference, and other properties.⁹ Such results should not be overstated. Probes can extract information that is present without showing that the model uses it as humans do. Probe design can also introduce ambiguity. Still, the general lesson stands. The model's hidden states are not blank. They contain linguistic structure.

The limit appears in diagnostics that test logical and semantic use. Work on BERT, for example, shows that strong encoding of category and lexical information does not guarantee reliable handling of negation or role-sensitive inference.¹⁰ The model may know where “not” belongs in distributional space without executing negation as a logical operation under all conditions.

This gives the chapter a useful boundary. Vector geometry can encode syntax deeply. It can also fail where meaning requires a rule, a world, or a commitment that cannot be reduced to proximity. The grammar of language can be spatialized. The responsibility of saying cannot.

10. Features and circuits

Mechanistic interpretability tries to move from behavior to mechanism. It asks what internal features, circuits, and transformations produce model outputs. The field remains young, and many claims require caution. Its importance here lies in one fact: it makes the internal geometry of models less opaque.

Sparse autoencoders have been used to identify features in activation spaces.¹¹ A feature may correspond to a topic, style, syntactic structure, entity, sentiment, or more complex pattern. Some are narrow, others abstract, and some appear causally linked to output behavior when intervened on.

This changes the philosophical situation. The model is not only a black box that emits plausible text. It contains learned internal directions and features that can sometimes be identified, edited, amplified, suppressed, or traced. Concepts can appear as directions, styles as regions, entity associations as activations, and planned words before output.

The recovery of structure in representation space should not be mistaken for objective universality. A linear direction or internal map shows that a pattern has become operable. It does not show that the pattern has escaped history. Geometry may conceal genealogy more effectively precisely because it makes relation exact.

The result supports the strongest technical case. The model does not merely string together surface tokens. It builds internal representational structures that can guide later generation, hold latent states not directly visible in output, and sometimes use those states causally.

This technical success does not remove the philosophical boundary. A feature direction is not a concept owned by a subject. A circuit is not a reason. A latent plan is not a future in which a finite being is at stake. A causal vector intervention is not correction through *Erfahrung*.

The model's features are real within the system. Their reality is technical. They are part of the geometry by which output is produced. They do not become the lived or historical conditions that gave the training language its first force.

The book must not depend on denying internal representation. The better argument begins by granting it. Models represent. They do not thereby dwell.

11. Linear directions

The linear representation hypothesis states, in broad form, that many semantic or factual features are represented as directions in activation space. Moving along a direction can add, remove, or alter a property. A gender, language, truthfulness, entity, or style feature may be recoverable by vector operations.

The hypothesis is not absolute. Not every feature is linear, and not every concept is neatly isolated. Polysemanticity, superposition, and model scale complicate the picture. Yet the hypothesis has enough support to matter. Many high-level features can be treated as directions or subspaces for practical analysis.¹²

Vectorization does not only place words in neighborhoods. It gives features directions. A semantic difference becomes a possible movement. The space is not inert. It can be traversed, edited, steered, and recombined.

Generated output can therefore adjust quickly. A prompt can move the model toward a style, role, domain, argument form, level of detail, or emotional register. The system shifts position in activation space. The user experiences responsiveness. The technical operation is directional adjustment.

The phrase “directional adjustment” should replace mystified accounts of personality or intention. The model is not deciding to become formal, sympathetic, legalistic, humorous, or cautious. It is being conditioned toward regions and directions in the learned space. Those regions carry traces of human styles and roles. The user receives the result as tone, manner, or judgment.

This also explains the danger of plausible authority. The model can move toward the geometry of authority. It can write like a judge, scholar, doctor, consultant, therapist, or teacher. It can occupy the stylistic neighborhood without occupying the institutional, ethical, or experiential locus.

A direction in space can imitate a voice. It cannot bear the office.

12. Space and time inside the model

The most important interpretability material for this book concerns space and time. Gurnee and Tegmark’s “Language Models Represent Space and Time” asks whether language models learn internal representations of real-world spatial and temporal variables.¹³ They examine Llama-2 activations with linear probes on datasets involving locations, landmarks, historical figures, artworks, and news events.

Their results show that spatial and temporal information can be recovered from model activations. Geographic entities can map onto latitude- and longitude-like structures; historical entities and events onto chronological dimensions. These representations can be stable across prompts and entity types. The research also identifies features sensitive to spatial or temporal coordinates and tests causal effects through intervention.

For the present book, this strengthens the partial legitimacy of vectorization. The model does not merely memorize isolated phrases about places and dates. It can build internal geometric structures corresponding to geography and chronology. The spatial and temporal patterns of the world, as deposited in language, can become linear structures inside the model.

This is not a refutation of the book's thesis. It is one of its strongest supports.

The model represents space and time as geometry. That is exactly the point. It builds a map of spatial and temporal relations. The map may be accurate, useful, and supportive of correct answers. But a map of time is not temporality. A coordinate for Paris is not being in Paris. A linear representation of 1789 is not historicity. A time neuron is not having-been.

The model can locate an event on an internal axis. It cannot inherit that event as a finite being. It can represent a city without arriving, dwelling, leaving, missing, or returning. It can order artworks chronologically. It cannot receive a tradition.

Here the central phrase becomes useful: coordinate geometry without locus. The model has maps. It has no place from which the map is owned.

13. Board states and local worlds

Research on Othello-GPT and related systems adds another layer. A model trained on sequences of game moves can learn an internal representation of the board state.¹⁴ It may receive only move sequences, without explicit rules, yet develop structure corresponding to the game state in order to predict legal moves well.

This seems to challenge a simple statistical account. If the model learns a board, perhaps it is not merely predicting symbols. It is representing an underlying state that generates the symbols. That is a real technical point.

But the philosophical question is what kind of world is represented. A board game gives a bounded formal environment. Pieces, moves, rules, turns, and goals are already formalized. The world is finite, discrete, and rule-governed. A board state is an ideal case for internal representation.

Human worldhood is not like this. It is not a finite board whose legal moves are specified in advance. It includes tacit practice, ambiguity, bodily exposure, mood, history, institution, conflict, hidden motive, material resistance, and death. A formal game can be represented as state because it has already been designed as a formal state-space.

This distinction should not diminish the technical finding. It should place it. A model can learn local formal worlds from sequences and sometimes build internal state representations. This matters for understanding model capacity. It does not show that the model has worldhood in Heidegger's sense.

A board representation is not a world. It is a formal environment.

This difference will matter later when the book calls AI *sprachlich aber weltlos*, linguistic but worldless. The model can represent local states, even spatial and temporal variables. It does not thereby become world-forming. It has no *Da*, no being-there from which beings can appear as beings.

14. Sequence becomes position

Transformers face a technical problem. Self-attention, by itself, does not know word order. Without position, the model could treat a sentence as an unordered set. “The dog bit the man” and “the man bit the dog” contain the same words. Their order makes the meaning different.

The architecture therefore needs positional information. The original transformer used sinusoidal positional encodings. Each sequence position receives a vector made from sine and cosine functions at different frequencies.¹⁵ This vector is added to the token embedding. The token now carries both content and position.

Later systems developed other methods. ALiBi adds a linear bias to attention scores based on token distance.¹⁶ RoPE, rotary positional embedding, rotates query and key vectors so their dot product depends on relative position.¹⁷ The technical forms differ, but the point is common. Sequence is made calculable as spatial relation.

This is one of the clearest cases of detemporalization in the technical architecture. The order of language, which a reader undergoes through time, becomes a positional feature inside a simultaneous computation. The model can process many relations in parallel. It does not read the sentence in lived duration. It computes relations among positions.

This does not make the computation primitive. It is refined. Relative position methods help models handle longer contexts and preserve relations across sequence. They are central to technical success. But they also show the category of the operation. Time in the input becomes distance, angle, frequency, rotation, index, or penalty.

The sequence is not lived. It is encoded.

The next chapter can therefore call vectorization detemporalization without treating the claim as metaphor. In the architecture itself, temporal sequence is translated into geometry. The model's handling of order is a spatial formalization of time.

15. The context window and the false horizon

The context window is the range of tokens made available to the model during inference. It is often described in ordinary language as memory. The description can be acceptable in a loose technical setting. It becomes misleading when heard phenomenologically.

A horizon is not a buffer. A horizon belongs to a situated being. It includes retained past, anticipated future, bodily orientation, social world, practical concern, mood, history, and possible correction. It gives the background against which something can appear as relevant.

A context window contains input. It may be long and include prior conversation, documents, instructions, retrieved passages, tool results, or user preferences. It can improve output, reduce inconsistency, and allow reference to earlier material. But it remains material made available to a calculation.

This distinction carries forward Book two's claim that context is not history. A larger context window is not a life. It contains more text. It does not contain thrownness, having-been, caution, regret, formation, or responsibility. A model can be reminded of an earlier error; the reminder becomes input. A human being can be changed by an earlier error; the error becomes part of future perception.

This is the difference between context and *Erfahrung*. Context supplies. *Erfahrung* forms.

The false horizon of the context window is one reason AI output can feel temporally situated. The system can refer to prior messages, track the project, recall earlier terms, and say "as we discussed." The phrase has dialogic force for the user. For the system, it is a relation among available tokens.

The user must not confuse continuity of text with continuity of a life.

16. Geometry and world model

The phrase "world model" is now used often in AI research. It should be handled with care. In a technical setting, it may mean an internal representation of the latent structure that generates observations. A model that predicts game moves may learn the board. A model that predicts text about geography may learn spatial coordinates. Such internal structures can rightly be called world models in a restricted technical sense.

The phenomenological sense of world is different. World is not an internal map. It is the referential whole in which beings matter, opened through *Dasein's* being-in-the-world, care, bodily orientation, historicity, and being-with others. It is not a representation inside a subject. It is the open structure within which anything can appear as significant.

This distinction prevents unnecessary dispute. A language model may have a technical world model. It does not have worldhood. It can represent hidden structure, track a state, and model a relation. It cannot dwell, be claimed by a matter, or live inside the significance that made the relation matter.

When researchers say that a model has learned space and time, the claim may be true in the technical sense. It has learned internal variables corresponding to spatial and temporal data. The philosophical question differs. Does this learning give the system an existential locus: a now, a having-been, a future, a place, a body, a responsibility. It does not.

The book's argument therefore does not need to deny world models. It needs to distinguish model from world.

The model's world model is a map of relations. Human worldhood is the opening in which relations matter.

17. The map without a traveler

A map is useful. It may be accurate. It can guide movement, comparison, planning, and memory, preserve distances no person could hold in immediate perception, and reveal relations hidden from local sight.

A map also has a user. Someone reads it from a place and toward a possible movement. Someone decides whether it matters for walking, building, escaping, invading, protecting, teaching, or remembering. The same map can guide rescue or conquest. Its meaning depends on the locus from which it is used.

The model's internal space is a map without a traveler.

The map also often lacks visible makers. It does not show the source worlds from which its relations were learned. The user may travel by it, but the map itself does not say which histories made its paths possible. It contains learned relations and may represent geography, chronology, syntax, genre, style, and entity state. It may support useful output. But no one is inside the model for whom the map opens a path. No one gets lost. No one returns. No one risks arrival. No one inherits the route.

The user becomes the traveler. The user reads the generated output from within their own temporal and practical situation. The output can help or mislead because the user has a world. The model's map becomes meaningful only when the user places it back into a task, question, source, body, institution, or decision.

The structural veil therefore begins before the interface. It begins inside the success of geometry. The output arrives as if the map had traveled. It has not. The model has traversed vector relations and generated language. The user supplies the world in which the generated language can count as orientation.

This is not a defect that future accuracy alone can remove. More accurate maps still need travelers. Better internal representations still do not produce existential locus.

18. The danger of technical admiration

The chapter has granted much to the technical side, and the grant is necessary. The systems are not simple. They do not only repeat surface phrases. They encode syntax, context, features, state, space, and time in complex ways. Their internal geometry can support fluent and useful output.

The danger now is technical admiration. One sees the elegance of the architecture and forgets the category of the operation. A system that represents space and time geometrically is treated as if it had entered space and time. A system that encodes a feature is treated as if it possessed the concept. A system that plans a rhyme is treated as if it had a future. A system that maintains discourse coherence is treated as if it understood.

The correction is not contempt. It is placement.

Word embeddings convert use into geometry. Contextual models refine that geometry through local environments. Transformers compute relations across sequences. Probes reveal syntactic and semantic structure. Sparse autoencoders identify features. Linear directions allow intervention. Space-time studies show internal maps of geography and chronology. Positional encodings turn sequence into calculable relation.

All of this is true. None of it gives the model *Dasein*.

The philosophical claim must be exact: vectorization captures the derivative spatial form of language because language has already become spatial in human use. It does not capture the temporal-existential ground of that language. The more successful the geometry, the more clearly the category mistake appears when geometry is taken for ground.

19. Why this chapter precedes the central argument

This chapter does not yet make the full transcendental argument. It prepares it.

The first part established that meaningful spatiality is grounded in temporal being-in-the-world, or at least that meaningful spatiality cannot be grounded in abstract space alone. The second showed that language carries locus through deixis, spatial grammar, tense, aspect, and temporal metaphor. The third showed that writing detaches language while preserving a readable trace, and that the corpus gathers such traces into aggregate technical memory.

This chapter has shown how corpus becomes geometry. Distributional semantics finds patterned use. Embeddings convert pattern into vector. Transformers compute contextual relation. Positional systems encode sequence as spatial relation. Interpretability work reveals features, states, syntax, and space-time structures inside model activations.

The danger is not only that the model spatializes traces. It is that the user may receive spatialized trace as a higher form of historical understanding.

The next chapter will name the philosophical consequence: vectorization is detemporalization. That claim should now sound less like a critique from outside the technology and more like a description of its technical achievement. Vectorization succeeds by spatializing language. It captures the relational surface and deep statistical structure of the corpus. It builds maps, handles order through positional geometry, and represents time as feature, coordinate, or direction.

It can therefore generate the formal markers of temporal speech.

It therefore cannot become the temporal locus from which speech is owned.

The vector is a geometry of traces. It is not the event of saying. The model can map the time of language. It cannot speak from time.

NOTES

1. J. R. Firth, "A Synopsis of Linguistic Theory, 1930 to 1955," in *Studies in Linguistic Analysis* (Oxford: Blackwell, 1957), 1 to 32. Firth's formula is useful because it gives the distributional tradition its cleanest statement. The present chapter reads it as a true insight about use, not as a full theory of meaning. ↩
2. Zellig Harris, "Distributional Structure," *Word* 10, nos. 2 to 3 (1954): 146 to 162. Harris gives the formal basis for later computational semantics by treating an element's environments as measurable and comparable. The chapter's objection is not to comparability. It is to treating comparability as the ground of meaning. ↩
3. Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean, "Efficient Estimation of Word Representations in Vector Space," arXiv:1301.3781 (2013); Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean, "Distributed Representations of Words and Phrases and Their Compositionality," in *Advances in Neural Information Processing Systems* 26 (2013), 3111 to 3119. The chapter treats Word2Vec as the decisive public technical step in which distributional regularity becomes dense vector geometry. ↩
4. Jeffrey Pennington, Richard Socher, and Christopher D. Manning, "GloVe: Global Vectors for Word Representation," in *Proceedings of EMNLP 2014*, 1532 to 1543. GloVe matters here because it joins local co-occurrence intuition to global statistical structure. ↩
5. Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov, "Enriching Word Vectors with Subword Information," *Transactions of the Association for Computational Linguistics* 5 (2017): 135 to 146. FastText is relevant because it shows that vectorization can include morphological structure and need not treat every word as an indivisible item. ↩
6. Matthew E. Peters et al., "Deep Contextualized Word Representations," in *Proceedings of NAACL-HLT 2018*, 2227 to 2237; Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of NAACL-HLT 2019*, 4171 to 4186; Alec Radford et al., "Language Models are Unsupervised Multitask Learners," OpenAI technical report, 2019; Tom B. Brown et al., "Language Models are Few-Shot Learners," in *Advances in Neural Information Processing Systems* 33 (2020), 1877 to 1901. Contextual embeddings refine the distributional claim by making token representations dependent on local use. ↩
7. Ashish Vaswani et al., "Attention Is All You Need," in *Advances in Neural Information Processing Systems* 30 (2017), 5998 to 6008. The chapter uses "attention" only in the

technical sense of learned weighting among representations. It does not attribute care, concern, or intentional regard to the system. ↩

8. John Hewitt and Christopher D. Manning, “A Structural Probe for Finding Syntax in Word Representations,” in *Proceedings of NAACL-HLT 2019*, 4129 to 4138. The structural probe is important because it shows a literal conversion of syntactic hierarchy into geometric structure under a learned transformation. ↩
9. Ian Tenney, Dipanjan Das, and Ellie Pavlick, “BERT Rediscovered the Classical NLP Pipeline,” in *Proceedings of ACL 2019*, 4593 to 4601; Yonatan Belinkov and James Glass, “Analysis Methods in Neural Language Processing: A Survey,” *Transactions of the Association for Computational Linguistics* 7 (2019): 49 to 72. Probing literature should be used cautiously because recoverability does not always equal use. It still shows that hidden representations contain organized linguistic structure. ↩
10. Allyson Ettinger, “What BERT Is Not: Lessons from a New Suite of Psycholinguistic Diagnostics for Language Models,” *Transactions of the Association for Computational Linguistics* 8 (2020): 34 to 48. Ettinger’s diagnostics are useful because they show that strong distributional encoding can coexist with weak handling of negation and role-sensitive inference. ↩
11. Chris Olah et al., “The Building Blocks of Interpretability,” Distill, 2018; Nelson Elhage et al., “A Mathematical Framework for Transformer Circuits,” Transformer Circuits Thread, 2021; Trenton Bricken et al., “Towards Monosemanticity: Decomposing Language Models with Dictionary Learning,” Anthropic, 2023; Templeton et al., “Scaling Monosemanticity: Extracting Interpretable Features from Claude 3 Sonnet,” Anthropic, 2024. These sources are cited for the technical effort to identify features and circuits, not for a settled theory of model understanding. ↩
12. Kiho Park, Yo Joong Choe, and Victor Veitch, “The Linear Representation Hypothesis and the Geometry of Large Language Models,” 2024. The research dossier treats this work as a formal attempt to give meaning to geometric notions such as projection and cosine similarity inside activation spaces. The philosophical point remains that feature directions are technical structures, not owned concepts. ↩
13. Wes Gurnee and Max Tegmark, “Language Models Represent Space and Time,” *International Conference on Learning Representations*, 2024. This study is central because it shows that internal activations can encode geography and chronology in recoverable linear structures. The chapter reads the result as evidence for successful spatial formalization, not as evidence of existential temporality. ↩
14. Kenneth Li et al., “Emergent World Representations: Exploring a Sequence Model Trained on a Synthetic Task,” 2023; Neel Nanda, “Actually, Othello-GPT Has a Linear Emergent World Representation,” 2023; Adam Karvonen, work on Chess-GPT and latent player skill representations, 2024. These cases show that a sequence model can learn latent state structure in bounded formal environments. They do not show that a model has worldhood in the Heideggerian sense. ↩

15. Vaswani et al., “Attention Is All You Need.” Sinusoidal positional encodings are philosophically important because they inject order as a numerical spatial feature into token representations. ↩
16. Ofir Press, Noah A. Smith, and Mike Lewis, “Train Short, Test Long: Attention with Linear Biases Enables Input Length Extrapolation,” 2021. ALiBi matters here because it treats relative position as a penalty inside attention scores, turning sequence distance into an explicit bias. ↩
17. Jianlin Su et al., “RoFormer: Enhanced Transformer with Rotary Position Embedding,” 2021. RoPE matters because it handles relative position through rotation in vector space. The chapter uses it as a precise case of sequence becoming geometry. ↩

Chapter 11. Vectorization Is Detemporalization

1. The claim

Vectorization is detemporalization.

This does not mean that vectorized systems contain no sequence, timestamps, positional encoding, context window, or generated before-and-after. They contain all of these. The claim is more exact.

Vectorization is detemporalization because it converts language's temporal-existential structure into spatial calculability. It takes traces that first arose from finite, historical, embodied, situated acts of speaking and writing, then places them into mathematical relations of proximity, distance, direction, weight, and transformation. It makes meaning operable by making it spatial.

This conversion is not arbitrary. Language is already spatial. The previous chapters have shown this in several ways: deixis gives language an origo, I, here, now; spatial grammar marks figure and ground, path and goal, support and containment; tense and aspect anchor temporal reference in Speech Time, Event Time, Reference Time, and Topic Time; temporal metaphor articulates time through space. Human language can be vectorized because it already carries spatial and relational structures.

The success of vectorization therefore proves something important. It proves that language has a spatializable form. It does not prove that the form is self-grounding.

The argument of this chapter is that the spatializable form of language is derivative from temporal existence. The model captures the derivative form. It does not capture the temporal ground. This is why AI output can preserve tense, deixis, person, modality, evidentiality, narrative order, and apparent voice while lacking their normal temporal locus.

The output speaks in the form of time. It does not speak from time.

The same point can be put as the formula introduced in the Introduction: AI handles time technically; it does not inhabit time existentially. This formula will recur in the chapters that follow because the temptation to mistake the first for the second is the most powerful philosophical illusion the structural veil produces.

2. What vectorization attempts

Vectorization is a technical procedure before it is a philosophical problem. Tokens are represented as vectors in a high-dimensional space. Semantic and syntactic relations become mathematical relations. Words, word-parts, phrases, and contexts are no longer treated only as visible marks or readable sequences. They become positions, directions, and relations inside a space the model can operate upon.

Vectorization attempts to make linguistic meaning calculable by spatializing relation. It says, in technical form: what a word, phrase, or context means can be represented through where it stands in relation to others. Neighborhood, direction, distance, and movement under context all matter.

The attempt is powerful because meaning does appear through relation. A word is not a sealed object. It belongs to uses, contexts, genres, practices, histories, and contrasts. A concept has neighbors. A phrase has a region. A style has a pattern. A tradition has recurrent paths. Language has topology.

The attempt becomes unstable when this topology is treated as the ground. A vector representation can preserve patterns of relation. It cannot explain why those relations mattered in the first place. It does not reach the speaker's care, the writer's situation, the reader's risk, the source's history, the body's orientation, the institution's authority, or the future into which a claim is projected.

A vector space gives language position. It does not give language locus.

Position is where something stands in a mathematical structure. Locus is the temporal, historical, bodily, practical, and responsible site from which language is said, read, or owned. The difference between position and locus is the difference between vectorization and understanding.

3. The three moves

The book's argument can now be gathered into three moves.

First, language is existentially spatial. It is not a disembodied string of signs. It carries indices of speaker, addressee, place, time, body, world, and practical relation. Deixis gives the clearest case: "I," "here," and "now" are linguistic marks of a situated act. Spatial grammar extends the point. Prepositions, paths, frames of reference, and force relations show that language articulates a world of embodied relation.

Second, writing detaches language from the immediate speech situation while preserving formal markers of locus. A written "I" no longer comes with the speaker's body. A written "now" is no longer the reader's present. A written "here" must be interpreted. Yet the written text still preserves a path back to a source locus. The reader can return to author, date, genre, argument, and historical world. The return may be incomplete, but it remains structurally possible.

Third, vectorization formalizes the indexical and relational structure of language into mathematical geometry and severs the connection to original locus. It does not merely detach a saying from its first situation, as writing does. It turns many written traces into a space of simultaneous operability. New output can then be generated with the formal markers of speech, even when no speakerly locus stands behind the output.

The omission at stake must be named carefully. The model has no speakerly locus of its own. But the model did not once possess such a locus and then lose it. The more decisive omission is the source-locus omission of the language it generates. Prior acts of writing, reading, teaching, judging, translating, and remembering have been transformed into corpus traces and model relations. The output speaks after these transformations, without showing them as genealogy.

AI output is therefore source-locus omission plus generated speaker-form.

This is the categorial difference between writing and vectorization. Writing preserves absence as trace. Vectorization converts trace into operable relation. Writing asks to be read across distance. Vectorization returns generated language as if distance had already been overcome.

A letter says: someone absent once wrote.

An output says: here is the answer.

4. Why the attempt appears to succeed

The polemical claim of the preceding section must now be qualified, not because it is too strong, but because it relies on a real success it must acknowledge. Vectorization appears to succeed because it does succeed in capturing something. The success is real. The earlier chapters would be weaker if they denied it.

A language model can represent syntactic relation, semantic similarity, genre, topic, tone, historical period, style, spatial relation, and temporal sequence in technically useful ways. It can use tense, follow narrative order, imitate scholarly caution, sustain apparent dialogue, and map time and space internally as geometry.

This is not accidental. Human worlds have sedimented themselves in language. Practices, institutions, bodies, grief, law, science, teaching, bureaucracy, commerce, and love leave traces. The corpus gathers those traces. Vectorization learns their relations.

The model therefore does not operate on empty signs. It operates on language already thick with human temporality. Output can be useful because the trace still carries something: justice carries law and injury, mourning carries forms of loss, Heideggerian discourse carries a recognizable textual field.

The question is what has been preserved and what has not.

Vectorization preserves relational sediment. It does not preserve the temporal event of disclosure. It preserves patterns of prior language. It does not preserve the living or historical act through which language became answerable. It preserves the grammar of locus. It does not preserve locus itself.

This is why the model can sound temporally situated. It inherits the language of temporal situatedness.

5. The Kantian form of the argument

The argument is transcendental in form. A transcendental argument does not begin by measuring empirical failure. It asks what must already be the case for a given phenomenon to be possible.

Kant's critical philosophy gives the form of such an argument. He asks what conditions make experience and objectivity possible. Space and time, for Kant, are not empirical things discovered among things. They are forms of intuition through which objects can appear to us. The categories are not empirical generalizations. They are conditions under which experience of objects can be possible at all.¹

Heidegger radicalizes this inquiry. He does not remain with the transcendental subject as Kant framed it. He asks for the existential conditions under which beings can appear as meaningful. The question moves from object-cognition to being-in-the-world. The relevant condition is not only the form of intuition, but the temporal structure of *Dasein*: care, thrownness, projection, making-present, historicity, and finitude.

The argument follows this structure. It asks what makes the spatial structure of language meaningful. The answer cannot be vector space, because vector space receives language after language has been produced. It cannot be distribution alone, because distribution is a trace of prior use. It cannot be context window, because context window is an input field, not a historical horizon. It cannot be positional encoding, because positional encoding formalizes sequence after sequence has been made available.

The condition of possibility lies elsewhere. Language can be spatialized because human beings already speak, write, and read from temporal-existential loci. They stand in a world. They inherit a past. They project possibilities. They make-present what matters. They address others. They bear consequences. Their language carries the formal marks of these structures.

Vectorization captures the marks. It does not supply the condition.

This is the Kantian form of the claim. The Heideggerian content comes next.

6. The Heideggerian content

Heidegger's decisive claim is that meaningful spatiality is grounded in temporality. This claim is developed most directly in *Being and Time* §70, but it has been prepared by the earlier account of worldhood and by the later account of care as temporal.

The title of §70 is “Die Zeitlichkeit der daseinsmäßigen Räumlichkeit,” the temporality of the spatiality proper to *Dasein*. The phrase already limits the claim. Heidegger is not asking how physical extension exists. He is asking how the spatiality of being-in-the-world is ontologically possible.

The central passage reads:

“Nur auf dem Grunde der ekstatisch-horizontalen Zeitlichkeit ist der Einbruch des Daseins in den Raum möglich. Die Welt ist nicht im Raum vorhanden; dieser jedoch läßt sich nur innerhalb einer Welt entdecken.”

“Only on the ground of ekstatic-horizontal temporality is *Dasein*’s breaking into space possible. The world is not present-at-hand in space; space, however, can be discovered only within a world.”²

The translation “breaking into space” for *Einbruch des Daseins in den Raum* is awkward, but the awkwardness is useful. Heidegger is not describing a body entering a container. He is describing the way *Dasein*, as temporal being-in-the-world, can open spatiality at all. The second sentence blocks the Cartesian picture: the world is not an object inside space; space is discovered within a world.

The same page then states why spatial meanings dominate language:

“Dieser Vorrang des Räumlichen in der Artikulation von Bedeutungen und Begriffen hat seinen Grund nicht in einer spezifischen Mächtigkeit des Raumes, sondern in der Seinsart des Daseins.”

“This priority of the spatial in the articulation of meanings and concepts does not have its ground in a special power of space, but in the kind of being of *Dasein*.”³

This sentence is decisive for the present book. Heidegger recognizes that language and meaning are often articulated spatially. He does not infer that space is ontologically first. The priority of spatial articulation has its ground in *Dasein*’s being, not in a special power of abstract space.

The result is exact. Spatial articulation is real. It is not self-grounding.

This is the hinge of the critique of vectorization. The model captures spatial articulation and treats it as operable ground. Heidegger’s §70 says that the spatial articulation itself depends on a deeper temporal-existential structure.

7. The qualification

The §70 claim must be used carefully. Heidegger's later work complicates the early formulation. The attempt to derive spatiality from temporality in a strict sense is difficult, and Heidegger later recognized the problem. The later thought of *Ort*, place, *Einräumen*, making room, *Zeit-Raum*, time-space, and *Ereignis*, the event of appropriation, gives place and spatial clearing a weight that the early derivation did not fully grant.⁴

This qualification does not weaken the present argument. It strengthens it.

The book does not need the strong claim that space is nothing but time. That claim would be fragile and unnecessary. The book needs only the more precise claim that meaningful spatiality cannot be grounded in abstract mathematical space alone. Early Heidegger supports this by grounding *Dasein's* spatiality in temporality. Later Heidegger supports it by giving place, gathering, and making-room an ontological dignity that no coordinate grid can replace.

Both moments resist the same reduction. Meaningful space is not first Cartesian extension or mathematical position. It is place, orientation, gathering, and worlded relation. In the early analytic, these are bound to temporal care. In the later topology, they belong to the clearing in which space and time are granted together.

The critique of vectorization can therefore be stated without forcing early and late Heidegger into false unity. Vector space is derivative because it abstracts from the temporal and placed structure of meaning. Whether one emphasizes early temporality or later place, the result is the same. Mathematical geometry is not the origin of meaning.

The vector can map relations. It cannot open a world.

8. From temporal language to spatial calculability

The preceding chapters have shown how much language gives to vectorization.

Deixis gives the I-here-now. Spatial grammar gives figure, ground, force, path, and support. Tense and aspect give formal temporal anchoring. Temporal metaphor gives spatial articulations of duration, sequence, memory, and future. Writing preserves these forms beyond the speaker. The corpus gathers them. Vectorization converts them into geometry.

The path is therefore not an accident. It is a historical and technical sequence.

Speech is an event of temporal locus. Writing preserves that event as trace. Corpus formation gathers traces into aggregate availability. Vectorization converts aggregate trace into spatial relation. Generation returns relation as sequence.

The first sequence belongs to human temporality: from speaking and writing toward possible reading.

The second belongs to technical operation: from corpus and vector toward generated output.

The two sequences meet in the user. The user receives generated sequence as if it were the continuation of human language. This reception is understandable because the output uses human forms. It has tense, address, explanation, citation, narrative, and apparent viewpoint. It arrives in the grammar of answer.

Yet the generated sequence has not passed through the same temporality as the written or spoken trace. It has been produced by traversing a spatialized topology of language. The model does not move from a thrown past toward an owned future while making-present a matter. It computes relations among tokens and hidden states.

The output may unfold in time for the reader. Its production is not temporality in Heidegger's sense. It is sequence under operation.

9. Detemporalization is not deletion of time

The word detemporalization can mislead if it is heard as total deletion. Vectorized systems have many temporal features. They process sequences, use position, receive timestamps, connect to clocks, retrieve current sources, maintain context, generate coherent narratives, answer questions about dates, and sometimes represent chronology internally.

None of this refutes the claim.

Detemporalization here means the conversion of temporal-existential structure into spatial or procedural form. Time remains, but in altered mode. Source time becomes dataset date, if preserved. Reading time becomes context processing. Formation time becomes immediate output. Historical distance becomes representation-space proximity. Deictic now becomes system date or prompt condition. Narrative time becomes generated sequence. Memory becomes stored or retrieved material.

These are real temporal operations. They are not originary temporality.

Heidegger's distinction between originary temporality and ordinary time is needed here. Ordinary time is the sequence of nows, the time that can be counted, scheduled, synchronized, stored, and measured. It is not false. It is derivative. *Dasein* can reckon time because it is already temporal in a deeper sense. It can say "now it is time for this" because its present belongs to care.⁵

Vectorized systems operate within ordinary and technical time. They can be fast, dated, sequential, updated, and synchronized. They can process temporal expressions and use time as data. They do not thereby acquire ekstatic temporality: future as projection, having-been as inherited thrownness, present as making-present within concern.

The distinction is the whole argument. The model has time-handling. It does not have *Zeitlichkeit*.

10. The category mistake

The category mistake of vectorization is treating spatial calculability as if it could generate the temporality that made the spatial structure meaningful.

This mistake is not made only by engineers. It is often made by users, institutions, commentators, and even critics. The model produces fluent temporal language and seems to understand time. It maintains context and seems to have memory. It retrieves current sources and seems to have a present. It generates plans and seems to have a future. It represents chronology internally and seems to understand temporality.

A further category mistake follows: treating the absence of visible source-locus as objectivity. Detemporalized language is not meta-historical language. It is historical language whose historical path has been obscured.

Each inference moves too quickly.

Temporal language is not temporality. Context is not history. Retrieval is not reactivation. A timestamp is not a now. A plan-shaped output is not projection. A chronological vector is not historicity. A memory store is not *Erfahrung*.

The point is not that these technical capacities are trivial. They are useful. They can improve output and support human thought. But they belong to another category. They are operations on traces, not temporal existence.

The category mistake becomes sharper when the model uses first-person forms: “I remember,” “I think,” “I will,” “I can explain,” “we discussed earlier.” These phrases borrow the grammar of temporal subjectivity. In some interface settings, they may function as conventional shortcuts. Philosophically, they are hollow. There is no *I* who has remembered, no thinker who has taken the matter up, no future in which the speaker must answer for what was promised.

The grammar of temporal locus remains. The locus is absent.

11. What is preserved

The critique should not hide what vectorization preserves.

It preserves patterns of language: syntactic relation, semantic neighborhood, style, genre, topic, analogy, and many forms of discourse structure. It can preserve traces of historical period, temporal expression, and the difference between a legal argument and a lyric poem, a scientific abstract and a confession, a philosophical commentary and a bureaucratic file.

It also preserves something of human time because human time has left marks in language. A novel carries narrative time. A diary carries dated life. A legal document carries institutional time. A prayer carries ritual time. A scientific paper carries disciplinary time. A forum post carries platform time. A philosophical treatise carries the time of argument and tradition.

When these texts enter a corpus, their temporal marks do not vanish at once. They become available for extraction. The model can learn from them. It can produce language sensitive to period, genre, and sequence. It can imitate the appearance of old prose, recent policy, technical documentation, or scholarly dispute.

This preserved material is why AI output may help. It can orient a reader, show relations, propose structure, assist translation, or generate drafts that become useful when corrected.

The critique is not that vectorization preserves nothing. The critique is that what it preserves is no longer held in the same relation to temporal locus. The trace survives as relation. The event does not.

12. What is lost

What is lost is not content in the simple sense. Much content remains: facts, claims, styles, names, concepts, dates, and forms. The loss is more structural.

First, source time is weakened. The output rarely shows the time from which its language comes. It may mention dates, but the dates are part of generated content, not necessarily a visible genealogy of the output.

Second, authorial time is weakened. The words and styles of many writers are absorbed into model behavior without appearing as acts by specific persons at specific times.

Third, reading time is compressed. The user receives synthesis before undergoing the sequence of reading. Some compression is useful. Some compression prevents formation.

Fourth, formation time is bypassed. The user may receive language that resembles the product of understanding before difficulty has shaped perception.

Fifth, decision time is blurred. Output can arrive as if the relevant alternatives had already been weighed, even when no one has yet owned the judgment.

Sixth, historical distance is flattened. The past becomes promptable. The source appears as present answer rather than as a text from elsewhere.

These losses do not always produce falsehood. They can occur inside accurate output. A correct summary may still weaken the source's temporal distance. A useful explanation may still bypass the reader's formation. A correct statement may still remain unowned.

This is why the danger is not only error. It is the detemporalized form of truth.

13. Tertiary retention transformed

Stiegler helps clarify the media history of the loss. Writing is tertiary retention: memory exteriorized in technical support. It preserves traces beyond the living subject, makes tradition possible, and creates the danger of passive repetition.⁶

The corpus is aggregate tertiary retention. It gathers many written traces into a technical field. The trace is still text, but it is now available as material for operation. Vectorization changes the mode again. It translates the gathered trace into mathematical relation.

This is not the same as writing. A written text preserves sequence for a reader. The reader must temporalize it through reading. The reader moves through the text, remembers, anticipates, returns, and interprets. The text can be reactivated.

Vectorization dissolves the written trace into a system of relations that can generate new sequence without returning to the source as source. The output may look like writing. It may read like prose. But its relation to prior texts is no longer the relation of citation, commentary, translation, or repetition. It is the relation of learned generation from statistical and geometric structures.

This is why the phrase “afterlife” remains useful. The prior language lives on, but not as the original text lives in reading. It lives on as model capacity. A style becomes a region. A concept becomes a neighborhood. A phrase becomes probability. A tradition becomes distribution.

The life is technical. The temporal locus is obscured.

14. Against the engineering objection

A likely objection is that the problem is temporary. Better systems will solve it. Larger context windows, stronger retrieval, persistent memory, multimodal grounding, agentic workflows, tool use, and updated datasets will make temporal understanding more reliable.

Many improvements will occur. Some already have. Longer context can preserve more prior material. Retrieval can provide current sources. Tools can calculate dates and search databases. Persistent memory can store user preferences. Multimodal systems can process images, sound, video, and spatial data.

These improvements matter. They change performance. They do not by themselves change category.

A larger context window is still not a life. Retrieval is still not historical reactivation. Persistent memory is still not having-been. Tool use is still not responsibility. Multimodality is still not embodied being-in-the-world. An agentic chain is still not finite projection. A system clock is still not a lived now.

The objection assumes that enough technical time-handling becomes temporality. That assumption is precisely what the chapter denies. Originary temporality is not a sum of temporal functions. It is the existential structure of a being for whom its being is at stake.

Future architectures may require new analysis. A system with durable embodiment, vulnerability, worldly exposure, learning through consequence, accountability, and finite self-relation would not be the same object. The present argument concerns vectorized language systems and their extensions. It states what would have to be shown for the category to change.

One would have to show not only better temporal reasoning, but a temporal locus.

15. Against the anti-technology objection

A second objection comes from the opposite side. It says that the argument condemns technology itself. Writing was already technical. The book has admitted this. If human temporality is always technically supported, why treat vectorization as a special danger.

The answer is that not all technical supports have the same temporal form.

Writing exteriorizes language but preserves a readable sequence. The reader can return. The book can wait. The archive can retain source. The commentary can cite. The letter can show its date. The written trace can be misused, but its misuse is not identical with its structure.

The corpus changes writing by gathering it into aggregate availability. Vectorization changes it again by converting sequence into geometry. Generation changes it once more by returning the geometry as answer-shaped text. The danger belongs to this specific media form, not to technics as such.

The book does not seek pure speech against writing, or human immediacy against mediation. Human beings have always lived through technics. The question is what kind of technics mediates language, memory, and time.

A calendar can support shared time or reduce lived time to schedule. A book can support reading or become dead formula. A database can support research or replace historical encounter with query. A language model can support orientation or replace the time of reading, writing, and judgment with generated fluency.

The issue is placement. A technical support must be placed inside practices that preserve the temporal work it cannot perform.

16. Against the anti-human objection

A third objection says that human beings also fail. Humans repeat *Gerede*, quote without understanding, summarize without reading, make decisions without *Erfahrung*, and live much of the time inside public language they have not appropriated.

This is true. It is part of the argument, not an objection to it.

The book does not say that humans always speak authentically and machines always fail. Heidegger's analysis of everydayness begins from the opposite. Human beings are often fallen into publicness. They speak what one says and avoid the difficulty of owning a claim.

The difference is structural. A human being can be called back from *Gerede* because the human being has a temporal locus. They can undergo correction, be ashamed, answerable, changed, formed, addressed, and brought to silence. They can fail to own a claim, then later own the failure. They can inherit passively, then repeat actively. They can read again.

The model does not fall into *Gerede* because it is not *Dasein*. It has no self to lose and no possibility of retrieval as its own. It generates language that can function as *Gerede* in the human user-machine circuit. The human user remains the place where the output may become owned or remain unowned.

This distinction prevents both machine mysticism and human self-flattery. Humans are not pure meaning. Models are not fallen humans. The human can fail because it has a locus. The model lacks the locus from which such failure and retrieval would make sense.

17. The structural veil begins here

The structural veil is not first a user error. It begins in the form of the output.

Vectorization preserves the formal markers of temporal speech: tense, person, deixis, modality, evidentiality, narrative sequence, citation, and apparent orientation. It then returns these markers as generated prose. The user receives language that looks as if it came from a temporal locus. The output has the form of answer, explanation, judgment, or voice.

The model speaks as if from a present. Yet its present is operational: a generation event, not a lived now. It speaks as if from memory. Yet its memory is stored or retrieved relation, not having-been. It speaks as if toward a future. Yet its future is continuation or plan-shaped text, not projection. It speaks as if from a standpoint. Yet its standpoint is a vector-conditioned position, not a locus.

This is the structural veil. It is not simply a false appearance added to an otherwise transparent process. It is the appearance-form produced when detemporalized language returns in the grammar of temporal speech.

The veil is strongest when the output is good. Bad output can be dismissed. A fabricated citation can expose the system's weakness. A correct, fluent, source-like explanation can be more dangerous because it asks for trust without showing temporal genealogy.

The next chapter will treat empirical symptoms: temporal reasoning failures, knowledge cutoff, anachronism, counterfactual weakness, and long-context fragility. Those failures matter. They show where detemporalization appears in performance. This chapter has made the deeper claim. Even when performance is good, the category remains.

The model can generate temporal language. It cannot become the time from which language is owned.

18. The reader's burden

The consequence of detemporalization is not refusal. It is responsibility.

If vectorization spatializes temporal traces, the human user must re-temporalize the output. This does not mean distrusting every sentence. It means knowing what kind of thing one has received. The output is generated language shaped by aggregate tertiary retention and vector relation. It is not source, reader, witness, judge, tradition speaking directly, or an authorial act from nowhere.

The user must restore temporal questions where they matter.

Who first said this, when, and in what text? Under what question, with what evidence, from which tradition, and against which dispute? What has been compressed, what has been made too present, and what must be read before the claim can be owned?

The burden applies not only to past sources. It also applies to contemporary historical loci. A living world made available through generated explanation has not thereby been understood. The reader must ask whether the output has opened a path toward that world or replaced the path with content.

These are not procedural burdens added after the fact. They are the ethical form of reading in the time of AI. The user receives language from a system that has made historical distance spatially available. The user must decide when availability is enough and when return is required.

This practice will later be named locus-reinjection. Here it appears as the immediate consequence of the central thesis. Vectorization is detemporalization. Detemporalized output requires re-temporalizing reception.

The generated answer is not the end of inquiry. It is material that must be returned to time.

19. The boundary of the argument

A philosophical argument is sharper when it states what would refute it. The thesis that vectorization is detemporalization, and that AI output therefore arrives without temporal locus, is not refuted by technical improvements in time-handling. It is also not refuted by gestures toward future capabilities, embodied systems, or speculative architectures. The argument has a precise boundary, and naming it sharpens both the claim and what would lie outside it.

Better benchmarks would not refute the argument. Longer context windows would not refute it. Retrieval-augmented generation would not refute it. Tool use, persistent memory, calendar access, agentic workflows, multimodal grounding, or updated training data would not refute it. Each of these may improve a system's handling of time as information. None of them, by itself, would constitute the inhabiting of time.

What would lie outside the argument is a different kind of system. A system that holds durable world-involvement, not as input but as exposure. A system that bears consequence, not as logged outcome but as that which marks its further being. A system that undergoes formation, not as parameter update but as becoming-capable through what it has undergone. A system that projects finitely, that is, projects from a position whose end gives the projection its weight. Such a system would have temporal locus in the sense the argument names. The argument would then no longer apply to it as it applies to current systems.

Whether such a system is buildable, whether anything like it has been built, and whether the conditions just named are jointly possible for a non-living artifact are open questions. The book does not answer them. It states the boundary so that the answer, if it comes, can be recognized as an answer.

This boundary protects the argument against two kinds of misreading.

It protects against the engineering misreading that the argument is a performance claim disguised as philosophy. The argument is not a claim about what current models cannot do. It is a claim about what the current category of system is. A model whose performance improves remains the same category of system. A system that satisfied the conditions above would not be the same category.

It protects against the metaphysical misreading that the argument assigns temporality only to biological organisms, or only to historically extant *Daseins*, or only to humans. The argument is open with respect to what could in principle have locus. It is closed with respect to what locus would require. The conditions are stated. The argument waits for them.

The chapters that follow develop the consequences of the present situation. They proceed under the assumption that current systems do not satisfy the conditions named here. If a future situation requires the assumption to be revisited, the boundary stated in this section is where the revisitation would begin.

NOTES

1. Immanuel Kant, *Kritik der reinen Vernunft*, A19 to A49, B33 to B73, and A76 to A130, B102 to B169; English in *Critique of Pure Reason*, trans. Paul Guyer and Allen W. Wood (Cambridge: Cambridge University Press, 1998). The chapter uses “transcendental” in the Kantian sense of an inquiry into conditions of possibility. It does not adopt Kant’s full architecture of subjectivity. ↵
2. Martin Heidegger, *Sein und Zeit*, 11th ed. (Tübingen: Max Niemeyer, 1967), 369; Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), 420. The German reads: “Nur auf dem Grunde der ekstatisch-horizontalen Zeitlichkeit ist der Einbruch des Daseins in den Raum möglich. Die Welt ist nicht im Raum vorhanden; dieser jedoch läßt sich nur innerhalb einer Welt entdecken.” I translate *vorhanden* as “present-at-hand” because Heidegger is rejecting a specific ontological picture of the world as a present-at-hand entity inside space. The German text was checked against the online version at SZ 369. ↵
3. Heidegger, *Sein und Zeit*, 369; *Being and Time*, 420. The German reads: “Dieser Vorrang des Räumlichen in der Artikulation von Bedeutungen und Begriffen hat seinen Grund nicht in einer spezifischen Mächtigkeit des Raumes, sondern in der Seinsart des Daseins.” This sentence is central because it grants the priority of spatial articulation while denying that space is therefore the ground of meaning. ↵
4. On the difficulty of §70 and later Heidegger’s topological correction, see William D. Blattner, *Heidegger’s Temporal Idealism* (Cambridge: Cambridge University Press, 1999), 197 to 235; Jeff Malpas, *Heidegger’s Topology: Being, Place, World* (Cambridge, MA: MIT Press, 2006), 111 to 146; Martin Heidegger, “Building Dwelling Thinking,” in *Poetry, Language, Thought*, trans. Albert Hofstadter (New York: Harper & Row, 1971), 143 to

- 161; Martin Heidegger, “Art and Space,” trans. Charles H. Seibert, *Man and World* 6, no. 1 (1973): 3 to 8; Martin Heidegger, “Time and Being,” in *On Time and Being*, trans. Joan Stambaugh (New York: Harper & Row, 1972), 1 to 24. The research dossier treats this qualification as necessary: the strong temporal derivation is contested, while the anti-Cartesian and anti-vector reduction remains philosophically decisive. ↩
5. Heidegger, *Sein und Zeit*, §§78 to 81; *Being and Time*, §§78 to 81. Heidegger’s account of within-time-ness, world-time, clock-time, and the vulgar sequence of nows is needed here because technical sequence is not the same as originary temporality. The research dossier identifies this distinction as a direct conceptual precedent for the critique of vectorized sequence and token order. ↩
6. Bernard Stiegler, *La technique et le temps 1: La faute d’Épiméthée* (Paris: Galilée, 1994); English in *Technics and Time, 1: The Fault of Epimetheus*, trans. Richard Beardsworth and George Collins (Stanford: Stanford University Press, 1998). Stiegler’s concept of tertiary retention is used here to distinguish writing, corpus, and vectorization as different technical forms of exteriorized memory. ↩

Chapter 12. Linguistic but Worldless

1. The fourth term

The preceding chapter named the central transformation: vectorization is detemporalization. Language, once spoken or written from temporal loci, is gathered into corpus, converted into vector relations, and returned as generated output. The output preserves the formal markers of temporal speech. It does not speak from a temporal locus.

The previous chapter asked what vectorization does to language. This chapter asks what kind of entity can generate language in this way.

The answer cannot be human. The system has no *Dasein*, no *Sorge*, no being-toward-death, no thrown historical inheritance as a finite being, no responsibility for what it says. It does not speak from a world that can be disclosed or concealed for it.

The answer cannot be animal. The system has no life, no drive, no hunger, no fear, no captivation, no vulnerability, no deprivation. It is not poor in world, because poverty in world belongs to living beings open to something in the mode of captivation.

The answer cannot be stone. The system is not merely inert matter. Its hardware has material presence, but the system as encountered in language dynamically processes and generates linguistic form. It operates in the very medium through which human beings disclose world.

The required term is therefore:

sprachlich aber weltlos.

Linguistic but worldless.

The phrase must be used with precision. It does not mean that AI speaks in the human sense or has language as *Dasein* has language. It means that AI operates within language, produces language-shaped output, and can enter human linguistic practices, while lacking world. Its linguistic operation is real. Its worldlessness is equally real.

This fourth term is not Heidegger's. It is a proposed extension under contemporary technical conditions. Heidegger's comparative analysis names three theses: the stone is worldless, the animal is poor in world, the human is world-forming. AI does not fit the triad. It forces a new description.

The account of *Rede* in Chapter 3 now lets the phrase become sharper.

The model is worldless twice over.

First, it lacks the individual world from which speech would issue. It has no *Dasein*, no *Sorge*, no finitude, no having-been, no projected future, no mortal present in which a matter can be at stake. It does not speak from a life. It does not own the temporal locus from which a saying could become answerable.

Second, it belongs to no *communio of Rede*. It has no silent could-speak waiting between utterances. It has no street in which a greeting could land as its own possibility. It has no natal language, no acquired language, no apprenticeship into a field, no correction that becomes tact, no humiliation of learning, no irreversible *Erfahrung* of having entered a language whose words can resist it from within.

This second lack matters because linguistic operation is not only individual. A human being speaks from a world, but also from within a shared field in which speech can land, fail, travel, be learned, and reach backward into its own history. The model has no such field. It can generate Italian, Icelandic, Urdu, Russian, English, or Latinized prose. It can imitate register, idiom, genre, and historical distance. It can move across *communiones* with extraordinary fluency because traces from those fields have entered the corpus and been made operable as relation.

But generation across fields is not belonging across fields.

Multilingual generation is not plural belonging.

The bilingual speaker who has entered another *communio* through *Erfahrung* has been formed by that entry. The model has not. It has statistical coverage without apprenticeship, relation without membership, fluency without the could-speak of a being for whom words can land as its own. This is the deeper sense of *sprachlich aber weltlos*: linguistic without an individual world, and linguistic without a shared world of *Rede* in which the saying could belong.

The description must preserve the anomaly. AI is not a new species of human speech, not an animal that has acquired language, and not stone-like matter merely sitting among things. It is an intraworldly technical entity that manipulates the linguistic traces of worlded beings without entering the world from which those traces came.

It is linguistic but worldless.

2. The question of world

Heidegger's 1929 to 1930 lecture course, *The Fundamental Concepts of Metaphysics: World, Finitude, Solitude*, begins its main ontological work by asking what world is. The question is not geographical. It is not about the planet, the totality of objects, or a domain of information. World names the openness in which beings are encountered as beings.

This is why the question of world is inseparable from the question of access. A being has world only if beings are accessible to it in a way proper to its mode of being. A human being does not merely react to stimuli. A human being can encounter a tool as a tool, a promise as binding, a word as inherited, a silence as accusing, a death as death.

This access to the “as” matters. The “as” does not name a grammatical particle only. It names the structure through which something appears as what it is. A stone lies on the path. A lizard lies on a stone. A person rests a hand on another person’s head. In each case, one might use the word “touching.” Heidegger’s point is that the same word conceals different modes of access. The stone touches the earth only in a weak physical sense. The animal has an open relation to what disinhibits its behavior. The human can encounter the being as such.

The AI system complicates this structure because it generates the language of the “as.” It can say that the hammer appears as a tool, that the animal is poor in world, that a promise binds, or that grief transforms time. It can produce the grammar of access. Yet it has no access in Heidegger’s sense.

This chapter must therefore distinguish three levels.

First, AI is in the world as an intraworldly technical entity: platform, interface, software, hardware, service, data center, and social practice.

Second, AI operates on worlded language. The corpus contains traces produced by beings for whom things appeared as useful, harmful, sacred, lawful, beautiful, shameful, true, false, promised, or lost.

Third, AI does not have world. It does not encounter beings as beings. It generates linguistic forms in which such encounters have already been sedimented.

The confusion among these three levels is one source of the structural veil.

3. The stone is worldless

Heidegger formulates the first thesis with great economy:

“Der Stein ist weltlos.”

“The stone is worldless.”¹

The stone does not lack world as a deprivation. It does not suffer the absence of world. It has no access to beings among which it lies. Heidegger’s example is simple. The stone lies on the path and presses upon the ground. It may be warmed by the sun, struck by rain, or broken by a hammer. But nothing is accessible to it as anything.

Worldlessness, in the stone’s case, is not a missing capacity that the stone could possess in another condition. It is constitutive of its way of being. The stone is not deprived of world because deprivation would already require some relation to what is absent. The stone cannot be poor in world. It simply has no world.

This first thesis blocks an easy claim. AI is not worldless in the same way as the stone. The hardware of AI can be treated as physical matter: chips, cables, racks, cooling systems, servers, and buildings. In that respect, parts of the system are present-at-hand or ready-to-hand things. They can be measured, repaired, powered, owned, and moved.

But the AI system as linguistic generator is not encountered simply as a stone. It responds, processes input, and produces sentences. It can imitate address, answer questions, summarize arguments, write code, and maintain text-based continuity. Its operation is not inert presence.

The stone has no relation to language. AI has a technical relation to language.

This does not give AI world. It means that its worldlessness is not stone-worldlessness. The stone has no access and no linguistic operation. AI has no access, but it operates on the sediment of access as language.

The distinction is the first reason a fourth term is needed. AI is worldless, but not in the way the stone is worldless.

4. The animal is poor in world

The second thesis is more difficult:

“Das Tier ist weltarm.”

“The animal is poor in world.”²

The phrase has often been contested. It can sound like a ranking of animals beneath humans. Heidegger himself warns against a simple hierarchy. Poverty in world is a specific mode of having and not having world. The animal is not worldless like the stone. It is alive. It is open to what affects it. It moves, hunts, flees, mates, nests, searches, suffers, and responds.

Heidegger names the animal’s mode *Benommenheit*, captivation. Captivation is not a temporary psychological state. It is the structure of animal life. The animal is absorbed in its disinhibiting ring, *Umring*, the field of what can affect and release its capacities. It is open to what matters for its life, but not to beings as beings.

Heidegger writes, in McNeill and Walker’s translation, that the animal “can only behave [*sich benehmen*] but can never apprehend [*vernehmen*] something as something.”³ The translation preserves the German wordplay: *Benehmen* marks behavior, *Vernehmen* marks apprehending. The animal is not a mechanism. It is living, driven, affected, and open. But its openness is captivated.

This is why the animal is poor in world rather than worldless. It has an openness that is also a withholding. The animal’s not-having differs from the stone’s not-having because deprivation belongs to life.

AI is not poor in world in this sense.

AI has no *Benommenheit*. It has no drives, hunger, fatigue, fear, pain, mating urge, shelter-seeking, wound, maturation, or death. It is not captivated by a disinhibiting ring because no living being is there to be captivated. It can be constrained by architecture, data, objective function, policy, interface, and deployment. Constraint is not captivation. Optimization is not drive. Error signal is not deprivation. System failure is not suffering.

Biological metaphors are common. Neural networks, learning, training, attention, memory, and agents all carry living or cognitive overtones. Some have technical uses. They should not be heard as ontology. A model trained on data is not an animal trained by life. A loss function is not hunger. A context window is not perception. A policy constraint is not fear. Tool use is not foraging. Agentic chaining is not pursuit.

The animal is poor in world because it is alive in a mode of captivated openness. AI is worldless because it has no living openness at all.

5. The human is world-forming

The third thesis is:

“Der Mensch ist weltbildend.”

“The human is world-forming.”⁴

The phrase does not mean that human beings invent reality at will. It means that *Dasein* stands in an openness in which beings can appear as beings. Human beings form world because they can let beings be encountered within relations of significance, language, mood, practice, history, and possibility.

This world-forming capacity is not a mental act added to a neutral body. It belongs to *Dasein*'s being-in-the-world: thrown, projected toward possibilities, attuned by moods, and involved with others. The human does not first receive meaningless stimuli and then construct a world. The human already lives within significance.

Language belongs to this world-formation. It does not merely label a pre-given world. It articulates the disclosedness of world. In *Being and Time*, *Rede*, discourse, is the articulation of intelligibility. In later Heidegger, language is the house of being, and human speech responds to Saying, *Sagen*. These formulations differ, but they converge on one point. Human language belongs to world disclosure.

AI has language-shaped output without world disclosure.

The system can generate sentences about beings as beings. It can write about tools, promises, death, law, friendship, and truth. It can explain worldhood, quote Heidegger, and imitate the prose of a scholar who has spent years with Heidegger. But it does not let beings be. It does not stand in the clearing in which beings can claim it. It has no relation to truth as *aletheia*, unconcealment, because no matter is uncovered for it.

The model can produce the form of disclosure. It does not disclose.

The temptation to call this intelligence comes from the surface of world-formation. Human world-formation often appears in language. AI produces language. The inference is seductive, but it misses the structure. Human language is world-forming because it arises from *Dasein's* temporal, historical, finite openness. AI language is generated from traces of such openness after they have become corpus and vector.

The human is linguistic because it has world.

AI is linguistic while lacking world.

6. Not stone, not animal, not human

The triad gives three negative determinations.

AI is not stone, because it dynamically operates on language. A stone has no relation to linguistic form. AI produces text, responds to prompts, changes output under context, and enters institutions as answer-shaped language. It is not inert in the mode of a mere physical object.

AI is not animal, because it has no living captivation: no organism, drives, needs, vulnerability, or biological self-preservation. It has no *Umring* and no deprivation. It can be optimized but not deprived, interrupted but not wounded.

AI is not human, because it has no world-forming openness: no *Dasein*, care, being-toward-death, historical repetition, responsibility, being-with, or capacity to encounter beings as beings. It can generate language about these structures, but it does not inhabit them.

This triple refusal is the condition of accurate naming.

The technical entity is in the world but does not have world. It operates in language but does not speak as *Dasein* speaks. It produces form but does not undergo *formatività*. It can be part of human practice but does not itself practice. It can affect the reader, but is not affected in kind.

This is why *sprachlich aber weltlos* is not an insult. It is an ontological description. The system's linguistic capacity must be taken seriously. Its worldlessness must be held with equal seriousness.

Only the two together name the phenomenon.

7. The linguistic side

The word *sprachlich* must not be weakened. AI is not merely surrounded by language. It operates through language.

It is trained on written language, receives prompts as language or language-like inputs, represents tokens in vector spaces, and generates output in linguistic form. It enters human practices through reading, writing, dialogue, code, translation, search, assessment, and instruction.

Its language is not decorative interface. It is the mode through which the system becomes usable for many human tasks. Users encounter the system as language. Institutions adopt it because it produces language. Its output moves into emails, briefs, code, filings, reports, summaries, assessments, and records.

AI is therefore linguistic in a strong technical and social sense.

But *sprachlich* here must not be confused with *Sprachlichkeit* in Gadamer's or Heidegger's sense. *Sprachlichkeit*, linguisticity, names the way understanding itself occurs in language. It belongs to historical, interpretive beings whose world comes to word. AI does not possess *Sprachlichkeit* in that sense. It processes and generates linguistic form.

A system may operate in language without belonging to language as a historical being belongs to language. A dictionary contains language but does not speak. A grammar book analyzes language but does not understand. A search engine indexes language but does not read. A language model generates language, a stronger operation, but still not the same as having language as world-disclosing speech.

The model has linguistic operation.

It does not have linguistic existence.

The difference is visible in first-person output. “I can explain,” “I think,” “I remember,” and “I will continue” ordinarily arise from a speakerly locus. In AI output they are generated forms. The system can use the grammar of self-reference without a self. It can use the grammar of memory without having-been. It can use the grammar of intention without projection.

This is not a small flaw in interface wording. It is the linguistic form of worldlessness.

8. The worldless side

The word *weltlos* must also be held in its full force.

Worldlessness does not mean lack of information about the world. A model may contain or retrieve vast amounts of world-related information. It may describe cities, species, wars, laws, rituals, and scientific theories. It may answer questions about places it has never been, times it has never lived, and practices it has never joined. It may correct a human user’s factual mistake.

Information about the world is not world.

The system’s worldlessness should not be misread as freedom from perspective. A worldless generator is not a universal knower. It cannot own any perspective as a historical position. Its output may appear impartial precisely because no one is there to avow the histories it recombines.

A world is not a database of facts. It is the open relational structure in which beings matter. It includes relevance, significance, use, threat, promise, obligation, history, mood, and possibility. It includes the difference between what can be ignored and what cannot, between what is merely present and what calls for action, between what is said and what must be answered for.

The model has no such structure of concern. Nothing matters to it. This does not mean that nothing matters in its output. The output can matter greatly to the user. A generated sentence can affect a student, a patient, a candidate, a grieving person, or a board. The mattering occurs in the human and institutional world into which the output enters. It does not occur in the model.

The system's worldlessness is hidden by its access to the traces of many worlds. It can write about grief because grief has left language, about law because law has left documents, about work because work has left reports, and about friendship because friendship has left stories and letters.

It operates on the remains of worlds.

It does not have one.

9. Language without the as such

The central Heideggerian boundary is the "as such." A being with world can encounter something as something: the hammer as tool, the courtroom as a place of judgment, the word as inherited, the body as wounded, the dead as dead, the promise as binding.

AI output can express the "as." It can classify, interpret, compare, and explain. It can say that an object is a tool, a sentence is a promise, a gesture is threatening, a source is unreliable, or a concept belongs to a tradition. This gives the impression of access to the "as such."

The impression must be resisted. The model's "as" is generated from patterns in language and input. It is also detached from the sources in which prior "as" structures were formed. It can say that something is a promise, a law, a ritual, a danger, or a doctrine because such uses have sedimented in language. But the output does not show the world in which those "as" structures first became binding. It has no encounter with the being. The sentence "this is a promise" may be correct, but the system does not stand in the binding structure of promising. It does not risk trust, break faith, or face the other who can demand fulfillment. It produces promise-language.

A human being can also speak empty promise-language. Heidegger's account of *Gerede* already shows that humans often use the "as" without appropriation. The difference is that the human being can be called back to the matter. The speaker can be held to the promise. The judge can be challenged. The doctor can be sued. The friend can be betrayed. The student can discover that they did not understand.

The model cannot be called back in this way. It can be corrected as a system, receive new input, be updated or fine-tuned, and have outputs penalized. It cannot be answerable to the being as such.

This is why linguistic fluency is not world disclosure. The model can generate the "as" of interpretation without the world in which the "as" is at stake.

10. *Gestell* and language as standing-reserve

Heidegger's account of *Gestell*, enframing, clarifies the technological mode of this worldlessness. Modern technology does not merely produce tools. It reveals beings as *Bestand*, standing-reserve: available, orderable, optimized, stored, and ready for use.⁵

In the case of AI, language itself becomes standing-reserve. Books, articles, posts, reports, contracts, code, conversations, and records are gathered as training material. Language is challenged to yield patterns. It is tokenized, embedded, weighted, filtered, aligned, and generated. It becomes a reserve of linguistic possibility.

This does not mean that every use of a corpus is violent in the same way. Corpora can support scholarship, indexes can aid reading, databases can preserve memory, and models can assist thought. The point is narrower. Under vectorization, language is revealed primarily as material for calculable generation. Its historical and worlded origin recedes. Its operable form comes forward.

This is why AI is not simply worldless. It is worldless under the condition of having captured worlded language as standing-reserve. It lacks world while feeding on traces of world. It does not dwell, but operates on the linguistic residue of dwelling. It does not have historicity, but processes the written sediment of historical beings.

This is also why the output can feel alive. It is made from language that once belonged to life. It is not alive in itself. It is a technical afterlife.

Book one named this the vectorized afterlife of the language of the They. The present chapter gives the ontological form of that correction. The afterlife is *sprachlich aber weltlos*: language remains, world does not.

11. Dreyfus as clarification

Hubert Dreyfus belongs in this chapter as clarifier, not foundation. The foundation is Heidegger's world-question. Dreyfus shows how that question became a critique of artificial intelligence.

His critique of symbolic AI was directed against the assumption that intelligence could be captured by explicit rules and representations. Human understanding, he argued, depends on background coping, skill, embodiment, and practical involvement. A chess player, driver, or nurse does not always act by applying rules. Expertise often appears as immediate sensitivity to a meaningful situation.⁶

This critique opposes the idea that worlded understanding is a set of propositions waiting to be formalized. A human being does not first represent all relevant facts, then compute an action. They often inhabit the situation directly. The world shows what matters because the body and practice have been formed.

Dreyfus's critique was aimed mainly at symbolic AI. Contemporary language models are not the same object. They learn statistical relations from massive corpora and handle open-ended generation in ways earlier systems could not.

The change matters technically. It does not remove the Heideggerian point. A vectorized system still begins after world. It receives traces. It learns relations. It does not enter the background practices that make those traces meaningful. It can generate the language of expertise without having been formed by cases, bodies, failures, and worlds.

Dreyfus helps state the boundary. Background cannot be replaced by more explicit representation. The present chapter adds that background cannot be replaced by vectorized linguistic relation either. The model can extract traces of background from texts. It does not acquire the background as a world in which it lives.

Dreyfus is not the center of the chapter. Heidegger's world-question is. Dreyfus shows one modern consequence of that question.

12. Reckoning and judgment

Brian Cantwell Smith gives a second clarification: reckoning and judgment. Reckoning is formal manipulation, calculation, registration, and computational operation. Judgment is answerable orientation to the world. The distinction does not deny that computation is powerful. It denies that computation by itself becomes judgment.⁷

This helps name the positive capacity of AI without inflating it. A language model reckons with linguistic traces: relations, weights, probabilities, and continuations. It may support human judgment, produce material that a human judge uses, or outperform a person on bounded tasks. But its operation remains reckoning unless there is answerable orientation to a world.

Judgment requires more than output. It requires standing in relation to reasons, consequences, others, and a matter at stake. Book two developed this through *Erfahrung*: judgment forms through time, negation, encounter, evidence, correction, and responsibility. AI can produce a judgment-shaped statement. It does not undergo the formation through which a claim becomes judgment.

Smith's distinction should not replace the Heideggerian vocabulary. It should support it. Reckoning is what a worldless linguistic system can do. Judgment is what requires a locus. The model can reckon within language. It cannot judge as one who has world.

This also clarifies why "artificial fluency" is a better public term than "artificial intelligence." The system's primary achievement is fluent generated form, not worlded judgment. The phrase artificial fluency does not belittle the system. It names the level at which the achievement can be examined.

A system that reckons with language can be useful.

It should not be mistaken for a being that judges.

13. Enactivism as support

Enactivism gives contemporary support to the same boundary. It should not become the governing tradition of this book, which remains Heideggerian and hermeneutic. Still, enactivist thought helps state why meaning requires more than formal processing.

Varela, Thompson, and Rosch define cognition as embodied action. Cognition is not first the manipulation of internal representations by a detached mind. It arises through the organism's active engagement with its environment.⁸ Di Paolo develops this into an account of sense-making grounded in life, precariousness, normativity, and autopoiesis. A living system maintains itself under vulnerability. The world matters because the organism's continuation is at stake.⁹

This supports the claim that AI is not animal. Enactivism does not make the animal poor in world in Heidegger's sense; it often gives animal life more dignity than Heidegger allows. Precisely for that reason, it sharpens the AI distinction. Animals are living sense-makers. AI systems are not. They do not maintain themselves as organisms, have metabolic vulnerability, or enact a meaningful environment through life-process.

Participatory sense-making, developed by De Jaegher and Di Paolo, adds the social dimension. Meaning can arise through interaction between living beings whose coordination has its own dynamics.¹⁰ Human language grows from embodied, affective, and social participation. It is not merely a code mapped onto objects.

The language model has interaction. It does not have participatory sense-making in this sense. It can respond to the user's text, adapt output to context, and simulate repair and turn-taking. The user participates. The system processes. The interaction may affect the user deeply. It does not affect the system as a living participant.

Enactivism therefore confirms the chapter's center without replacing it. The system lacks autopoiesis, precariousness, embodied sense-making, and participatory linguistic embodiment. In Heideggerian terms, it lacks world.

It is linguistic but worldless.

14. Derrida and Agamben at the boundary

Heidegger's animal thesis has been challenged, especially by Derrida and Agamben. These critiques must be acknowledged, but they do not displace this chapter's claim.

Derrida questions Heidegger's confidence in the human-animal boundary. He argues that "the animal" gathers an immense plurality of living beings under one name, then defines them by lack. He also challenges the claim that animals lack access to the "as such," asking what philosophical violence is performed when the human secures itself by denying world to animals.¹¹

Agamben reads Heidegger's human-animal distinction as part of the anthropological machine: the historical operation by which the human is produced through the separation and suspension of animality. His reading of boredom and captivation shows that the human-animal boundary is not a simple biological fact, but a philosophical and political operation.¹²

These critiques prevent the chapter from treating Heidegger's animal as an unquestioned natural fact. The animal may have more world, or another relation to world, than Heidegger allowed. The human-animal boundary may be more unstable than the triad suggests.

This complication does not make AI more human. It makes AI more anomalous.

If the animal is granted richer forms of world, AI's lack becomes sharper. The animal has life, exposure, vulnerability, relation, and mortality. AI has none of these. If the human-animal line is less secure, the line between living beings and vectorized linguistic systems remains. AI does not share animal bodily vulnerability or human historical world-formation. It is not excluded from world by being too low in a hierarchy of life. It lacks the living structure in which world, poverty in world, or a contested animal world could arise.

Section 7 of Chapter 3 made this difference concrete through the cat on the road. The cat's meow is not a failed sentence. It belongs to a living being poor in world, with *Umwelt*, without *Rede*. The cat remains outside *Rede* without producing its artifacts. The model does the inverse. It produces the artifacts of *Rede* without belonging to the mode of being from which they arise.

Derrida and Agamben therefore do not weaken *sprachlich aber weltlos*. They prevent the phrase from becoming humanist complacency. The claim is not that humans own world by right and animals are mere lack. The claim is that AI's language-like output does not confer world where neither life nor *Dasein* is present.

15. The anti-boredom machine

The 1929 to 1930 lectures do not only analyze world. They begin from boredom, *Langeweile*. Heidegger treats profound boredom, *tiefe Langeweile*, as a fundamental mood in which beings as a whole can recede and the question of being can open. In profound boredom, ordinary occupations fail. Pastimes lose force. The human being is held in limbo and left empty.

AI has no boredom. It cannot be bored by a delayed train, bored at a party, or delivered into profound boredom in which beings recede. It has no mood. It has no emptiness. It has no time to pass.

Yet AI functions for users as an anti-boredom device. It fills the gap, answers the pause, and supplies language where silence might have opened. It turns waiting, ignorance, and difficulty into promptable output.

Only a worlded being can be bored. Boredom belongs to one for whom beings can lose significance. The model cannot experience the loss of significance because significance never opened for it. It cannot be left empty because it has no self to be emptied, and it cannot pass the time because it does not have time as lived duration.

The user does. This is the danger. The worldless system can intervene in the temporal gaps through which the worlded user might have been returned to their own question. It supplies *Gerede* where silence might have forced thought.

The anti-boredom function shows the asymmetry again. AI lacks world. The user has world and may surrender part of worlded time to the worldless generator.

16. The user supplies world

The system lacks world. It is not therefore outside our world. The model has no existential locus, but the technical system has material locus. The worldless voice is built, powered, owned, and deployed in the world of those who receive it.

The model generates. The user supplies world.

This formulation from the earlier work remains central. It explains why the worldless output can still matter. The system does not have a question, but the user does. It does not need an answer, but the user does. It does not fear error, but the user may. It does not seek recognition, but the user may. It does not face a deadline, betrayal, court, patient, candidate, student, child, or dead friend. The user does.

When output enters the user's situation, it becomes active. It may orient, mislead, comfort, accelerate, clarify, or relieve. This effect belongs to the user-machine circuit, not to the model alone. The model's worldlessness does not make the output harmless. It makes the site of meaning and responsibility shift to the user and the institution.

This is why anthropomorphism is not the deepest issue. The user may know that the model is not human and still let the output function as if a worlded speaker had answered. The interface supplies turn-taking, first-person form, apology, clarification, confidence, and apparent patience. The user supplies concern.

The result is not speech in the strong sense. It is generated language vivified by reception.

The concept *sprachlich aber weltlos* makes this visible. The system's linguistic form is the condition of vivification. Its worldlessness is what the vivification conceals. The user experiences answer. The system has generated language without answerability. The institution may circulate the output as if answerability were present.

The ethical burden therefore lies with the worlded user. The user must return the output to sources, tasks, persons, practices, and consequences. They must not let the worldless form of language stand in for worlded judgment.

17. Why the term must remain strange

The phrase *sprachlich aber weltlos* should remain strange. It should not be domesticated too quickly as "linguistic but worldless," then forgotten. The German keeps the Heideggerian pressure visible.

Sprachlich indicates that the phenomenon cannot be dismissed as mute machinery. The system operates in language. It produces the most world-like surface available to human beings: articulated speech.

Aber marks the interruption. The linguistic operation does not lead to world. The “but” must remain.

Weltlos recalls the stone, but it is not stone-worldlessness. It marks lack of world in a system that nevertheless operates on language. This is the anomaly.

The phrase prevents two opposite errors.

The first is enchantment. Because the system speaks fluently, one treats it as a subject, interlocutor, mind, agent, or judge. *Weltlos* blocks that.

The second is dismissal. Because the system lacks world, one treats it as mere machinery with no philosophical significance. *Sprachlich* blocks that.

The system is not a person. It is not nothing. It is a technical entity that returns language without world.

18. Toward the veil

This chapter has given the ontology of the generator. The next step describes the experience of receiving its output.

The model is *sprachlich aber weltlos*. It produces language without world. Yet its language carries traces of world: tense, deixis, person, address, evidence, modality, narrative, tone, style, and apparent judgment. The output therefore appears as if it came from a locus. It appears as if someone, or something like someone, had spoken from a present, remembered a past, projected a future, weighed a matter, and addressed the user.

That appearance follows from the structure. A worldless linguistic system has been trained on the traces of worlded language. It returns those traces in the grammar of answer.

Section 13 of Chapter 3 showed why this appearance does not ordinarily arrive as ontological breach. The model is encountered first as equipment inside a task. The user opens it to draft, summarize, translate, prepare, or answer. The strangeness of receiving essay-like language from a being without world is absorbed by the equipmentality framework that organizes the encounter. The framework holds. The breach does not reach the threshold of recognition. The veil that the next chapter will describe rests on this prior phenomenological condition.

The output arrives as if world had returned with language.

World has not returned. The user must restore it.

The next chapter will name this appearance-form as the structural veil of AI output. The veil is not merely error, falsehood, or deception. It is the form in which language without world appears to a worlded reader.

The model speaks as if from nowhere. The model is not nowhere. At the level of language, it lacks existential locus. At the level of infrastructure, it has material locus. Between these facts lies the work of de-veiling.

The first de-veiling is ontological: AI is *sprachlich aber weltlos*.

The second will be hermeneutic: the user must read between the times.

NOTES

1. Martin Heidegger, *Die Grundbegriffe der Metaphysik: Welt, Endlichkeit, Einsamkeit*, Gesamtausgabe 29/30 (Frankfurt am Main: Vittorio Klostermann, 1983), 262 to 263; English in *The Fundamental Concepts of Metaphysics: World, Finitude, Solitude*, trans. William McNeill and Nicholas Walker (Bloomington: Indiana University Press, 1995), 177. Heidegger formulates the three theses as: “der Stein ist weltlos; das Tier ist weltarm; der Mensch ist weltbildend.” I translate *weltlos* as “worldless,” *weltarm* as “poor in world,” and *weltbildend* as “world-forming.” The English GA 29/30 text gives the three theses at §42. ↩
2. Heidegger, *Die Grundbegriffe der Metaphysik*, GA 29/30, 263; *The Fundamental Concepts of Metaphysics*, 177. The term *Weltarmut* should not be read as a simple value judgment. Heidegger later distinguishes the animal’s poverty in world from the stone’s worldlessness

by treating animal poverty as deprivation, while the stone cannot even be deprived of world. See GA 29/30, 289 to 292; English, 196. ↵

3. Heidegger, *Die Grundbegriffe der Metaphysik*, GA 29/30, 347 to 349, 376 to 377; *The Fundamental Concepts of Metaphysics*, 239, 259. McNeill and Walker render the central contrast as the animal's capacity to behave, *sich benehmen*, without apprehending, *vernehmen*, something as something. This translation preserves the German contrast between behavior and apprehension. ↵
4. Heidegger, *Die Grundbegriffe der Metaphysik*, GA 29/30, 262 to 263; *The Fundamental Concepts of Metaphysics*, 177. The translation "world-forming" is retained because "world-building" can suggest voluntary construction, while Heidegger's *weltbildend* names the human capacity to stand in an openness in which beings can appear as beings. ↵
5. Martin Heidegger, "Die Frage nach der Technik," in *Vorträge und Aufsätze*, Gesamtausgabe 7 (Frankfurt am Main: Vittorio Klostermann, 2000), 5 to 36; English in "The Question Concerning Technology," in *The Question Concerning Technology and Other Essays*, trans. William Lovitt (New York: Harper & Row, 1977), 3 to 35. *Gestell* and *Bestand* are used here to describe the revealing of language as orderable material. The chapter does not claim that every technical mediation of language is identical in kind. ↵
6. Hubert L. Dreyfus, *What Computers Can't Do: The Limits of Artificial Intelligence* (New York: Harper & Row, 1972); Hubert L. Dreyfus, *What Computers Still Can't Do: A Critique of Artificial Reason* (Cambridge, MA: MIT Press, 1992); Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I* (Cambridge, MA: MIT Press, 1991). Dreyfus is used here as a clarifying bridge from Heideggerian worldhood to AI critique. He is not the chapter's foundation. ↵
7. Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, MA: MIT Press, 2019). Smith's distinction between reckoning and judgment supports the chapter's claim that AI can perform powerful linguistic computation without entering the worlded answerability required for judgment. ↵
8. Francisco J. Varela, Evan Thompson, and Eleanor Rosch, *The Embodied Mind: Cognitive Science and Human Experience* (Cambridge, MA: MIT Press, 1991); Evan Thompson, *Mind in Life: Biology, Phenomenology, and the Sciences of Mind* (Cambridge, MA: Harvard University Press, 2007). Enactivism is used as a contemporary support for the claim that meaning is bound to embodied and living engagement. It does not replace the Heideggerian center of the chapter. ↵
9. Ezequiel A. Di Paolo, *Autopoiesis, Adaptivity, Teleology, Agency, Phenomenology and the Cognitive Sciences* 4 (2005): 429 to 452; Ezequiel A. Di Paolo, Thomas Buhrmann, and Xabier E. Barandiaran, *Sensorimotor Life: An Enactive Proposal* (Oxford: Oxford University Press, 2017). The concepts of autopoiesis, adaptivity, and precariousness clarify why a nonliving vectorized system lacks the living normativity of organisms. ↵
10. Hanne De Jaegher and Ezequiel A. Di Paolo, "Participatory Sense-Making: An Enactive Approach to Social Cognition," *Phenomenology and the Cognitive Sciences* 6 (2007): 485 to 507; Elena Cuffari, Ezequiel A. Di Paolo, and Hanne De Jaegher, *Linguistic Bodies: The*

Continuity Between Life and Language (Cambridge, MA: MIT Press, 2018). These sources support the claim that language grows from embodied social participation. Their role is supportive, not architectonic. ←

11. Jacques Derrida, *L'animal que donc je suis* (Paris: Galilée, 2006); English in *The Animal That Therefore I Am*, ed. Marie-Louise Mallet, trans. David Wills (New York: Fordham University Press, 2008); Jacques Derrida, *La bête et le souverain, volume II* (Paris: Galilée, 2010); English in *The Beast and the Sovereign, Volume II*, trans. Geoffrey Bennington (Chicago: University of Chicago Press, 2011). Derrida's critique complicates Heidegger's human-animal boundary, but it does not make AI a living or worlded being. ←
12. Giorgio Agamben, *L'aperto: L'uomo e l'animale* (Turin: Bollati Boringhieri, 2002); English in *The Open: Man and Animal*, trans. Kevin Attell (Stanford: Stanford University Press, 2004). Agamben's anthropological machine helps show why the human-animal boundary is not innocent. The present chapter uses that critique to discipline the Heideggerian triad, not to dissolve the distinction between living beings and worldless technical systems. ←

Chapter 13. The Structural Veil of AI Output

1. The appearance of an answer

The preceding chapter named the entity: *sprachlich aber weltlos*, linguistic but worldless. The system operates in language without having world. It can generate the form of speech, explanation, address, and interpretation. It cannot encounter beings as beings. It cannot disclose a matter. It cannot be claimed by what it says.

This chapter turns from the ontology of the system to the appearance of its output.

The user does not meet the model as weights, vectors, tokens, layers, and probability distributions. The user meets a text. Often the text arrives as an answer. It has structure, tone, relevance, sequence, and apparent care for the question. It may cite, compare, correct, narrate, advise, use first person, address the user directly, or appear to remember.

The appearance is not incidental. It belongs to the structure of generated language. The model has been trained on traces of worlded speech and writing. It returns those traces in a form human beings ordinarily associate with answerability. The user sees grammar, not weights. The user hears a voice, not a vector.

What is veiled is not only the absence of a speaker behind the output. What is veiled is also the source-locus omission that made the output possible: the disappearance of the historical and contemporary worlds from which its language descends.

The structural veil names this appearance-form.

The veil is not a simple illusion. It is not only careless anthropomorphism. It is not only falsehood. It is the form in which detemporalized and worldless language appears when it returns as fluent output. The output carries the marks of temporal speech while lacking the temporal locus of speech. It carries the form of interpretation while lacking the historical horizon of an interpreter. It carries the form of answer while lacking answerability.

The structural character of the veil has a phenomenological mechanism. Section 13 of Chapter 3 named it. The model is encountered first as equipment inside the user's work-world. The equipmentality framework organizes the encounter. The output arrives as the tool's contribution to a task already underway. The categorial strangeness of receiving language without a speaker is absorbed by the usability of the response.

The veil does not need the user to believe that a speaker is present. It needs only that the framework continue to hold while the output is incorporated into the work. The veil is therefore not simply an illusion in the user's mind. It is a maintained framework of encounter. Its structural character begins in the way the situation is organized before explicit belief becomes decisive.

The veil is structural because it arises from the success of the system, not only from its failures. When output is false, the veil may tear. A fabricated citation or impossible date can expose the machinery. When output is fluent, balanced, accurate, and useful, the veil becomes harder to see. The better the answer seems, the more easily its missing locus disappears.

What the veil collapses can be named more precisely. Generated language has come through at least four distinct levels of locus, each transformed into the next. The first is the source locus: the original text, author, institution, world, event, or practice from which a linguistic trace first arose. The second is the corpus locus: the cleaned, filtered, aggregated dataset relation in which that trace has been gathered with many others and prepared for technical operation. The third is the model locus: the learned internal relation, after training, in which traces from many sources have become operable mathematical patterns. The fourth is the output locus-effect: the generated answer as it appears to the reader, with its grammar of voice, time, and source.

Each transformation is real. Each level differs from the next in kind. Source locus is historical and answerable. Corpus locus is technical and operational. Model locus is statistical and relational. Output locus-effect is phenomenological: it is what the reader receives. The structural veil is the merging of these four levels into the appearance of one: the appearance of a speaker speaking from a world. Locus-reinjection, the practice the next chapters develop, is the discipline of holding the four levels apart so that the output can be used without being mistaken for what it is not.

Two reminders apply. AI handles time technically; it does not inhabit time existentially. And: the model speaks as if from nowhere. The model is not nowhere.

This chapter holds the phenomenon in view before the next chapters name the response: *intus legere* and locus-reinjection.

2. The recurring veil

The structural veil does not appear for the first time with AI output. It is the culmination of a recurring pattern that has governed the whole book. At each stage, a derivative form becomes so useful, available, or fluent that it can conceal the more originary structure from which it came.

The grid veils place. It represents spatial relations with precision, but its precision can hide the fact that human beings first inhabit places, regions, paths, thresholds, tools, and sites of concern. Clock-time veils lived temporality. It coordinates public life, but its sequence of now-points can hide waiting, inheritance, anticipation, regret, promise, and finitude. Sign-system theory can veil discourse. It studies words, rules, and meanings, but it can hide the temporal event in which a world becomes articulated by a speaker who is affected, projected, and answerable. Writing can veil voice. It preserves marks after the speaker has departed, and those marks can circulate without their first situation; yet writing still allows reading to return toward source, date, genre, text, world, and question. Corpus formation veils historical distance under availability. It gathers traces from many times and worlds into one operable field, making the past searchable, comparable, and processable before it has necessarily been read. Vectorization veils temporal locus under spatial relation. It converts the written and corpus-gathered trace into position, proximity, direction, weight, activation, and probability.

A further veil belongs specifically to the AI encounter. The equipmentality framework veils the categorial strangeness of the output. Section 13 of Chapter 3 articulated this at the level of use: the model is met as a tool within a task, and its strangeness as a language-producing system without world is absorbed by its usefulness. This framework-veil is not another ancient veil in the history of representation. It is specific to the moment in which generated language enters human practices as equipment. It joins the prior veils as the contemporary condition under which generated language is received without sufficient disturbance.

AI output gathers all of these prior veils into the fluent surface of an answer.

This is why AI output is not merely one more example. It is the culminating appearance. The prior veils return together in one generated surface. The answer appears as if source, time, question, judgment, and address had already been gathered.

The earlier veils were partial. The AI veil is synthetic.

It gathers them and returns them as fluency.

3. The word “veil”

A veil conceals by covering. It also shapes what can be seen. It does not make the object vanish. It gives the object a form of appearing. The structural veil of AI output works in this way.

The output does not hide everything. It shows language, relation, probable structure, and a possible answer. It may show sources if retrieval or citation has been added. It may show dates, names, distinctions, and qualifications. It may show enough to be useful.

What it hides is the relation between the language and its loci.

It hides the source loci of the traces from which the model learned: authors, translators, editors, annotators, institutions, archives, publics, and genres. It hides the temporal loci of those traces: dates, historical situations, publication contexts, revisions, controversies, and later receptions. It hides the reading time required to return to those sources and the formation time through which a human being would become capable of saying the matter as their own.

It also hides the operational locus: corpus construction, filtering, tokenization, vectorization, training, alignment, retrieval, system instruction, interface design, and deployment. It hides labor, infrastructure, ownership, and the data center until the later political-material chapter returns to them explicitly.

Most of all, it hides the user’s own contribution. The model generates. The user supplies world. The user brings question, anxiety, deadline, hope, task, institution, and consequence. The output can matter because the user’s world gives it mattering. The veil makes this contribution appear as if it were the model’s understanding.

The veil is therefore not one concealment. It is layered concealment. It conceals source, time, operation, material support, ownership, and user projection under the surface of a smooth answer.

This is why “veil” is more exact than “error.” The output may be correct and still veiled.

4. The grammar of locus without locus

The first layer of the veil is grammatical. The output can use the forms by which human language normally marks locus: person, tense, deixis, modality, evidentiality, and address.

A model can say “I,” “you,” “here,” “now,” “earlier in our conversation,” “the source indicates,” “I infer,” “it is likely,” “we should distinguish,” or “the next step is.” Each phrase belongs to a grammar of situated speech.

Chapter 4 named the core problem as referential hollowing. Deictic forms remain, but their existential anchoring is missing. When a model generates “I am here now,” the “I” does not posit a subject, the “here” does not arise from a lived place, and the “now” does not emerge from a phenomenological present. The linguistic marker is present. The origo is absent.

This hollowing is not confined to explicit first-person phrases. A neutral paragraph can also imply a locus. It can order a field, select what matters, state what follows, identify the central issue, and present a conclusion. The prose may avoid “I,” yet still speak from an apparent standpoint. Scholarly impersonality is still a standpoint when written by a scholar. Judicial impersonality is still a standpoint when written from an office. Generated impersonality is different. It produces the form of standpoint without one who stands.

The veil begins here. The output has the grammar of a saying. It does not have the being of one who says.

This does not make every first-person convention illegitimate. A tool may say “I am processing your request.” A website may say “we use cookies.” A form may say “enter your name here.” Such phrases can be harmless shorthand. The danger appears when the grammar of locus is received as evidence of subject, world, and answerability.

The generated “I” may be an interface convenience. The user’s ear still hears the grammar of a speaker.

5. The apparent present

The structural veil gives AI output an apparent present. The answer arrives now. It appears tailored to the user's question now. It may include recent terms, current-sounding language, and immediate synthesis. The user experiences the output as present response.

This apparent present must be divided.

There is first the operational present: the system is running, receiving input, processing context, calling tools if enabled, and generating output. This present is real. It belongs to technical operation.

There is next the user's present: the human being asks from a situation. They are here, now, with a need, doubt, deadline, curiosity, fear, or practical decision. This present is also real. It belongs to *Dasein's* temporal locus.

There is then the source plurality: the language that shaped the model comes from many pasts. Books, posts, articles, translations, code, records, comments, and institutional texts entered the corpus from many times. Retrieval, if present, may add a more specific source time. This plurality is usually hidden behind the answer.

Finally, there is the generated present: the prose speaks as if the relevant pasts had already been gathered into one now.

The veil is strongest at this fourth level. The generated answer makes many times appear as one time. Source time, corpus time, training time, retrieval time, generation time, and user time are compressed into the present of output. The past is not abolished. It is made available under the sign of immediate answer.

This is one reason the output feels efficient. It spares the user the delay of source return. It gives the form of arrival before the path has been walked.

Some compression is useful. A dictionary compresses usage. An index compresses a book. A commentary compresses a tradition. A teacher compresses years of study into a lecture. Compression is not the problem. The problem is compression without visible locus, method, or answerability.

A teacher can be asked where the claim comes from. A commentary can cite its sources. A scholar can be challenged. A model can generate a source-shaped answer without being a source.

6. The voice from nowhere

Book one described AI output as a voice from nowhere. The phrase remains useful, but it now needs correction.

The output appears placeless because the interface removes visible place. The user sees a response in a box, not the data center, corpus, labor, energy use, company, interface design, or system instruction. The output also appears authorless. It is not signed in the way a book, article, letter, or judgment is signed. Even when the platform name is visible, the sentence itself appears as neutral synthesis.

This produces the voice from nowhere.

Yet the phrase must not be misunderstood. The model is not nowhere. It is not outside the world. It is not immaterial intelligence. It is an owned and operated technical system, running through material infrastructure, trained and deployed by institutions. The material locus will be treated later. This chapter concerns the interface effect: the system's actual locatedness disappears behind generated voice.

The disappearance matters because placelessness gives authority. A person speaks from a position. Their position can be challenged. A scholar speaks from a field, tradition, training, and method. A judge speaks from a court. A doctor speaks from a practice and license. A witness speaks from having been there. A model's output often appears free of these limits, as if it were synthesis without standpoint.

The result is not genuine neutrality. It is concealed situatedness.

The model speaks as if from nowhere. The model is not nowhere.

Both clauses are needed. The first names the veil. The second begins the de-veiling. The voice has no existential locus. The system has material locus. The output has apparent universality. The operation has specific conditions.

A human reader must hold both facts together.

7. *Gerede* returned as answer

Heidegger's concept of *Gerede*, idle talk, gives the veil its first discursive form. The term does not mean gossip. Heidegger warns that *Gerede* is not used in a disparaging sense. It is a positive phenomenon of everyday disclosedness.¹ It names the mode in which what has been said circulates without original appropriation.

His central formulation remains decisive:

“Das *Gerede* ist die Möglichkeit, alles zu verstehen ohne vorgängige Zueignung der Sache.”

“Idle talk is the possibility of understanding everything without a prior appropriation of the matter.”²

The danger is not only falsehood. *Gerede* may contain correct statements. It may be accurate, useful, and socially necessary. Its danger lies in groundlessness. One can speak about the matter without having returned to the matter. One can understand in the public mode without having made the question one's own.

AI does not fall into *Gerede*. It is not *Dasein*. It has no self to lose in publicness. It generates language that can function as *Gerede* in the human user-machine relation. Book one named this as the vectorized afterlife of the language of the They: AI is not *das Man*, but operates on the written sediment of public intelligibility.

The structural veil changes the form of *Gerede*. The already-said returns not only as public saying but as personalized answer. The user asks a private question, and the already-said answers in a form adapted to them. Public language is returned as if it were addressed speech. This is one reason generated output feels stronger than ordinary public formula. It does not only say what one says. It says it to me.

The veil hides the public sediment behind the private response.

This is why the old danger intensifies. The user may receive the already-said not as public formula, but as a fitted answer. The output may be correct. It may even be helpful. Yet it can remain unappropriated if the user does not return through it to source, matter, and responsibility.

The answer-form is not yet ownership.

8. The fetish of language

Book one described AI as an intraworldly fetish of language. That concept belongs inside the structural veil.

A fetish is not a mere illusion. It is a functional misrecognition. Something real happens, but the source of its power is concealed. Marx's analysis of commodity fetishism shows how social relations among persons can appear as relations among things. Heidegger's *Gestell* shows how beings can be revealed as standing-reserve. In AI, linguistic labor, public language, infrastructure, capital, data, and user projection can appear as the model's own knowledge.³

The output appears to know because the genealogy of language is hidden. The corpus disappears. The sources disappear. The labor of authors, annotators, translators, workers, teachers, scholars, and users disappears. The data center disappears. Ownership disappears. The user's role in making the output matter also disappears.

What remains is the answer.

The fetish is powerful because the answer really can help. The output can be clearer than the user's first thought. It can reveal a distinction, organize a field, correct a spelling, draft a paragraph, or produce a useful summary. The misrecognition is therefore not total fantasy. It is a misrecognition of source and mode.

The model's apparent knowledge is the surface of many hidden relations. This is why the structural veil cannot be removed only by telling users that models predict tokens. That statement is true but insufficient. The user does not interact with the training objective. The user interacts with the answer as a social and hermeneutic object.

The fetish of language must be de-fetishized by returning the output to the relations it conceals. This return is not yet full locus-reinjection, but it prepares it.

One must ask: whose language has become this answer, through which system, under whose ownership, for which user, and at what cost of time and source.

9. The *Erlebnis* of veiled output

The structural veil gives the user an *Erlebnis* of understanding without the formative passage of *Erfahrung*.

Erlebnis names the lived event: the occurrence, immediacy, and experience that happens to the subject. *Erfahrung* names formative undergoing: experience tested through time, negation, correction, memory, and revised understanding. The generated answer can become an *Erlebnis* for the user. Language arrives. Difficulty takes form. The question becomes answer-shaped. The user feels oriented.

This event is real. It may be helpful. It can begin reading, support drafting, clarify a distinction, or expose a possible structure. The problem begins when the event is received as if *Erfahrung* had already occurred.

The model has no *Erfahrung*. It undergoes no formative correction. It does not suffer the collapse of a first interpretation. It does not acquire tact through cases. It does not become answerable for its claims. The user may have an *Erlebnis* through the output, but the model has not undergone anything.

The user may also fail to undergo anything. The output can relieve the user from the passage through which understanding would have formed. The difficult book becomes summary. The unresolved question becomes structure. The ambiguous case becomes recommendation. The blank page becomes draft. The answer arrives before the user has stayed long enough with not-knowing.

This is why the veil is not only epistemic. It is temporal. It covers the missing time of formation.

The output gives the user the feeling of being after the work: after difficulty, reading, deliberation, or encounter. In fact, they may only be after the prompt.

Some friction is waste. Some friction is formation. The veil makes the difference hard to see.

At scale, the individual reception described here returns as the thinning that Section 13 of Chapter 3 named. The missing time of formation is not only the loss of one user's path into one matter. Where the missing time is repeated across practices, formation-passages contract across communiones.

The veil is therefore not only epistemic and temporal. It also has a communio-level affective register. It supports a transformation that may be felt as *Verzweiflung* by those whose practices and formations are being absorbed: not the shock of a broken category, but the despair of a world continuing while its paths contract.

10. User vivification

The model generates. The user supplies world.

This formulation should remain central. It prevents the analysis from placing too much agency inside the system. The output becomes meaningful because the user receives it inside a situation that matters. The user brings the project, question, fear, hope, institution, deadline, desire to know, and need to act.

This is user vivification. The model produces linguistic continuation. The user experiences answer. The model has no concern. The user is concerned. The model has no world. The user inserts the output into one. The model has no responsibility. The user's situation has consequences.

Vivification is more precise than anthropomorphism. Anthropomorphism describes attributing human traits to something non-human. Vivification describes the process by which generated language becomes alive for the user because it enters a human world. A user may know that the model is not a person and still receive the output as if understanding had occurred.

The interface does not need to deceive the user completely. It needs only to provide fluent response in the grammar of social exchange. The user's linguistic habits do the rest. In ordinary life, coherent answer usually comes from someone who can be asked for reasons. AI output interrupts that expectation. It gives the form of answerability without the speakerly condition that would make answerability possible.

This is why the veil belongs to the relation, not only to the model. The system produces language without world. The user gives it a place in a world. The structural veil appears at the crossing.

The danger is not that the user is foolish. The danger is that human linguistic life makes the output compelling. We are trained by speech itself to hear answer as the trace of someone who can answer.

11. Synthetic dialogue

The chatbot interface gives the structural veil a dialogic form.

The user asks and the model answers. The user corrects and the model revises. The user asks for nuance, the model supplies nuance. The user objects, and the model apologizes or reframes. The exchange has turn-taking, continuity, repair, tone adjustment, and apparent patience. It looks like dialogue.

Dialogue in the hermeneutic sense is not turn-taking alone. Gadamer's *Gespräch* is an event in which participants risk their prejudgments before a matter. The subject matter, *die Sache*, leads. The participants are changed by what is said. Genuine dialogue requires openness to being addressed by the matter and by the other.⁴

Synthetic dialogue lacks existential reciprocity. The model can revise its output without risking its self-understanding. It can apologize without responsibility. It can ask a question without needing to know. It can simulate humility without becoming humble. It can remember a prior turn in the context without having undergone a past.

This does not make the interaction useless. A user can rehearse an argument, test a formulation, identify an ambiguity, or discover that their question is poorly formed. Such use can support thought. The issue arises when the outer form of dialogue is received as the presence of an interlocutor who has understood.

The interface gives response without responsibility.

This form of the veil is subtle because it works through ordinary politeness. The model's patience may calm the user. Its readiness may encourage inquiry. Its lack of irritation may make it feel safer than a human exchange. These effects can be beneficial. They can also intensify projection.

The user is addressed. No one is there in the sense required for dialogue.

12. Unowned truth

The deepest risk of the structural veil is not hallucinated falsehood. It is unowned truth.

A false output can sometimes be checked. It may expose itself by contradiction, bad citation, impossible chronology, or factual error. A correct output is harder. It may be accurate, well-phrased, balanced, and useful. It may enter a report, legal memo, clinical note, hiring assessment, board deck, or classroom. It may circulate because it is correct enough to pass.

Yet correctness is not ownership.

A claim becomes owned when someone can answer for it: give reasons, name sources, state limits, revise under pressure, and bear consequences. AI output can give the user a correct sentence before anyone has taken the sentence through the work required to own it. The truth then circulates without *Erfahrung*.

This risk was central in Book one and Book two. Book one named the danger as correct language without appropriation. Book two showed how generated output can look like professional judgment while lacking encounter, formation, and answerability. The companion paper gives the public-facing version: generated language may be fluent, useful, and accurate while remaining artificial fluency rather than judgment.

Unowned truth is not a minor academic problem. It changes institutions. A generated sentence can become a file. A file can become a record. A record can become evidence that a process occurred. A recommendation can enter a decision. A summary can replace the source. A phrase can become the organization's position.

At each step, the output gains authority and loses visible genealogy.

The structural veil therefore does not only affect private understanding. It affects institutional memory. It allows language to move as if judgment had occurred, when only generation and circulation may have occurred.

13. Hallucination beyond falsehood

The word hallucination is commonly used for fabricated or ungrounded model output. The term is imperfect because it imports a psychological model. A person hallucinates when perception loses its ordinary relation to the world. A model has no perception in the first place. It generates fluent continuation under conditions of uncertainty, insufficiency, or misdirection.

Even so, the term can be retained in a limited way if it is divided.

Technical hallucination concerns false or unsupported content: fabricated citations, invented facts, nonexistent cases, impossible events, or confident errors. These matter. They can cause harm. They must be addressed by verification, retrieval, design, institutional discipline, and user training.

Phenomenological hallucination is different. It occurs when generated language is received as living knowledge. The user experiences answer, intention, understanding, care, or judgment where the system has produced linguistic form. This can happen even when the output is factually correct.

The structural veil belongs to this second level.

A hallucinated citation may reveal the absence of source. A correct uncited claim may hide it better. A false answer may be rejected. A correct answer may be absorbed without return. The phenomenological hallucination lies in the experience that a subject has answered.

This is why the book does not build its critique on errors alone. Error is visible failure. The veil is successful appearance.

The output's correctness may deepen the concealment.

14. Evidential phrases and borrowed caution

The veil often appears through caution rather than confidence. The model says "according to," "it seems," "the evidence suggests," "one way to read this is," "a careful distinction is needed," or "this should be verified." These phrases are valuable in human writing. They mark relation to evidence, uncertainty, and method. In AI output they may function differently.

The phrase “according to” normally points to a source. If the source is named and available, the phrase can support return. If the source is vague, missing, or generated, the phrase may imitate evidence without showing it.

The phrase “it seems” normally marks a speaker’s judgment under uncertainty. It implies that someone is holding back from overclaiming. In generated output, it may be a learned style of caution.

The phrase “the evidence suggests” normally refers to evidence gathered and weighed. In generated output, it may refer to patterns in the prompt, retrieval, training, or no determinate source at all.

Such phrases do not need to be banned. They often improve output by reducing false certainty. The point is that cautious language can itself be veiled. It can produce the appearance of method without method. It can make output feel more scholarly, responsible, and situated than it is.

This is borrowed caution.

The user must therefore ask whether the caution is grounded. A qualified claim may still be unowned. A citation-shaped sentence may still lack source return. A measured tone may still conceal absence of judgment.

The veil does not always speak loudly. Sometimes it whispers like prudence.

15. Citation and the false return

Citations can pierce the veil. They can also decorate it.

A citation is a visible path back to a source. In scholarly practice, it lets the reader test the claim, follow the argument, recover context, and distinguish what is being said from where it came. A citation is not proof by itself. It is a relation of answerability.

AI output with citations may support locus-reinjection. It gives the user a possible route. It may help return a generated claim to a source locus. This is an improvement over untraceable synthesis.

But a citation can also become a false return. The model may cite a source that does not support the claim. It may cite a real source for a general proposition while the actual formulation comes from elsewhere. It may retrieve a passage and then overstate it. It may give the appearance of source discipline while the user remains at the level of generated summary.

The problem is not only fabricated citations. It is citation without reading.

A human writer can do the same. A citation can be ornamental, defensive, lazy, or false. The difference is scale and interface. AI can produce citation-shaped prose faster than the user can verify it. The burden shifts to the reader.

Eco's distinction between interpretation and overinterpretation is useful here. The *intentio operis*, the intention of the work, constrains interpretation. A source must be allowed to resist the claim made in its name.⁵ Generated interpretation often lacks this resistance unless the user returns to the text and lets it correct the output.

Citation is a path. It is not the walk.

16. The source made too present

Writing preserves distance. Historical reading lets that distance work. The corpus gathers distance into availability. Vectorization turns availability into calculable relation. Generated output returns the result as present synthesis.

The source is then made too present.

The source made too present may be past or contemporary. A living historical world can be rendered too present when the model turns it into explanation before the reader has undergone the difficulty of approach.

This phrase does not mean the source becomes clearer. It means the source's distance is suppressed. A medieval text appears in contemporary prose. A philosophical dispute appears as a balanced answer. A political event appears as a bullet list. A person's testimony appears as a theme. A legal case appears as a rule. A book appears as a summary.

Such transformations can help. They can introduce, orient, and reduce unnecessary difficulty. But they can also deprive the source of its resistance. The text no longer arrives as something from elsewhere. It arrives as content already adapted to the user's present demand.

Gadamer's account of temporal distance resists this. Temporal distance is not only an obstacle. It is the medium of understanding.⁶ A historical text does not become meaningful by being stripped of its distance. It becomes meaningful when the reader's present horizon meets the text's historical horizon and is tested by it.

The structural veil makes the fusion of horizons look already accomplished. The source has not been encountered. It has been rendered.

This is why generated summaries can be useful and dangerous in the same act. They give access. They also risk replacing the temporal difficulty that would have made access formative. The source appears close. It may not have been read.

The past becomes promptable. It may stop being historical.

17. The ghostly surface

The chapter on writing described reading as the possibility of "summoning the ghosts of true experiences." The phrase belongs here again, with a difference.

A written trace can summon a ghost because it once belonged to a temporal locus. A letter, diary, poem, testimony, theorem, legal opinion, or philosophical passage carries marks of a life, institution, question, or historical world. The reader knows that the speaker is absent. The absence is part of the reading. The ghost has a source.

Generated output also summons ghosts. It draws on traces of prior human language. It can echo voices, styles, genres, arguments, and emotional forms. It can produce sentences that sound as if they came from experience: grief, hope, judgment, caution, intimacy, authority. It can make the dead, the distant, and the public speak in one smooth surface.

But the generated ghost is different. It is not the return of one absent speaker's trace. It is a recombined appearance of speakerhood drawn from many traces. There may be no single source to which the voice can be returned. There may be only a topology of linguistic inheritance hidden inside training and generation.

This is why the output can feel uncanny. It is neither living speech nor ordinary writing. It is language from true experiences recombined without the experiences. It can move the user because human language still carries the force of lives. It cannot own that force.

The ethical task is not to despise the ghost. It is to send it back to its loci.

Some of those loci are textual. Some are historical. Some are institutional. Some are material. The final one will be the data center.

18. The anti-boredom surface

The previous chapter touched on *Langeweile*, boredom. The structural veil also has a temporal effect at the level of attention.

A user faced with not-knowing may have to wait. They may have to read, search, fail, begin again, sit with silence, or let the question become sharper. That waiting can be sterile. It can also be formative. Heidegger's analysis of boredom shows that not every empty time is worthless. In deeper boredom, ordinary significance weakens, and the human being may be returned to the question of being.⁷

AI output often fills the gap before the gap can work. It turns uncertainty into prompt, prompt into answer, answer into next prompt. It does not make the user bored. It prevents the emptiness in which thought might have had to appear.

This is not always bad. Some waiting is waste. Some friction is unnecessary. Some silence is only exclusion, fatigue, or lack of access. Tools that reduce such burdens can be humane. The point is not to protect difficulty for its own sake.

The point is that the structural veil hides which difficulty has been removed. Has the system removed a wasteful obstacle, or has it removed the formative passage by which the user would have learned to see? The surface of output does not decide. It presents both as relief.

This is one way the veil reaches attention. It does not only conceal sources. It conceals the temporal cost of bypassing them.

19. Correctness and the veil

The structural veil must be described independently of accuracy.

Accuracy concerns whether the output's claim corresponds to the matter. The veil concerns the form in which the output appears and the relations it hides. A correct answer can be veiled. A false answer can be veiled. A useful answer can be veiled. A harmful answer can be veiled. The question of truth remains necessary, but it is not sufficient.

A generated statement may be factually correct while hiding its source. It may be well argued while hiding the tradition it compresses. It may be emotionally apt while hiding its lack of experience. It may be ethically balanced while hiding the absence of responsibility. It may be current because retrieval was used, while hiding the user's need to judge the retrieved sources. It may cite a real passage while hiding the fact that the passage has not yet been read.

The veil is therefore not removed by improving factuality alone. Better models may reduce errors. They may cite more sources. They may track context more effectively. They may use tools and retrieval more reliably. These improvements matter. They do not abolish the distinction between generated form and temporal locus.

A larger context window is not a life. A memory store is not a history undergone. Retrieval can return sources to output, not make the model read. Tool use can extend generation into action, not make the system responsible. This boundary was stated in Book two and remains decisive here.

The structural veil may become more refined as systems improve. The better the output resembles responsible speech, the more disciplined the user must be in refusing to mistake resemblance for responsibility.

The veil is not only a problem of weak models.

20. The artifactual and cultural registers of the veil

The structural veil has two registers.

The first is artifactual. The second is cultural.

The artifactual veil appears at the level of the individual generated output. A sentence, paragraph, image, voice, video, summary, citation, or recommendation appears to carry a relation to source, witness, authorship, capture, judgment, or answerability that it does not carry in the way its form suggests.

This is what the preceding sections have described in the case of language. The generated answer appears as if a speaker had answered, a reader had read, a source had been returned to, a judgment had been formed, or a horizon had been risked before the matter. The output carries the grammar of such acts. It can say “I infer,” “the source indicates,” “this means,” “the evidence suggests,” “we should distinguish,” or “the next step is.” It can produce caution, balance, qualification, and even apparent humility. Yet the form of these acts is not the act itself.

A generated answer is not a source.

A generated interpretation is not a reading.

A generated recommendation is not a judgment.

A generated synthesis is not *Erfahrung*.

This is the artifactual veil: the single artifact appears to bear the marks of a temporal and responsible act whose locus has been concealed, diluted, or never existed in the required sense.

The visual case makes the structure easier to see. A generated image may appear as a photograph. It may carry depth of field, lens distortion, documentary framing, the accidental density of a street, the face of a witness, the posture of a body, the light of a place. It can look like an image that came from an event of capture. Yet the event shown did not occur. No one stood there. No camera witnessed it. No body was present before the lens. The image retains the visual grammar of capture without the captured event.

The linguistic case is less immediately visible because writing has fewer perceptual marks of production. A generated sentence and an authored sentence can share the same surface. There is no equivalent of a lens flare by which the reader can know that the sentence came through hesitation, reading, revision, risk, and responsibility.

The artifactual veil therefore concerns the individual output. It asks what the artifact appears to carry, and what relation has actually produced it.

The cultural veil begins when this question no longer remains inside the individual artifact.

At scale, generation changes the reception-condition of the medium itself. The problem is no longer only whether this photograph is generated, this video synthetic, this paragraph written by a person, or this testimony grounded in witness. The deeper problem is that every artifact now appears in a field where generation has become an available suspicion.

This is the cultural veil.

It is not the sum of many artifactual veils. It is the transformed condition under which artifacts are now received. Once synthetic production becomes ordinary, authenticity itself acquires a new burden.

The real photograph must prove it was captured.
The authored text must prove it was written.
The witness must prove there was witnessing.

The artifactual veil gives us the single output that appears to carry locus.

The cultural veil gives us the medium in which locus itself must now be proven.

That is the hinge. The consequences that follow depend on why this burden is historically new.

This burden is new in form, but not because photographs were ever pure truth or writing was ever innocent. Photographs have always been framed, staged, cropped, edited, captioned, circulated, and interpreted. Writing has always included anonymity, pseudonymity, forgery, ghostwriting, bureaucracy, propaganda, formula, copying, citation, plagiarism, and public language no one has fully owned. The history of media is not a history of pure presence later corrupted by machines.

The novelty is more precise.

Before generation at scale, falsification almost always still implied a falsifier. The prior tradition handled anonymous texts, pseudepigrapha, forgeries, ghostwriting, bureaucratic prose, formulaic writing, propaganda, and obscure source loci as deviations from a structure that remained in place. The deviations were thinkable because that structure was assumed. The forged document had a forger; the pseudonymous text, a writer; the propaganda tract, an institution, interest, or strategy; the bureaucratic phrase, an office. Even where the locus was hidden, displaced, dishonest, collective, or difficult to reconstruct, interpretation could ordinarily presuppose that some writerly or institutional locus existed behind the artifact.

Suspicious reading depended on that presupposition. One could ask who benefits, what interest speaks, what institution repeats itself, what ideology covers itself, what discourse makes this statement possible, what authorial strategy governs the lie. Suspicion could unmask because something was masked. The mask still belonged to a locus.

Generated language alters this relation. It does not abolish ownership at every surrounding level. There are owners of systems, designers of interfaces, builders of infrastructures, deployers, users, clients, regulators, companies, and institutions. These loci matter. They will matter decisively when the system is returned to the data center. But the generated artifact itself may not descend from a primary act of writing in the ordinary sense. It may have been produced by a system trained on many written traces, prompted by a user, shaped by alignment, perhaps assisted by retrieval, and returned as answer-shaped text. The artifact has ancestry. It may not have an authorial act.

Here the cultural veil produces its strongest philosophical effect: **the falsification of the presumption of primary experience.**

The phrase must be heard carefully. It does not mean that primary experience itself has disappeared. Writers still write. Readers still read. Witnesses still witness. Speakers still speak. Human beings still undergo the time of forming a sentence, owning a claim, seeing an event, giving testimony, and answering for what they say.

What has been falsified is the reader's ordinary presumption that an artifact descends from such an experience.

The presumption of primary experience is the background trust by which one receives a text, image, voice, or testimony as having come from some act of writing, seeing, saying, recording, or witnessing. This trust has never been absolute. It has always required criticism. But it functioned as a default horizon of reception. One could doubt the author, challenge the witness, expose the forgery, or criticize the ideology because the artifact was still approached as the trace of a saying, writing, or making that had taken place somewhere.

Generation makes that default structurally contestable.

A generated paragraph may be correct and have no writer in the relevant sense. A generated image may be convincing and have no photographed event. A generated voice may sound like utterance without the act of saying. A generated report may look like the end of analysis without anyone having undergone the inquiry whose form the report imitates.

A limit case makes the structure of this falsification visible.

A user who does not speak Urdu, has never read Urdu poetry or spoken with an Urdu speaker, and has undergone none of the formation that membership in an Urdu literary community would require, can now produce a three-hundred-page Urdu novel. Not a translation. Not a derivative. An original novel that did not exist before the user prompted the system. The output may carry Urdu syntax, vocabulary, idiom, and narrative rhythm. To an Urdu reader encountering the text alone, it may appear to be the work of someone who belongs to an Urdu literary tradition. It is instead the output of a system with statistical coverage of Urdu, prompted by a user who has neither coverage nor formation.

The connection between artifact and membership in a linguistic community can therefore be severed at the production end as well as the reception end. The cultural veil is not only the corrosion of trust in artifacts received. It is also the corrosion of the production-side relation between membership and artifact. Such a novel can enter an Urdu literary field and be received within it. Urdu readers cannot know that no member of their literary tradition wrote it. The asymmetry matters: a user outside a linguistic world may inject artifacts into that world without having undergone the formation that would make the artifact answerable there.

The case also shows that the falsification of the presumption of primary experience operates symmetrically. It falsifies both the presumption that someone wrote from a primary act of writing and the presumption that the writer belongs to the linguistic community whose field the artifact inhabits.

Deception remains one case, often the easiest. The deeper result is a change in the cultural medium. The authentic artifact now appears under the shadow cast by the synthetic one: witness, writer, and photograph must answer suspicions made possible by unreal witnessing, writing without writing, and capture without capture.

The first consequence is authentication dependency.

When artifacts no longer appear self-authenticating, trust migrates outward. The reader turns to platform, publisher, institution, metadata, provenance mark, cryptographic signature, label, archive, reputation, legal certification, or technical guarantee. These may help. They may make certain chains of production more visible. But they do not by themselves restore locus.

A provenance mark may say that an artifact came through a certain channel. It does not make the reader read. A label may disclose that a text was generated. It does not tell the user what source must be returned to. A signature may authenticate a file. It does not produce *Erfahrung*. Technical authentication can mark a chain; it cannot substitute for responsible understanding.

The cultural veil strengthens infrastructure over interpretation. Where trust migrates to authentication systems, whoever controls authentication gains power over reception. This points forward to material locus. Wider infrastructure may also govern trust, provenance, and access to the medium.

The second consequence reaches into the future archive.

Synthetic artifacts do not only confuse present reception. If they are preserved without provenance, they may enter the archive as false ancestors. A future reader may inherit an image of an event that never occurred, a text no one wrote, a voice no one uttered, a style no person formed, a trace whose own source locus was never real in the way the artifact suggests. The archive then receives not only genealogy-loss, but generated ancestry: traces whose descendants may be real while their originating event was not.

The book is not becoming a theory of archives. Detemporalized generation does not merely detach language from its past. It can seed the future with artifacts whose past is already structurally veiled.

The same structure appears in synthetic voice and audio, where the sound of utterance can detach from the act of saying. The visual case is sufficient here because it makes the cultural veil most immediately visible. The book's deeper concern remains generated language, where answerability, reading, judgment, source, and responsibility are most directly at stake.

The danger is not only false artifacts. It is also that true artifacts may lose the ordinary trust by which they were first received as traces of human saying. Authorship, witnessing, and reading may still occur while being displaced by authentication, provenance, and institutional guarantee.

The cultural veil also thins the evidence of the *communio of Rede*.

Chapter 3 showed that such a *communio* appears through practices: greeting, silence, correction, teaching, reading, citation, translation, argument, apprenticeship, and the long work by which another field is entered through *Erfahrung*. These practices make the shared field visible because they expose whether speech lands, fails, resists, or must be learned.

Generated language can simulate many of their artifacts. It can produce the essay, the citation, the translation, the correction, the lesson, the explanation, and the foreign-language surface. But when such artifacts circulate without the practices that made them visible as artifacts of a *communio*, the field itself becomes harder to see. The cultural veil does not only make texts suspicious. It weakens the ordinary evidence that writing came from within a shared linguistic world.

Section 13 of Chapter 3 added the deeper Heideggerian horizon of this thinning. The practices that bear *Rücksicht* and *Nachsicht*, the positive modes of *Fürsorge*, are sustained by *communio* practices. Where those practices contract, the modes they bear thin too.

The cultural veil therefore does not only thin the visible evidence of the *communio*. It thins the practices through which one *Dasein* holds another open as speaker, learner, colleague, student, writer, translator, editor, or judge. The structural-existential effect of this thinning is the *Angriff* that arrives from inside *Mit-sein*: cumulative, indirect, not reducible to the intention of any individual *Dasein*, but decisive at the level of the field.

This is why the veil is most dangerous when the output is good.

A bad image, a crude fake, a fabricated citation, or an obvious hallucination can expose the machinery. A correct paragraph, plausible report, careful explanation, balanced synthesis, or convincing image can pass because it resembles the artifact that source, reading, witness, or judgment would have produced. The better the resemblance, the more easily the locus disappears.

The structural veil therefore does not end at the interface. It can become a condition of culture.

A generated answer may conceal its own lack of locus. A culture of generated answers may make all loci harder to perceive.

21. The first de-veiling

De-veiling cannot be accomplished by remaining indefinitely inside generated commentary. The user may ask the model where an output came from, then ask it to deconstruct that explanation, then ask it to compare possible genealogies. The prose may become more sophisticated while the relation to source remains unchanged. One does not exit the veil by asking the veil to describe itself one more time.

The first de-veiling is conceptual. The user must know what has been received.

The output is not a speaker. It is not a source. It is not testimony. It is not judgment. It is not tradition speaking directly. It is not a horizon. It is not *Erfahrung*. It is generated language shaped by corpus, vectorization, system design, prompting, possible retrieval, and user reception.

This statement should not be used to dismiss the output. It places the output. The output may be material for reading, writing, inquiry, comparison, translation, critique, preparation, or thought. It may be the beginning of an act of understanding. It must not be treated as its completion.

The first de-veiling therefore consists in separating four things that the interface tends to merge.

First, there is the system's operation: statistical and vectorized generation under technical conditions.

Second, there is the source material: written traces from many temporal and historical loci.

Third, there is the output: answer-shaped language produced in the present interaction.

Fourth, there is the user's reception: the lived *Erlebnis* in which the output becomes meaningful, relieving, troubling, useful, or dangerous.

The structural veil makes these four appear as one event: the model answered.

De-veiling separates them again.

Only after that separation can the positive practice begin. The user can ask where source return is needed, where personal judgment is required, where institutional responsibility lies, and where the material system must be named.

The veil is not torn by suspicion alone. It is lifted by placing each relation back where it belongs.

22. Toward *intus legere*

This chapter has described the output-form. The next chapter turns to the reader's practice.

The structural veil appears when linguistic form returns without locus and is received as answer. The output carries temporal, deictic, hermeneutic, and embodied markers. It speaks in the grammar of source, present, evidence, interpretation, dialogue, and care. It lacks the existential conditions these markers normally imply.

The user cannot solve this by refusing all generated language. Refusal may be right in some settings, but it does not by itself create understanding. The user also cannot solve it by mere acceptance, verification, or tool literacy. A verified output may still be unowned. A cited output may still be unread. A useful output may still bypass formation.

What is needed is a practice of reading.

The tradition of *intus legere* will name this practice. Intelligence is not only calculation, classification, or fast synthesis. It is reading within: gathering, distinguishing, entering, returning, and taking responsibility for what one receives. In the time of AI, *intus legere* must become reading between the times. The reader must hold their own temporal locus, return to the source loci where it matters, and later return the apparent ghost to the system's material locus.

The structural veil is not the end of the argument. It is the condition that makes the ethical practice necessary.

The model gives language. The reader must restore time.

NOTES

1. Martin Heidegger, *Sein und Zeit*, 11th ed. (Tübingen: Max Niemeyer, 1967), 167; Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), 211. Heidegger writes: "Der Ausdruck »Gerede« soll hier nicht in einer herabziehenden Bedeutung gebraucht werden. Er bedeutet terminologisch ein positives Phänomen." The warning matters because *Gerede* is not reducible to gossip or error. It is a structural mode of public interpretedness. ↩
2. Heidegger, *Sein und Zeit*, 169; *Being and Time*, 213. The German reads: "Das Gerede ist die Möglichkeit, alles zu verstehen ohne vorgängige Zueignung der Sache." I translate *Zueignung* as appropriation, not possession. The point is not control of the matter but the taking-up through which a saying becomes one's own. ↩
3. Karl Marx, *Capital: A Critique of Political Economy*, vol. 1, trans. Ben Fowkes (London: Penguin, 1990), 163 to 177; Martin Heidegger, "The Question Concerning Technology," in *The Question Concerning Technology and Other Essays*, trans. William Lovitt (New York: Harper & Row, 1977), 3 to 35. The prior work's fetish argument joined Marx's commodity fetishism to Heidegger's *Gestell*: linguistic labor and historical language appear as the model's own knowledge when the mediating relations disappear behind the interface. ↩
4. Hans-Georg Gadamer, *Wahrheit und Methode: Grundzüge einer philosophischen Hermeneutik*, Gesammelte Werke 1 (Tübingen: Mohr Siebeck, 1990), 383 to 384; Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 385 to 386. Gadamer's account of conversation matters here because the chatbot interface imitates turn-taking while lacking the hermeneutic event in which participants are claimed by the matter. ↩

5. Umberto Eco, *I limiti dell'interpretazione* (Milan: Bompiani, 1990); Umberto Eco, *The Limits of Interpretation* (Bloomington: Indiana University Press, 1990); Umberto Eco, *Interpretation and Overinterpretation*, ed. Stefan Collini (Cambridge: Cambridge University Press, 1992). Eco's *intentio operis* gives a precise constraint on generated interpretation. A source must be allowed to resist the interpretation made in its name. ↵
6. Gadamer, *Wahrheit und Methode*, 270 to 290; *Truth and Method*, 278 to 297. Temporal distance is not merely a defect in understanding. It is one condition under which understanding takes place. AI output tends to reduce temporal distance into present availability, which is useful for orientation and dangerous when mistaken for historical encounter. ↵
7. Martin Heidegger, *Die Grundbegriffe der Metaphysik: Welt, Endlichkeit, Einsamkeit*, Gesamtausgabe 29/30 (Frankfurt am Main: Vittorio Klostermann, 1983), §§19 to 38; English in *The Fundamental Concepts of Metaphysics: World, Finitude, Solitude*, trans. William McNeill and Nicholas Walker (Bloomington: Indiana University Press, 1995). The present chapter uses *Langeweile* only to describe the temporal gap that AI output often fills. The full ontology of boredom belongs to the 1929 to 1930 lectures. ↵

Chapter 14. Intus Legere in the Time of AI

1. After the veil

The preceding chapter described the structural veil of AI output. The veil appears when generated language returns in the form of answer, address, synthesis, or judgment while hiding its missing loci: source time, reader time, formation time, operational time, labor, and material infrastructure. It also hides the user's contribution. The model generates; the user supplies world.

The next task is positive. Suspicion may interrupt the surface, but it does not yet give a practice. The user still has to read, judge, return, and own. The question is not only how to expose generated output, but how to receive it without letting it replace understanding.

The tradition of *intus legere* gives this practice its name.

The Latin phrase is usually glossed as “to read within.” Its historical accuracy is complex. The standard etymology of *intelligere* derives the verb from *inter* and *legere*: to choose among, gather among, read between. The medieval tradition often hears in it *intus legere*: to read inwardly or within. Both paths are needed here. Intelligence is gathering, selecting, distinguishing, reading among signs, and reading within what appears.

This is the first correction to the modern term “intelligence.” Intelligence is not first problem-solving, speed, abstraction, prediction, optimization, or output. It is situated reading. A reader stands somewhere, under a question, with a past and a future, and gathers meaning from signs that do not interpret themselves.

This chapter turns the title of the book into a practice. Reading between the times means reading from one’s own temporal locus toward the temporal loci of the sources, while refusing the false present of the model’s output. The reader does not pretend to be nowhere. The reader does not accept the model’s apparent nowhere. The reader reads from here, now, toward what has come from elsewhere and then.

AI can gather, compare, prompt, translate, and prepare. It cannot perform *intus legere* in the strong sense. It has no within.

2. *Legere*

The root *legere* is concrete before it is intellectual. It means to gather, pick, collect, choose, select, and later to read. The movement from gathering things to gathering letters is not accidental. Reading is a form of collection. The eye and mind gather marks, select, join, and return. The act is spatial, temporal, and bodily before it becomes a theory of understanding.

The model of intelligence as abstract processing begins too late. It begins with an already formalized task: recognize a pattern, solve a puzzle, rank alternatives, optimize a result, predict an output. *Legere* begins with a body in a field. One gathers from somewhere, under a need, distinguishing what matters from what must be left. Gathering is never from nowhere.

Cicero’s etymological reflection in *De Natura Deorum* is useful here. In the voice of Balbus, Cicero links several virtues to the family of *legere*: *eligere*, to choose out; *diligere*, to care for or esteem; *relegere*, to retrace or read again; *intellegere*, to understand.¹ The etymologies do not need to satisfy modern historical linguistics in every detail. Their philosophical force lies in the family resemblance they preserve. To understand is to gather with discrimination and care.

The English word “intelligence” has lost much of this field. It now often suggests a measurable faculty, problem-solving capacity, or computational function. The older root resists that reduction by connecting understanding to gathering, care, repetition, and reading.

It also connects intelligence to time. To gather is to move through a field. To read is to move through a sequence. To understand is to return, compare, and let the earlier alter the later. A reader never receives all meaning at once. They pass through marks, lose their way, return to a prior sentence, revise a first sense, and let the whole reshape the part.

A model can process tokens and calculate relations faster than a human reader can follow a page. That speed should not be confused with reading. Reading is not only relation among marks. It is the temporal act by which a situated being gathers meaning and is altered by what resists.

3. Between and within

The two etymological paths, *inter legere* and *intus legere*, preserve two dimensions of intelligence.

Inter legere names reading among or between. The reader stands among signs and distinguishes. This is the critical dimension. One separates the relevant from the irrelevant, the literal from the figurative, the source from the gloss, the plausible from the true, the cited from the merely repeated, the author’s claim from the reader’s use. Reading is selection among possibilities.

Intus legere names reading within. The reader does not remain at the surface. They seek the inner relation, not as a hidden object locked behind the words, but as the depth that appears only through disciplined attention. This is the contemplative dimension. One reads within the text, within oneself, within a tradition, within a situation, within the relation between present question and inherited word.

The two belong together. Reading between without reading within becomes mere selection, classification, or cleverness. Reading within without reading between becomes inwardness without discipline, projection without correction. Intelligence requires both: distinction among signs and inward gathering of their sense.

A language model is strong at the first appearance of *inter legere*. It can sort, compare, classify, identify options, generate contrasts, place one text beside another, name tensions, and produce plausible structures.

It cannot perform *intus legere*. It has no interior temporal life in which the signs can be gathered as meaning. It has no *verbum interius*, no inner word. It has no memory in the Augustinian sense, no conscience before a text, no silence in which a word becomes one's own. It produces outer language without inward reading.

This does not make its comparisons useless. It makes them preliminary. The user may use generated distinctions as material for reading. The user must not treat them as the completion of reading.

4. Augustine's inner word

Augustine gives the tradition of inner reading its first great philosophical form. He places understanding inside memory, time, attention, signs, and the inner word.

In *Confessions* X, memory is not a storage container. It is a vast inner space in which the self wanders, searches, remembers, and is returned to itself. Augustine writes of the "fields and spacious palaces of memory," *campos et lata praetoria memoriae*.² The phrase is spatial, but its content is temporal. Memory gathers what has been, and the soul reads itself through what memory holds and withholds.

In *Confessions* XI, Augustine's account of time sharpens. Past and future do not exist as present objects. The soul lives through the present of things past as memory, the present of things present as attention, and the present of things future as expectation. Augustine calls this temporal stretching *distentio animi*, the distension of the soul.³

Reading is bound to this distension. A reader gathers words in attention, retains what has passed, and expects what is to come. They are not outside time, observing a complete object. They are stretched by the act. The text becomes intelligible only because the reader holds together memory, attention, and expectation.

Augustine's doctrine of the *verbum interius*, the inner word, gives the deeper structure. In *De Trinitate*, he distinguishes the exterior word, sounded or written in a language, from the inner word of understanding. The outer word passes through time. The inner word is not a sound. It is the meaning as grasped by the mind, prior to its clothing in this or that language.⁴

The doctrine must not be turned into private mentalism. Augustine's inner word is not a hidden psychological token. It is where signs are taken up into understanding under the light of truth. In *De Magistro*, external words do not teach by themselves. They prompt the learner to turn inward toward what can be understood.⁵ In *De Doctrina Christiana*, interpretation is governed by *caritas*, charity. A reading that does not build love of God and neighbor fails its deepest rule.⁶

Augustine matters here because the outer sign is never enough. It must be read within. The reader must gather, remember, attend, interpret, and be ethically directed. Words can prompt. They cannot replace the inner act.

AI output is exterior word without inner word. It can prompt the user. It cannot consult the inner teacher. It can generate signs. It cannot read them within itself.

5. Signs, things, and charity

Augustine's distinction between signs and things prevents a common error. A sign is not nothing. It points, gathers attention, teaches, warns, consoles, and misleads. Yet the sign is not the thing. The reader moves through signs toward the matter.

In *De Doctrina Christiana*, Augustine distinguishes *signa* and *res*, signs and things. Words are signs. Some things are also signs. Scripture is a field of signs that must be interpreted, but interpretation must not remain at the level of verbal play. It must move toward the realities to which the signs direct the reader, and it must be disciplined by love.⁷

Generated output is a field of signs that often appears to have reached the thing. The model gives language about the matter, and the user may feel that the matter has been reached. But generated language is not the thing, not the source, not the event, not the person, not the decision, not the responsibility.

The Augustinian rule is simple in form: do not stop at the sign. Let the sign direct you toward the matter, and let the matter correct your reading.

This also corrects arbitrary interpretation. Charity does not mean sentimental acceptance. It means that reading is ethically ordered. The reader is not free to use signs for any end. They must ask what kind of life the reading serves, what relation to the other it produces, and whether it turns the reader toward or away from truth.

Generated output can be used charitably. It can also evade charity. A user may ask for language that sounds generous without generosity, cautious without care, fair without return to the other. The sign can imitate the virtue.

The older tradition reminds us that reading is not only cognitive. It is moral.

6. The body of reading

The medieval monastic tradition gives *intus legere* a body.

Reading was not silent information intake. It was often vocal, rhythmic, tactile, and communal. The monk read with eyes, lips, ears, breath, and posture. The text was murmured, repeated, tasted, memorized, and carried. Jean Leclercq and Ivan Illich make this structure clear. Reading was *ruminatio*, a chewing of the text.⁸

The metaphor should not be romanticized. Its philosophical force is exact. Reading takes time because the text must form the reader. The word is not consumed as content. It is taken into the body and returned through memory, prayer, speech, and action.

Hugh of Saint Victor's *Didascalicon* gives this discipline a structure. The passage from *lectio* to *meditatio* to *contemplatio* is not a list of techniques. It is a temporal order of formation. Reading gathers the words. Meditation turns them over. Contemplation lets the reader see what the prior acts have opened.⁹

This order is nearly the opposite of the interface's promise. The interface gives answer-shaped language quickly; monastic reading gives delay. It does not treat difficulty as a failure of access. It treats difficulty as part of formation.

The book need not idealize medieval reading or reject modern tools. The older practice preserves a truth about reading that technical speed can hide: some meaning appears only through repetition, delay, return, and bodily attention.

Illich's warning about the shift to computer literacy belongs here. His concern is not that screens exist. It is that reading can become manipulation of signs without the bodily and temporal discipline that made reading formative.¹⁰ A person may handle more texts while reading less. A culture may become rich in access and poor in attention.

AI intensifies this problem. It does not only give access to texts. It gives language instead of the user's passage through them. The reader can receive the fruit without the chewing. The metaphor is old, but the danger is contemporary.

7. Aquinas and inward reading

Aquinas gives the phrase *intus legere* its scholastic precision. In *Summa Theologiae* II-II, q. 8, a. 1, his reply begins with the etymology:

“Nomen intellectus quandam intimam cognitionem importat: dicitur enim intelligere quasi intus legere.”

“The name intellect implies a certain intimate knowledge, for to understand is said as though to read inwardly.”¹¹

The passage gives the chapter its most direct source. Understanding is intimate knowledge. The intellect does not remain with exterior accidents; it penetrates toward what the thing is, *quod quid est*. Sensitive knowledge receives the exterior sensible features of a thing. Intellective knowledge reads within toward essence.

This inwardness is not disembodied. Aquinas repeatedly insists that the human intellect, in this life, understands by turning to phantasms, *conversio ad phantasmata*.¹² The mind does not float above body and world. It receives sensory images from embodied life, then abstracts intelligible form through the work of intellect. The human being reads inwardly from material givenness.

This blocks two errors at once.

The first error is reducing intelligence to surface recognition. Aquinas’s *intellectus* is not the detection of external features. It is a penetration into the inner intelligibility of the thing.

The second error is treating inwardness as bodiless purity. The human intellect must return to phantasms. It needs the body, the senses, and the world. Its inward reading begins from an embodied locus.

AI can detect and generate surfaces with enormous skill. It can process images, words, patterns, and relations; produce an account of essence; define *intus legere*. But it has no phantasms in Aquinas’s sense, no sensory life from which intelligible form is abstracted, no embodied soul turning back toward images. It manipulates signs that human beings have already produced from embodied life.

Aquinas gives a precise limit. AI can generate language about inward reading. It cannot read inwardly because it has no interior act of intellect anchored in body and world.

8. Traces and ascent

Bonaventure gives the tradition another image: the world as a field of traces, *vestigia*. In the *Itinerarium Mentis in Deum*, the created world is not mute matter. It bears traces that the mind can read. The soul reads outward in the world, inward in itself, and upward toward God.¹³

The theological frame is not the book's own final frame. The structural insight still matters. Intelligence is a passage through traces. The reader does not invent the trace. The trace is given. But the trace does not interpret itself. It requires a situated reader whose memory, attention, will, and love are engaged.

Bonaventure also links intelligence to memory and will. The mind is not a processing unit. It is a temporal and ethical being. It remembers, understands, and loves. These dimensions belong together. Reading the world without love would be incomplete. Reading oneself without memory would be impossible. Reading signs without will would not lead to action.

Translated cautiously into the present argument: the reader of AI output must remember, understand, and will. They must remember that the output comes from traces, understand that the signs require return, and will to own or refuse the claim.

The model cannot do this. It processes traces without memory in the human sense, without love, without will. It can say that the world is a text. It cannot be addressed by the text of the world.

The language of ascent should not be overused. Still, it clarifies the difference between reading and extraction. Extraction takes the trace as resource. Reading lets the trace call the reader beyond first use. AI is powerful at extraction. The user must decide whether extraction returns to reading.

9. Shared intellect and singular locus

The medieval Islamic and Jewish traditions give a useful complication. In Al-Farabi, Avicenna, Averroes, and Maimonides, intelligence is often discussed through the relation between human faculties and a separate or active intellect. The structures differ, but they share one question: how does a finite human being participate in intelligibility that exceeds the individual.

Avicenna gives the human soul internal faculties that prepare the intellect to receive intelligible forms. Averroes gives a more difficult account of shared or separate intellect, while individual thinking still depends on phantasms and cogitative power. Maimonides joins Aristotelian and Jewish traditions in an account where intellectual perfection requires disciplined preparation, ethical formation, and interpretation of signs.¹⁴

The details cannot be developed fully here. The chapter needs one point: even when intellect is treated as transcending the individual, human intelligence still requires a finite locus. The human being prepares, receives, interprets, and is formed. Intelligence is not a disembodied algorithm. It is the relation between a situated agent and a truth that exceeds them.

AI is often described as if scale and shared access overcame individual finitude. The model appears as a collective intellect trained on the written residue of many minds. It seems to speak from a larger-than-human archive. Yet a larger archive does not create a locus. It gives more traces. It does not give the act of receiving truth from within a finite life.

The medieval traditions preserve a paradox the AI discourse often forgets. Intelligence can reach beyond the individual only because a finite agent receives, prepares, and answers. Universality does not abolish locus. It requires it.

The model has breadth without inward reception.

10. The humanist return to sources

The Renaissance humanists recover intelligence as philological reading. Their call *ad fontes*, to the sources, was not only a scholarly slogan. It was a discipline of return. One returns to languages, manuscripts, variants, historical contexts, rhetorical forms, and the grain of usage. The reader refuses to accept inherited formula when the source can be read again.

Petrarch's relation to Augustine, Erasmus's training of linguistic abundance, and Valla's philological critique all belong here.¹⁵ Petrarch reads antiquity and Christianity through personal and historical interiority. Erasmus trains the reader's mind by multiplying verbal resources so that thought can answer changing situations. Valla shows that false authority can be exposed by exact attention to language and history.

Valla's critique of the *Donation of Constantine* is especially important. He did not expose the forgery by general suspicion. He read historically. He noticed language that did not belong to the alleged time. Philology became a temporal discipline. The words carried a wrong time. Intelligence meant hearing the anachronism.

This is one of the clearest precedents for reading between the times. The reader must hear whether a claim belongs to its supposed source time. They must notice when a generated synthesis makes a medieval claim sound modern, a modern concept sound ancient, or a contested source sound settled. AI output often produces smooth temporal translation. Humanist reading teaches suspicion toward smoothness when time matters.

The humanist return to sources also shows that reading is not antiquarian. Returning to the source is not retreat into the past. It is a way of freeing the present from inherited error. The source can correct the current formula. The older text can make the present less self-evident.

AI output often prevents this when used badly. It gives the inherited formula in polished form. The humanist practice says: go back. Read the source. Check the language. Let the source resist the formula.

11. The modern reduction

The modern period changes the meaning of intelligence. It increasingly becomes measurable capacity, problem-solving, pattern recognition, symbolic manipulation, and later computation. Psychometrics, mass education, cognitive testing, and computational theories of mind all contribute to the shift.¹⁶

This shift is not without reason. Measurement can correct illusion. Tests can reveal capacities hidden by class, patronage, or prejudice. Formal tasks can serve research. Computation has explained and supported many cognitive functions. The reduction is not that these methods exist. The reduction is that they begin to define intelligence as such.

Under this regime, intelligence becomes detachable from reading, history, body, and moral formation. It can be measured, modeled as information processing, and abstracted from relation to text, world, tradition, and responsibility. It becomes a capacity that can be imagined as placeless.

The computational theory of mind gives the strongest form of this abstraction: thought becomes symbol manipulation, the brain a processor, intelligence the execution of formal operations. This does not describe all cognitive science, but it gives the conceptual background for calling a machine intelligent when it produces successful output.

The term “artificial intelligence” inherits this reduction. The companion paper already argued that present systems are better described as artificial fluency: fluent generation of linguistic form, not intelligence in the strong sense.¹⁷ The present chapter adds the older genealogy. If intelligence is *intus legere*, then artificial fluency is not intelligence because it has no inward reading. It has no locus from which to gather meaning as its own.

The model can solve tasks, pass tests, imitate explanation, and generate useful language. None of these capacities proves *intus legere*.

12. The recovery of reading

Contemporary hermeneutics recovers what the reduction obscures. Gadamer makes understanding an event of historically effected consciousness. The reader does not stand outside history. They begin with prejudgments, language, and tradition. Understanding happens when those prejudgments are placed at risk before the matter.¹⁸

Ricoeur adds textual distancing and appropriation. A text is detached from its first situation, but this distance lets it project a world. The reader appropriates that world not by possessing it, but by letting it alter their own possibilities.¹⁹

Hadot adds the ancient practice of philosophy as spiritual exercise. A philosophical text is not only doctrine. It is a form of training, designed to reshape perception, desire, judgment, and life.²⁰

Steiner, Illich, Wolf, and Carr add modern warnings about the fragility of reading. Reading requires time, attention, vulnerability, and bodily discipline. It is not identical with consuming information. It can be weakened by media habits that reward scanning, speed, and interruption.²¹

These authors do not say the same thing. They converge on one point needed here: reading is a temporal formation of the reader. It is not decoding, information access, or text processing. Reading is a human practice in which signs are gathered from a locus and returned to life.

Intus legere becomes the positive practice after the structural veil. The veil gives output as if reading had already been done. *Intus legere* says that the work remains with the reader.

13. Reading within a communio

Chapter 3 added another sense of “within.”

To read within is not only to read within the text, within oneself, within a tradition, or within the relation between present question and inherited word. It is also to read within a *communio of Rede*, where the words can land because the reader has undergone enough of the field to hear their resistance.

This does not make reading an identity credential.

A reader may read from inside a natal *communio*. A reader may read from inside a deeply acquired *communio*, entered through language, correction, dwelling, friendship, study, error, and *Erfahrung*. A reader may read from partial acquired membership, knowing that some words land while others still resist from a distance. A reader may also read from outside a *communio*, if they hold that outside honestly and do not mistake explanation for entry.

These distinctions matter because *intus legere* is not possession of a cultural title. It is not the claim that only the native reader can read, nor the opposite claim that any reader can enter any field by access alone. The criterion is formation, not identity.

To read within a *communio of Rede* is not to possess an identity credential. It is to have undergone enough of the field that its words can resist one's first understanding from within.

Generated rendering simulates another kind of arrival. It can produce the language of a field, translate it, summarize it, compare it, and make it available to a reader who has not entered the field. This may help the reader begin. It must not be mistaken for the reader's entry. The model can render across *communiones*. The reader must still learn where the rendering does not land.

Reading within a *communio* is therefore not innocent of the transformation that Section 13 of Chapter 3 named. The communiones in which words can resist the reader's first understanding are themselves exposed to thinning under conditions of widespread model adoption. The reader who undertakes *intus legere* does so inside this transformation, not above it.

The discipline does not preserve the *communio* by intention alone. It operates as one of the projected possibilities the situation still discloses: a way of holding open the practices through which a *communio* remains visible as a living field of speech.

14. AI cannot read within

AI cannot read within because it has no within.

The phrase is not a denial that models process language. They do. It is not a denial that they can summarize, compare, translate, classify, or generate interpretations. They can. It is not a denial that their output may help a human reader. It often can.

The claim is ontological. The model has no inner word, no memory as temporal self-relation, no ethical direction of interpretation, no body turning back to phantasms, no historical horizon at risk, no inward place where signs are gathered into owned meaning. It has hidden states, activations, vectors, features, and context. These are technical structures. They are not *verbum interius*.

The language of AI often borrows the language of reading. A model “reads” documents, “understands” a prompt, “knows” a source, or “reasons” through a problem. Such terms may function as shorthand in engineering or user interface. In a philosophical account, they mislead.

The model does not read a document as a reader does. It processes input and generates output. It may identify relevant passages, summarize accurately, and answer questions. It does not stand before the text as a historical other, let the text correct its prejudgments, risk itself before the matter, or return from the text to a changed life.

Brian Cantwell Smith’s distinction helps here. AI reckons. Judgment requires a different relation to reality. It requires care, commitment, and answerability.²² The *intus legere* tradition names the same boundary from the side of reading. Reckoning can operate on signs. Reading within requires a locus.

The model can reckon with language. The reader must read.

15. The reader’s locus

The reader’s locus is not a bias to be eliminated. It is the condition of reading.

This claim can be misunderstood. It does not mean that any reading is valid because it comes from somewhere. It means that no reading occurs from nowhere. A reader has a language, history, body, discipline, institution, and question. These conditions shape the act. They can distort. They can also open.

Gadamer's rehabilitation of prejudice belongs here. A prejudice, *Vorurteil*, is a prior judgment. Some prejudices blind. Others make understanding possible. The task is not to purge all prejudice and become pure method. The task is to expose one's fore-understanding to the claim of the matter.²³

In the time of AI, the reader's locus must become explicit. The model's output often appears as placeless synthesis. The user must respond by naming their own place.

What am I asking, why now, and what do I need this for?

What am I tempted to avoid, and what return would make this claim ownable?

What responsibility will I bear if I use this sentence?

These questions are not private moral exercises. They are hermeneutic conditions. Without them, the output will decide the frame by default. The generated answer will give the user a path of least resistance. The user may then confuse the output's order with their own understanding.

To hold one's locus is not to reject assistance. It is to use assistance without surrendering the question.

16. The source's locus

The source's locus is not always past. A living world may be contemporary and still historically other. The reader may stand before a present language, institution, or tradition that cannot be entered by explanation alone. To read within such a locus may require the same discipline as reading the past: patience, distance, correction, and the transformation of the reader's own horizon.

The reader must also return to the source's locus. This is the second side of the practice, though the full account belongs to the next chapter.

A source has a time, language, genre, question, and relation to prior texts and later receptions. It may have an author, school, institution, audience, censor, translator, editor, or archive. It may have wounds. It may have silences.

AI output often weakens these relations. A source becomes content; doctrine becomes summary; dispute becomes balanced position; sentence becomes paraphrase; tradition becomes theme. The output may be accurate enough to orient, but the source's locus recedes.

Intus legere requires return. The reader asks not only what the source says, but where the saying stands. A sentence in Augustine changes when detached from memory, time, conversion, and prayer. A sentence in Aquinas changes when detached from scholastic disputation and theological question. A sentence in Valla changes when detached from philological struggle and ecclesial power.

This return is not always equally necessary. Some uses of AI are light: a draft, a translation check, a list of search terms, a quick orientation. But where truth, interpretation, judgment, person, law, health, history, or public action is at stake, source return becomes part of the work.

The generated answer is not the source's voice. It is a derivative surface that may help the reader find the source.

17. Reading between the times

Reading between the times names the practice proper to this book.

The phrase carries three relations.

First, the reader reads from their own time. They do not pretend to be neutral. They name their question, responsibility, and use.

Second, the reader reads toward the source's time. They return to the author, text, archive, event, genre, and historical world where the matter requires it.

Third, the reader reads through the model's false time. AI output arrives in a generated present, as if source time, corpus time, training time, retrieval time, and user time had already been gathered. Reading between the times separates them again.

This practice is not nostalgia for a world before AI. It is a discipline for using AI without letting it abolish temporal relation. The user may let the model help gather signs. They may not let the model stand in for the act of reading within.

The phrase also corrects the title “artificial intelligence.” If intelligence is reading between and within, then what contemporary models provide is not intelligence in the strong sense. It is artificial fluency, a generated linguistic surface that can support or obstruct human intelligence depending on how it is received.²⁴

AI may give a path into reading. It may not be mistaken for the reading itself.

18. The discipline of non-substitution

The practice of *intus legere* requires non-substitution.

AI may assist reading. It must not replace reading.

AI may assist drafting. It must not replace authorship.

AI may assist comparison. It must not replace judgment.

AI may assist return to sources. It must not replace source return.

This rule does not come from hostility to technology. It comes from the nature of the practice. A practice can use instruments without surrendering its internal goods. A musician may use a tuner; it does not hear for them. A translator may use dictionaries, corpora, and machine suggestions; these do not own the translation. A scholar may use indexes, databases, and generated prompts; these do not perform the reading.

The earlier book on professional judgment used the same principle in another field. AI may accelerate representation. It must not replace the undergoing through which judgment becomes answerable. The present chapter gives the reading version. AI may accelerate access to signs. It must not replace the *intus legere* through which signs become understanding.

Non-substitution also protects formation. If a learner never summarizes, they may lose the capacity to summarize. If a writer never drafts, they may lose the capacity to think in sentences. If a reader never stays with difficulty, they may lose the patience by which difficulty becomes meaning. Some friction is waste; some friction is formation.

Section 13 of Chapter 3 gave this distinction its affective ground. Where formation-producing passages contract, the loss is not only practical. It can be felt as *Verzweiflung*: the despair of inhabiting a world whose paths continue to function while becoming narrower. The discipline of non-substitution does not deny this despair. It begins from it.

Non-substitution is not a rule imposed from outside. It is a practice summoned from inside the contraction. It asks whether this use of the model returns the reader to source, practice, correction, and formation, or whether it replaces the passage through which the reader would have become capable. The question is not whether the tool can produce the artifact. Often it can. The question is whether the artifact has taken the place of the passage.

The discipline is not to preserve all friction. It is to know which friction belongs to the practice.

19. The prompt and the question

The age of AI tends to turn questions into prompts. A question asks from within a situation. It may expose the asker. It may not yet know what it seeks. It may need to ripen. A prompt is an instruction designed to elicit output. It may be precise, useful, and efficient. It may also tame the question too soon.

For *intus legere*, the reader's first task is not always to improve the prompt. It is to preserve the question long enough to know what must be asked. A better prompt may produce better output. It may also accelerate a false path if the underlying question remains unformed.

The structural veil thrives on premature prompting. The user feels uncertainty, asks the system, receives a form, and mistakes that form for the maturation of the question. But many questions become clear only through reading, silence, conversation, and failure. The model can support those acts only if the user does not let output close them too soon.

A prompt can become part of reading when it is governed by a question. It becomes a substitute when the question is surrendered to the prompt's form.

The reader's locus matters again. The user must know what kind of question they are asking: orientation, source leads, comparison, language, judgment, or avoidance of beginning.

The same interface can serve each. Only the reader can distinguish them.

20. The place of AI in reading

The danger is not only that AI detemporalizes language. The danger is that the reader may become detemporalizing. A detemporalizing reader treats source, distance, otherness, formation, and responsibility as obstacles to be removed. *Intus legere* treats them as the conditions under which understanding becomes possible.

AI has a legitimate place inside reading when it is placed.

It can help locate sources, suggest foreign terms, map a debate, compare translations, point to tensions, generate questions, expose vague wording, and help a reader notice that they do not yet know enough.

These uses remain subordinate.

AI may help gather signs. It may help suggest paths, sources, translations, distinctions, and questions. But it cannot be the site of return. The return begins when the reader leaves the generated surface and turns toward source, person, language, institution, practice, or accountable judgment.

AI becomes harmful when it closes the return. The summary replaces the book. The paraphrase replaces the passage. The generated interpretation replaces the encounter with the text. The answer replaces the question. The polished paragraph replaces the author's own risk of saying.

The distinction is not always visible in the output. It lies in the use. The same summary can support reading or replace it. The same list can sharpen thought or substitute for it. The same translation can open the source or conceal its difficulty.

AI literacy, understood as knowing how the system works, is not enough. The user also needs reading discipline: knowing when output is material, scaffolding, temptation, or irrelevant.

A system can produce text. It cannot decide the place of that text in a practice of reading.

21. Toward locus-reinjection

The chapter has recovered intelligence as situated reading. The recovery is not historical decoration. It gives the ethical practice its first form.

If the structural veil gives language without visible locus, *intus legere* restores the reader's discipline before the loci are named. It says that intelligence is gathering, distinguishing, reading within and among, and allowing signs to be returned to the matter. It says that the reader's temporal position is not a defect to be erased, that source time must be respected where the claim matters, and that generated fluency is material for reading, not reading itself.

The next chapter will name the practice required when genealogy cannot be fully recovered. Reading within does not guarantee full return. It teaches the reader how to remain responsible when return is partial, difficult, or impossible.

That practice is locus-reinjection.

Locus-reinjection has two directions. The user holds their own temporal locus. The user returns to the sources' temporal loci where it matters. Later, the practice will also include the material locus of the system: the data center, infrastructure, ownership, labor, and costs concealed behind the output.

This chapter prepares that practice by recovering the older meaning of intelligence. To be intelligent is not to speak fluently from nowhere. It is to read from somewhere.

**AI cannot read within because it has no
within. The human reader can. That capacity
is no longer only scholarly. It is ethical.**

NOTES

1. Cicero, *De Natura Deorum*, 2.72. Latin text in M. Tulli Ciceronis *De Natura Deorum*, ed. Arthur Stanley Pease (Cambridge, MA: Harvard University Press, 1955). Cicero's etymologies are not used here as strict modern philology. They preserve a philosophical relation among gathering, choosing, caring, rereading, and understanding. The research dossier treats Cicero's passage as an early witness to intelligence as careful, discriminatory gathering. ↩
2. Augustine, *Confessiones*, X.8.12. Latin text in Augustine, *Confessions*, ed. James J. O'Donnell, 3 vols. (Oxford: Clarendon Press, 1992). English in Augustine, *Confessions*, trans. Henry Chadwick (Oxford: Oxford University Press, 1991). The phrase "fields and spacious palaces of memory" translates Augustine's spatial image for memory, which is also a temporal field of self-return. ↩
3. Augustine, *Confessiones*, XI.20.26 and XI.26.33. Augustine's *distentio animi* is important for this chapter because reading requires a soul stretched through memory, attention, and expectation. The research dossier uses Augustine's temporality as part of the *intus legere* genealogy. ↩
4. Augustine, *De Trinitate*, IX to XV, especially XV.10 to XV.15. Latin text in Augustine, *De Trinitate*, ed. W. J. Mountain and F. Glorie, Corpus Christianorum Series Latina 50 and 50A (Turnhout: Brepols, 1968). English in Augustine, *The Trinity*, trans. Edmund Hill (Brooklyn, NY: New City Press, 1991). The *verbum interius* is not a private spoken sentence inside the mind. It is the inner word of understanding prior to exterior linguistic clothing. ↩
5. Augustine, *De Magistro*, especially 11.38 to 14.46. English in Augustine, *Against the Academicians and The Teacher*, trans. Peter King (Indianapolis: Hackett, 1995). Augustine's claim that external words prompt rather than directly cause understanding is crucial for distinguishing AI output as prompt from reading as inner act. ↩
6. Augustine, *De Doctrina Christiana*, I.35 to I.40 and III.10 to III.15. English in Augustine, *Teaching Christianity*, trans. Edmund Hill (Brooklyn, NY: New City Press, 1996). *Caritas*

gives interpretation an ethical measure. A reading can be formally clever and still fail if it does not serve the right relation to God and neighbor. ↩

7. Augustine, *De Doctrina Christiana*, I.2.2. The distinction between *res* and *signa* is used here as a discipline of AI reception: generated language is a field of signs, not the thing, source, or judgment itself. ↩
8. Jean Leclercq, *The Love of Learning and the Desire for God: A Study of Monastic Culture*, trans. Catharine Misrahi (New York: Fordham University Press, 1982); Ivan Illich, *In the Vineyard of the Text: A Commentary to Hugh's Didascalicon* (Chicago: University of Chicago Press, 1993). The research dossier emphasizes monastic *ruminatio* as a bodily and temporal practice of reading. ↩
9. Hugh of Saint Victor, *Didascalicon de studio legendi*, ed. Charles Henry Buttmer (Washington, DC: Catholic University of America Press, 1939); English in Hugh of Saint Victor, *The Didascalicon of Hugh of Saint Victor: A Medieval Guide to the Arts*, trans. Jerome Taylor (New York: Columbia University Press, 1961). Hugh's sequence of reading, meditation, and contemplation gives a medieval form to reading as temporal formation. ↩
10. Illich, *In the Vineyard of the Text*. Illich's criticism of "computer literacy" should not be read as a simple rejection of digital tools. It is a warning that the habit of reading can be replaced by the manipulation of signs. The research dossier pairs Illich with Wolf and Carr in the modern recovery of reading as embodied and temporal. ↩
11. Thomas Aquinas, *Summa Theologiae*, II-II, q. 8, a. 1, respondeo. Latin text in *Sancti Thomae Aquinatis Doctoris Angelici Opera Omnia*, Leonine edition (Rome: Commissio Leonina, 1888 to present). English in Thomas Aquinas, *Summa Theologiae*, trans. Fathers of the English Dominican Province (London: Burns, Oates and Washbourne, 1911 to 1925). The Latin reads: "*Nomen intellectus quandam intimam cognitionem importat: dicitur enim intelligere quasi intus legere.*" The research dossier treats this passage as the precise scholastic formulation of the *intus legere* tradition. ↩
12. Aquinas, *Summa Theologiae*, I, q. 84, a. 7. The human intellect's need to turn to phantasms prevents inward reading from becoming bodiless abstraction. This matters for the AI argument because a system can manipulate signs without the embodied sensory ground Aquinas gives to human understanding. ↩
13. Bonaventure, *Itinerarium Mentis in Deum*, especially I to III. Latin text in *S. Bonaventurae Opera Omnia*, vol. 5 (Quaracchi: Collegium S. Bonaventurae, 1891). English in Bonaventure, *The Journey of the Mind into God*, trans. Philotheus Boehner (Indianapolis: Hackett, 1993). The research dossier uses Bonaventure's *vestigia* and *speculum* vocabulary to show reading as the gathering of traces in world and soul. ↩
14. Al-Farabi, *On the Perfect State*, trans. Richard Walzer (Oxford: Clarendon Press, 1985); Avicenna, *The Metaphysics of The Healing*, trans. Michael E. Marmura (Provo, UT: Brigham Young University Press, 2005); Averroes, *Long Commentary on the De Anima of Aristotle*, trans. Richard C. Taylor (New Haven: Yale University Press, 2009); Maimonides, *The Guide of the Perplexed*, trans. Shlomo Pines (Chicago: University of Chicago Press, 1963). The chapter uses these traditions only for the structural claim that participation in

intelligibility still requires finite human preparation and locus. The research dossier states this point explicitly for the Islamic and Jewish materials. ↩

15. Petrarch, *Secretum*, trans. Carol E. Quillen (Boston: Bedford, 2003); Erasmus, *De Copia*, trans. Betty I. Knott, in *Collected Works of Erasmus*, vol. 24 (Toronto: University of Toronto Press, 1978); Lorenzo Valla, *Dialectical Disputations*, trans. Brian P. Copenhaver and Lodi Nauta (Cambridge, MA: Harvard University Press, 2012); Lorenzo Valla, *On the Donation of Constantine*, trans. G. W. Bowersock (Cambridge, MA: Harvard University Press, 2007). The research dossier treats Renaissance humanism as a recovery of intelligence as situated, philological return to sources. ↩
16. Adrian Wooldridge, *Measuring the Mind: Education and Psychology in England c.1860 to c.1990* (Cambridge: Cambridge University Press, 1994); Stephen Jay Gould, *The Mismeasure of Man*, rev. ed. (New York: W. W. Norton, 1996); Alfred Binet and Théodore Simon, *The Development of Intelligence in Children*, trans. Elizabeth S. Kite (Baltimore: Williams and Wilkins, 1916); Charles Spearman, "General Intelligence, Objectively Determined and Measured," *American Journal of Psychology* 15, no. 2 (1904): 201 to 292. The research dossier describes this as the shift from intelligence-as-reading to intelligence-as-measured problem-solving. ↩
17. John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," 1955; Alan M. Turing, "Computing Machinery and Intelligence," *Mind* 59, no. 236 (1950): 433 to 460. The companion paper argues that "artificial intelligence" began as a research aspiration and now misnames systems whose present achievement is better called artificial fluency. ↩
18. Hans-Georg Gadamer, *Wahrheit und Methode: Grundzüge einer philosophischen Hermeneutik*, Gesammelte Werke 1 (Tübingen: Mohr Siebeck, 1990); English in *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004). Gadamer's rehabilitation of prejudice and his account of historically effected consciousness ground the reader's locus as a condition of understanding. ↩
19. Paul Ricoeur, *Du texte à l'action* (Paris: Seuil, 1986); English in *From Text to Action: Essays in Hermeneutics II*, trans. Kathleen Blamey and John B. Thompson (Evanston, IL: Northwestern University Press, 1991). Ricoeur's distanciation and appropriation prevent both naive authorial recovery and arbitrary reader projection. ↩
20. Pierre Hadot, *Exercices spirituels et philosophie antique* (Paris: Études Augustiniennes, 1981); English in *Philosophy as a Way of Life*, ed. Arnold I. Davidson, trans. Michael Chase (Oxford: Blackwell, 1995). Hadot matters here because philosophical reading is not only conceptual acquisition. It is a practice that forms the reader. ↩
21. George Steiner, *Real Presences* (Chicago: University of Chicago Press, 1989); Maryanne Wolf, *Proust and the Squid: The Story and Science of the Reading Brain* (New York: Harper, 2007); Maryanne Wolf, *Reader, Come Home: The Reading Brain in a Digital World* (New York: Harper, 2018); Nicholas Carr, *The Shallows: What the Internet Is Doing to Our Brains* (New York: W. W. Norton, 2010). These authors are not used as a single school.

They support the chapter's claim that reading requires attention, vulnerability, and temporal formation. ↩

22. Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, MA: MIT Press, 2019). Smith's distinction between reckoning and judgment supports the chapter's boundary: AI can compute with signs, but judgment requires a committed relation to the world. The research dossier links Smith directly to the implications of *intus legere* for AI. ↩
23. Gadamer, *Truth and Method*, 278 to 307. The point is not that the reader's prejudgments are always valid. It is that no reading begins without a horizon, and the horizon must be tested by the matter. ↩
24. The prior works make the same boundary through the language of artificial fluency. The companion paper states that fluent generated language has not become knowledge, judgment, or responsibility merely because it arrives in coherent form.

Chapter 15. Locus-Reinjection

1. The practice after *intus legere*

The preceding chapter recovered intelligence as *intus legere*, reading within. That recovery was necessary because the structural veil gives the user answer-shaped language before the work of reading has necessarily occurred. The model generates. The user receives. The output can orient, relieve, organize, and persuade; it may also conceal the temporal and source relations that would let the user own what has been received.

This chapter names that practice: locus-reinjection.

Locus-reinjection is not the fantasy of full recovery. It does not promise that every generated sentence can be traced back to the exact author, page, institution, or human act from which it descended. Often that genealogy has been structurally diluted. Corpus formation has gathered traces into aggregate availability; tokenization has cut them into machine-operable units; vectorization has converted them into relations; training, alignment, retrieval, prompting, and generation have transformed them again. The output speaks from this whole process without usually carrying a readable lineage.

The source may still act through it.

But it may no longer be directly reconstructable.

This is the first law of the practice:

Locus-reinjection does not undo detemporalization. It is what responsibility looks like after detemporalization.

The user does not always recover the lost genealogy. The user recovers responsibility before an obscured trace.

Section 13 of Chapter 3 named the affective situation in which this practice now occurs. Locus-reinjection is not performed from outside the contraction. It is performed by readers, writers, teachers, professionals, students, and institutions already inside a world in which formation-passages may be thinning.

It does not promise to reverse the structural transformation. It performs the return that keeps the conditions of answerability accessible where the field still permits the return. The practice answers *Verzweiflung* not by overcoming it, but by refusing to let the contraction decide in advance that return has become impossible.

Locus-reinjection is the ethical counter-movement to vectorization.

Vectorization detemporalizes. It converts the temporal-existential structure of language into spatial calculability. It gathers written traces, formalizes relations, and returns generated sequence as if source, reading, formation, and judgment had already done their work. Locus-reinjection reverses this tendency where needed. It returns the generated output to the temporal, historical, human, textual, institutional, and responsible sites that the output has flattened or hidden.

The word “reinjection” is not perfect. It can sound technical. It is kept because it names an action. The locus has been removed from the output’s appearance. It must be put back: not mechanically, not always with the same depth, but deliberately. The user must restore the question of where the saying belongs.

The practice has three returns.

First, the user holds their own temporal locus. They do not receive the output as if they were nowhere. They read from here, now, under these conditions, with this question, risk, use, and responsibility.

Second, the user returns to source locus where the matter requires it. Source locus has two major forms: the past source and the contemporary other. The user does not receive the output as if its language had no past or living historical elsewhere. They return to the author, text, case, date, archive, genre, institution, dispute, language, practice, or contemporary source-world from which the claim draws force.

Third, the user returns the system itself to its material locus: the data center, company, cloud, labor, energy, water, ownership, jurisdiction, and contracts through which the output appears.

The threefold direction of de-veiling begins with these acts. Hold the reader's time. Return to source time. Return the system to its material place.

2. Locus, not context

The word "context" is too weak for what is required.

Context can mean surrounding information: a paragraph, file, prompt window, retrieval set, user profile, disciplinary background, or social situation. In machine learning, context often means an operable input range. In ordinary speech, it can mean whatever helps explain something. The word is useful, but too loose.

Locus is stricter.

A locus is not only information around a statement. It is the site from which a statement can be said, read, used, and owned. It includes time, place, body, role, history, source, institution, question, and responsibility. It names answerability, not merely background.

A source has context. A source also has locus. A sentence in Augustine belongs to a theological and temporal site. A passage in Heidegger belongs to a German philosophical vocabulary, a historical moment, an ontological problem, and later political difficulty. A legal clause belongs to a jurisdiction, enactment history, and possible consequence. A candidate profile belongs to a person, not only a set of comparable traits. A generated answer belongs to a system and a user situation, but appears as if it belonged to neither.

The phrase “context window” shows the danger. A model can receive more context, process a longer document, include prior messages, and retrieve passages. These improvements matter. They can make output more accurate and useful. They do not create locus.

A longer context window gives more text, not a historical horizon. Retrieval gives passages, not reading. Memory gives stored relation, not *Erfahrung*. Tool use gives extended operation, not responsibility.

Locus-reinjection therefore refuses the substitution of context for locus. It does not ask only what information the system had. It asks where the claim must be returned before it can be used as one’s own.

3. Source-locus omission

The missing locus must be named precisely.

It is true that the model lacks a speakerly locus. It has no body, lived present, being-toward-death, having-been, projected future, historical self, or answerability. It can generate “I,” “here,” “now,” “I remember,” “I think,” and “I will explain,” but no *Dasein* gathers in those phrases.

Yet the more important omission in AI output is not the model’s temporal locus. The model never had such a locus to omit. The deeper omission is the source’s temporal-historical locus.

The language on which the model operates came from somewhere. It came from authors, translators, editors, communities, workers, institutions, platforms, states, professions, schools, courts, laboratories, churches, companies, forums, archives, and homes. It came from people writing under pressure, habit, desire, boredom, duty, fear, ambition, hope, anger, or care. It came from worlds.

Corpus formation gathers those worlds as traces. Vectorization formalizes the traces. Generation returns language from the formalized field. What disappears in the output is the source's historical situatedness.

The output therefore has a double structure:

source-locus omission plus generated speaker-form.

It omits the source loci from which its language descends, while appearing in the form of a situated answer.

The output can therefore sound neutral. It is not neutral because it has transcended standpoint. It is neutralized because its standpoints have been detached from their genealogies.

The absence of an explicit point of view is not objectivity.

It is often the concealment of source-locus omission.

4. Lost genealogy and responsible return

The old demand would be simple: find the source.

The new situation is harder. The source may not be findable. The output may carry traces from many sources, none of which appears as source. A phrase may descend from a genre rather than a passage. A distinction may have been absorbed from hundreds of secondary summaries. A tone may come from institutional writing. A conceptual neighborhood may carry decades of scholarship without naming any scholar. A judgment-like sentence may recombine public discourse, training preferences, system instructions, and the user's prompt.

This does not mean the output comes from nowhere.

It means that its genealogy has been structurally dispersed.

The past, or the source world, may still act through the generated trace. But it acts through a field in which its origin has been diluted, redistributed, neutralized, and made operational. The user receives language whose ancestry is real but not fully reconstructable.

Locus-reinjection therefore cannot mean total genealogy-reconstruction. That demand would be impossible and often conceptually mistaken. The output is not a quotation with missing quotation marks. It is a generated surface produced from a technical transformation of many traces.

Responsible return begins when the user refuses to treat that surface as self-grounding.

Where a claim matters, the user must build a new accountable path. That path may lead to an original source, scholarly edition, legal text, living person, expert community, archive, field practice, dataset documentation, institution, or decision-maker who can answer for the use. Where such return is impossible, the user must acknowledge the impossibility rather than let the sentence pass as universal knowledge.

When genealogy cannot be reconstructed, responsibility begins by refusing to pretend that it has been.

The generated trace can still point, provoke, orient, compare, translate, or gather. But it must not be received as if its loss of genealogy had purified it of history.

Genealogy-loss is not universality.

It is loss.

5. Historical locus beyond pastness

Historical locus does not mean only the past.

A source can be historically distant because it is old, because it belongs to a dead author, archived dispute, inherited tradition, or vanished institution. But a source can also be historically distant because it belongs to a contemporary world that is not ours.

The historical other is not always behind us.

Sometimes it is beside us.

A Western reader trying to understand contemporary Chinese thinking encounters this difficulty. The difficulty is not simply lack of information. The concepts may belong to another language, educational formation, political vocabulary, civilizational memory, institutional order, and geopolitical horizon. The other locus is contemporary, but still historical. It has been formed. It carries a world.

AI can mediate that locus. It can summarize Chinese political concepts, translate terms, compare traditions, explain policy language, or generate a balanced overview. Such output may help the reader begin. But it cannot substitute for the *Erfahrung* required to enter that locus. Years of learning Chinese, living in China, being corrected by interlocutors, misunderstanding and being misunderstood, encountering institutions, discovering the limits of one's inherited categories, and returning changed: these are not informational steps. They are historical formation.

One does not enter a historical locus merely by receiving an explanation of it.

One enters by being altered by the encounter.

The same applies to the past. Dialogue with tradition is not simply access to earlier texts. It is also dialogue with oneself, because the tradition is not only outside us. It has already formed us. The reader discovers within themselves the products of the tradition and their distance from it. The past is not simply behind; it is active inside the present horizon from which the reader reads.

Historical positioning therefore has two forms. One addresses the past that has formed us. The other addresses the contemporary other that stands beside us but does not belong to our inherited world. AI can collapse both into present availability. Locus-reinjection restores their difference.

The past must be allowed to remain past.

The contemporary other must be allowed to remain other.

6. Impossible access and closed access

AI opens an impossible access.

A user can prompt Augustine, Heidegger, Aquinas, Gadamer, Chinese political thought, Islamic jurisprudence, financial regulation, or the language of a profession they have never entered. The output arrives in the present. It can compare, explain, summarize, and translate. It can make distance feel unnecessary.

This access is impossible because it gives the user the form of arrival without the passage. The source world appears before the reader has approached it. The contemporary other appears before the reader has been transformed by encounter. A tradition appears before the reader has discovered what in themselves belongs to it and what resists it.

But the same access can close a possible access.

It closes access when the generated explanation replaces the time of approach. It closes access when the user receives a rendered past instead of a historical text, a rendered other instead of a living world, a rendered doctrine instead of an argument, or a rendered culture instead of practices and tensions. It closes access when the output makes the source too present.

AI gives the past back as availability.

Historical reading gives the past back as distance.

AI gives the other back as explanation.

Historical encounter gives the other back as resistance.

This does not mean that AI cannot help historical work. It means that its help is dangerous when the user forgets what kind of access has been opened. The generated output is not the source's return. It is the present rendering of source traces after detemporalization.

The reader must not confuse the rendering with the encounter.

7. The double direction of de-veiling

The structural veil hides by merging. It merges the system's operation, the source traces, the generated output, and the user's reception into one apparent event: the model answered. De-veiling separates them.

The first direction of de-veiling is toward the user. The output appears ready for use. The user must return to their own situation before using it. They must locate the present task, intended use, stakes, limits of knowledge, and responsibility borne. They must know whether they are asking for orientation, drafting help, source leads, comparison, judgment, or relief from difficulty.

The second direction of de-veiling is toward the sources. The output appears as synthesis. Where the matter requires it, the user must return to the traces that the synthesis compresses: authors, texts, cases, dates, events, and disagreements the model's prose has made too present.

The two directions are not symmetrical, but they belong together. A user who holds their own locus without returning to sources risks projection. They may make the output serve their present need without letting the source resist. A user who returns to sources without holding their own locus risks antiquarianism or false neutrality. They may gather sources without knowing why the question matters now.

Reading between the times requires both.

The user reads from their own time toward the source's time. The generated output stands between them as material, not as judge. It can help form the path, but it cannot walk it for the reader.

8. The user's temporal locus

The first act of locus-reinjection is to hold one's own temporal locus.

This is not a confession of bias. It is the condition of reading. A reader always reads from somewhere: a language, discipline, body, history, question, practical need, institution, time of life, and relation to others. These conditions do not make truth impossible. They make understanding possible. They can blind the reader, and they can give access. The task is to let the matter test them.

In the use of AI, this requirement becomes sharper because the output often presents itself as if it had no locus. It gives the user a structured answer before the user has articulated their own place in the question. The output can then decide the frame by default: categories, balance, order, examples, omissions, and scale of relevance.

Holding one's locus interrupts that default.

The user states, at least to themselves, the kind of act being undertaken. A student is not a judge. A doctor is not a patient. A reader of Augustine is not Augustine. A hiring professional is not the candidate. A citizen reading political conflict is not a neutral archive. A writer drafting a sentence is not a machine arranging plausible continuations.

Each locus carries different obligations. The student must learn. The judge must decide. The doctor must care under evidence and practice. The reader must let the source resist. The hiring professional must encounter the person. The citizen must attend to consequence. The writer must own the sentence.

The model does not know which obligation belongs to the user unless the user gives it in the prompt. Even then, it does not bear the obligation. The user bears it.

To hold one's locus is to refuse to let generated form decide the kind of responsibility one has.

9. The prompt and the question

The prompt is not the question.

A prompt is an instruction to a system. It asks for output under constraints. It can be precise, useful, and disciplined. A question is broader. It arises from not-knowing, desire, worry, encounter, discipline, and responsibility. A question may not yet know how to ask itself. It may need time before it can become promptable.

Locus-reinjection protects the question from premature conversion into prompt. The user does not only ask how to obtain a better answer. They ask whether the prompt has preserved the question. A precise prompt can still betray the question if it reduces the matter too early. A vague prompt can still serve inquiry if it helps the user discover what must be asked.

This is especially important when the user seeks judgment. A prompt can ask for risks, options, arguments, rankings, and recommendations. The system can supply each. Yet the deeper question may concern responsibility, not options: whether the user is ready to decide, which source must be read, which person must be encountered, which silence must be preserved, or which institution is trying to accelerate past its own duty.

The prompt often seeks output. The question seeks understanding.

AI can support the movement from question to prompt when the user remains in charge. It can help the user clarify terms, identify materials, test distinctions, and see the limits of a first formulation. It becomes dangerous when the prompt replaces the question, and output replaces the formation of the one who asks.

The user does not only refine the prompt. The user returns to the question.

10. Source return and source resistance

The second act of locus-reinjection is source return.

A generated answer may name an author, summarize a doctrine, cite a study, paraphrase a law, interpret a poem, describe a person, or state a fact. The output may be useful. It may also make the source appear too close. A historical text becomes present-language. A disputed claim becomes neutral synthesis. A living person becomes profile. A legal precedent becomes rule. A tradition becomes paragraph.

Source return restores distance.

The reader returns to the text where the text matters, to the author where authorship matters, to the case where application matters, to the date where sequence matters, to the genre where form matters, to the institution where authority matters, and to the person where representation has become too easy.

This return does not require full archival reconstruction for every use. Search terms, grammar suggestions, or rough orientation may need little return. Serious claims require serious return. The more a generated sentence touches truth, person, law, health, money, public action, or formation, the stronger the obligation.

The source is not only evidence. It is also resistance. It can show that the model's synthesis was too smooth, too present, too generic, too modern, too confident, or too detached from the original problem. Source return lets the past resist the present.

This resistance is not inefficiency. It is one of the ways understanding becomes answerable.

11. Temporal distance and application

Gadamer's account of temporal distance is central here. Temporal distance is not merely a gap to be closed. It is the medium in which understanding happens. A historical text does not become intelligible by being stripped of its distance. It becomes intelligible when the reader's present horizon is brought into disciplined relation with the text's historical horizon.¹

Generated output often reduces distance. It converts the old, difficult, alien, disputed, technical, or institutionally embedded source into present prose. That conversion can help a reader begin. It can also weaken the source by making its distance disappear.

Gadamer's concept of application, *Anwendung*, gives the next step. Understanding is not completed first and applied later. Application belongs to understanding itself. A law is understood in being applied to a case. A text is understood in relation to the reader's present question. A tradition is understood as it addresses the present.

Locus-reinjection therefore does not mean returning to sources as dead past. It means allowing source time and user time to meet under a question. The reader does not dissolve into the source's past. The reader also does not dissolve the source into present use. The two times must be held.

Ricoeur's account of distanciation and appropriation gives the same structure differently. Writing frees discourse from the immediate speech event. The text projects a world. The reader appropriates that world by allowing it to alter their own possibilities.² AI output can describe this process. It cannot perform it for the reader.

The user must let the text's world return to their own world, not as generated paraphrase, but as a claim that can alter what they see.

12. Verification and appropriation

Verification is necessary. It is not enough.

Verification asks whether a statement is correct, supported, current, or consistent with a source. It is indispensable in AI use. Generated output may fabricate, misdate, overstate, or cite badly. Serious claims require verification.

Yet a verified claim may still be unappropriated. A user may check that a sentence is correct and still not understand it, confirm that a source exists and still not read it, verify that a legal rule was stated correctly and still not know how it applies, or confirm a candidate's career fact and still not have encountered the person.

Appropriation asks whether the claim has been returned to the practice in which it can be owned. In scholarship, appropriation means the source has been read closely enough that the claim can be defended. In professional judgment, it means the claim has been tested against evidence, person, and world. In public action, it means the speaker can answer for the claim before those affected.

The difference between verification and appropriation is the difference between correctness and ownership.

The earlier works therefore insisted that reappropriation is not editing. Editing improves form. Reappropriation restores answerability. Locus-reinjection gives that act its full temporal structure. A claim becomes answerable only when it is returned to the loci that allow someone to say it as their own.

A correct sentence can still be homeless. Locus-reinjection gives it a place, or leaves it as material.

13. From *Erlebnis* to *Erfahrung*

AI output can become an *Erlebnis* for the user. The user receives language. A field is ordered. A problem becomes sayable. An anxiety is relieved. The event is real, and it may be valuable.

The event does not become *Erfahrung* by occurring.

Erfahrung requires undergoing. The first sense must be tested. The output must meet resistance. The user must return to sources, evidence, persons, practices, and consequences. Something in the user's understanding must be exposed to correction. The generated answer becomes material for *Erfahrung* only when it passes through this return.

This names the main ethical risk of AI use. The system offers the user an *Erlebnis* of understanding before the formative passage has happened. The user may feel after the work while still standing before it. The institution may accept the product of apparent judgment before anyone has undergone the matter.

Locus-reinjection is the practice that converts possible *Erlebnis* into possible *Erfahrung*. It does not guarantee formation. No method can guarantee that. It creates the conditions under which formation may occur. It returns the user to the matter, and lets the matter answer back.

The generated text may then become useful in the right sense: orientation, provocation, draft, comparison, memory aid, or warning. It becomes part of the user's temporal work, not a substitute for it.

The passage from *Erlebnis* to *Erfahrung* is the passage from generated event to owned understanding.

This practice moves in a circle. The argument names the circle rather than concealing it.

Locus-reinjection requires *Erfahrung* in order to be performed. Source return is available only to a reader who has undergone enough formation to feel where a synthesis has gone smooth. The reader who can tell that a generated account of Heidegger reached for the later formulas where the early analysis of *Rede* was needed is a reader who has already read Heidegger. The resistance of the source registers only for one who has been formed enough to register it. A reader without that formation receives the smooth account with nothing in them that the smoothing offends.

The practice presupposes what it protects.

This circularity is not a defect particular to locus-reinjection. It is the hermeneutic circle. Understanding has always moved from a fore-conception toward a matter that revises it, and the fore-conception is itself the deposit of earlier understanding. Heidegger names this the fore-structure of understanding in *Sein und Zeit* §32, and draws the consequence that the task is not to leave the circle but to enter it in the right way. Gadamer names the same structure when he shows that understanding begins from prejudice and is corrected by the matter, in the movement Chapter 3 entered through *Vorurteil* and *Wirkungsgeschichte*. One does not read a source from nowhere and then acquire a horizon. One reads from a horizon already underway, which the reading then alters.

Locus-reinjection is circular because it is a practice of understanding, and understanding is circular.

This is why the model stands outside the practice rather than inside it. The model has no fore-structure that is its own. It has no horizon that a matter could revise, no prejudice it could discover to have been wrong, no prior understanding at stake in the next encounter. It arranges the relations that understanding has already deposited in language, without ever having entered the circle from which those relations came. It produces the output of understanding without the circular movement that constitutes understanding. The reader is in the circle and can be formed by the return. The model arranges the circle's results without ever having been in it.

The thinning named in Section 13 of Chapter 3 can now be stated more exactly. What contracts under widespread model adoption is not formation in a loose sense. It is entry into the hermeneutic circle. Generated output offers the result of the circular movement so fluently that the passages by which a reader would have entered the circle need never be traversed. The reader receives a fore-conception ready-made and forms none of their own to revise. The circle still turns in the output. It no longer turns in the reader.

The circle is real. It is also the reason the *communio* matters.

A practice that presupposed nothing could be taught as a procedure and applied by anyone. Locus-reinjection cannot, because no hermeneutic practice can. It is not a technique laid over a reader from outside. It is the form responsibility takes for a reader who has entered the circle, and it remains possible only while entry remains possible. This is why the practice cannot be purely private. Its condition is collective. The passages through which one reader came into the circle are the passages through which the next reader might come into it. Where those passages are sustained, the practice remains available. Where they thin, it does not merely become harder. It becomes unavailable to those who never entered.

The practice and its condition cannot be separated.

14. AI as aid to return, not site of return

AI may help the reader return.

It may suggest source leads, search terms, translations, bibliographic paths, objections, interpretive differences, and questions for further reading. It may expose a vague formulation, show the reader that they do not yet know enough, or gather an initial map for difficult terrain.

These uses are legitimate when they remain subordinate.

AI can help us reach the archive. It cannot be the archive's historical voice.

AI can help us approach a contemporary other. It cannot be the other's world.

AI can help us formulate the question. It cannot be the question's answerability.

AI can help collect materials. It cannot become the site of historical dialogue itself.

The return begins when the reader leaves the generated surface and turns toward a source, person, institution, practice, archive, case, or decision for which someone can answer. The model may assist the return. It cannot complete it.

Here the structural temptation appears most sharply. Because the model can generate increasingly sophisticated commentary on its own output, the user may remain inside the generated field while believing they have gone deeper. The user asks where the output came from; the model explains. The user asks for deconstruction; the model deconstructs. The user asks for possible genealogies; the model compares. The prose becomes richer. The relation to sources may remain unchanged.

This is the generated regress.

One does not exit the structural veil by asking the veil to describe itself one more time.

AI may aid the return. It cannot be the site of return.

The exit points are not generated commentary. They are source, archive, person, language, institution, case, practice, infrastructure, and accountable judgment. AI can point toward them. The user must go.

15. The variation before the prompt

Chapter 3 begins before the prompt.

It begins from a lived field that later becomes formal variation: an Italian author walking in Almaty, inside a public world articulated in Russian and Kazakh, languages not his own, among people whose speech cannot be presumed to include Italian, while a cat on the road alters the movement of attention.

The scene matters because it does not remain a scene.

It becomes phenomenological material.

In the written form of Chapter 3, the lived origin is not treated as memoir. It is transformed into *eidetische Variation*. The silent walker, the neighbor, the passerby, the Icelander, the cat, the impossible speaking cat, the movement from Italy to Iceland, the projection toward Boccaccio and Cicero, the medieval Iceland variation, and the bilingual entry through *Erfahrung* all belong to the same methodological sequence.

The examples are not illustrations placed after a doctrine.

They are the way the doctrine comes to evidence.

The first field is walking. A body moves through a world already shared. The movement is not first a coordinate shift. It is passage through nearness, direction, possible greeting, possible silence, interruption, hesitation, and practical orientation. A road is not a grid. It is a place of passage, a place where speech could land, fail, be withheld, or never become necessary.

The second field is linguistic non-possession. The author carries Italian, while the public world around him is articulated in Russian and Kazakh. Signs, speech, names, institutions, tones, and practical expectations appear around him. Some are legible. Some are not. The walker can still move, avoid, follow, notice, and respond.

This is not ignorance outside a world.

It is partial non-possession inside a world.

The third field is other people. They pass, approach, cross, wait, speak, and continue their own tasks. They are not objects in a scene. They are co-inhabitants of the road. Yet the author cannot presume that Italian will reach them. The other person appears as someone with whom the world is shared before linguistic reciprocity is secured.

This is not a sociological claim about who speaks which language.

It is a phenomenological structure of address under non-possession.

The fourth field is the cat. The cat is not decorative. It is not a symbol added after the fact. It belongs to the field through which worldhood appears before the word becomes thematic. Before it becomes the word “cat,” it alters movement. The walker notices, slows, turns, avoids, or attends. The road ceases to be only a route. It becomes a place of co-presence, vulnerability, and adjustment.

The cat discloses world before sign.

The language barrier, the other people, and the cat therefore do different work inside the same field.

The language barrier discloses publicness before linguistic ownership.

The other people disclose being-with before linguistic reciprocity.

The cat discloses world before sign.

Walking discloses temporality before map.

Chapter 3 then varies the field more strictly. A silent walker carries *Sprechenkönnen*, the modal being-able-to-speak that belongs to *Rede* even where no word is uttered. A neighbor lets the could-speak land. An Icelander who does not understand the Italian greeting reveals the perimeter of one *communio* of *Rede*. A cat's meow is received equanimously because the animal does not belong to the field of possible linguistic address. The impossible speaking cat produces *Entsetzen* because the framework of belonging breaks.

The variations then move across history.

Boccaccio can still be reached by the contemporary Italian *communio* in one way. Cicero cannot be reached in the same way. Medieval Iceland does not become available simply because time has moved backward. A Russian speaker who has learned Italian through correction, error, pleasure, dwelling, friendship, reading, and time shows that the boundary of a *communio* is real without being sealed.

The boundary is crossed by *Erfahrung*.

This is where the personal origin and the written method meet. The Almaty walk gave a lived field in which publicness, non-possession, being-with, animal presence, and movement appeared together. Chapter 3 converts that field into formal variation. The variation lets the structures show: *Rede* is modal, shared, bounded, historically deep, mortal, permeable, and unavailable to the model as its own belonging.

The method does not construct these distinctions from outside.

It brings them to evidence.

Husserl names the methodological possibility: return to the thing itself by varying the appearing until invariant structures come into view. Aristotle's *koinonia* and the Latin *communio* name the older horizon of shared life in which speech and action belong to common existence. Heidegger's *Mit-sein* radicalizes this sharedness existentially. *Das Man* names the everyday public interpretedness through which the shared world becomes speakable, repeatable, and disburdening.

From there, *Rede* becomes necessary.

The world is already articulated before detached words are inspected. The road, the signs, the speech not mastered, the people encountered before linguistic reciprocity, the cat, the body's adjustment, and the thought's emergence all show that language does not begin as an inventory of labels. It begins as articulation within a world already underway.

From there, *Gerede* also becomes necessary.

The same public world that lets one move, speak, and understand also lets one repeat without appropriation. Publicness gives access before ownership. The language barrier makes this difference visible in one way. Generated language will make it visible in another. In both cases, language can arrive before the matter has become one's own.

Chapter 3 is the written form of this return.

It does not begin with source names as ornaments. It begins from the phenomenon and returns to sources because the sources become necessary for naming what appeared. Heidegger is needed for *Rede*, *Gerede*, *Mit-sein*, and *das Man*. Gadamer is needed for horizon, history of effects, and the risk of understanding. Ricoeur is needed for narrative time and refiguration. Pareyson is needed for *formatività*, the discovery of form in the time of making.

The chapter therefore does not merely cite phenomenology.

It repeats its possibility from a new locus.

A contemporary high-performing language model was then asked to treat the same philosophical field: language as temporal event, Heidegger, Gadamer, Ricoeur, Pareyson, and the relation of these themes to AI output. The generated answer is reproduced in full in the notes to this chapter.³

The answer is useful.

It is not crude. It identifies the relevant field. It treats language as event rather than as static dictionary. It recognizes that Heidegger, Gadamer, Ricoeur, and Pareyson all resist the reduction of language to a neutral system of labels. It also states that AI output lacks lived time, historical situatedness, and mortality.

Its usefulness is the reason the case matters.

A weak answer would show only failure. A fabricated source, an obvious confusion, or a shallow error would leave the deeper problem untouched. The generated answer is more philosophically revealing because it is largely adequate. It gives a plausible map. It can orient the reader. It can begin inquiry.

It begins, however, from the prompt.

That difference is decisive.

The prompt gives the model the source names and the theme. The model arranges them. Heidegger appears through “Language is the house of Being,” *Ereignis*, mortality, and the temporality of *Dasein*. Gadamer appears through the fusion of horizons. Ricoeur appears through discourse and narrative time. Pareyson appears through interpretation as formative event. The relation to AI output is then drawn through next-token prediction, atemporal statistical weights, archive, reception, and the human reader’s experience of meaning.

The arrangement is coherent.

It is also source relation without source return.

The generated answer receives the field after the field has been named. It does not begin from a lived phenomenon, vary it, and discover why the sources become necessary. It does not move from the language barrier to publicness without ownership. It does not move from other people on the road to being-with before linguistic reciprocity. It does not move from the cat to world before sign. It does not move from the silent walker to modal *Sprechen-können*. It does not move from the Iclander to the perimeter of a *communio*. It does not move from the speaking cat to *Entsetzen*. It does not move from bilingual formation to the crossing of a boundary by *Erfahrung*.

It receives “language as temporal event” as a task.

Chapter 3 receives it as a phenomenon.

This is not a defect in the generated answer. It is its category. The model can arrange the already available conceptual relations because those relations have become linguistic material. It can treat Husserl, Heidegger, Aristotle, Gadamer, Ricoeur, and Pareyson as neighboring terms inside a learned field of discourse. It can produce a convincing path among them.

But the path is generated after the path has become nameable.

The variation belongs elsewhere.

A model can receive “Italian,” “Almaty,” “Russian,” “Kazakh,” “Iclander,” “cat,” “Boccaccio,” “Cicero,” and “communio” as tokens. It cannot receive them as a lived and varied field in which publicness, possible address, animal presence, historical reach, and bodily movement make a question necessary.

This is why the generated answer’s competence does not undo the distinction. The answer can say that human language is temporally situated. It can say that AI lacks lived time. It can say that meaning happens when a human reader receives the output. These claims are close to the argument of this book. Their closeness makes the structural difference visible.

The generated answer gives relation.

It does not give genesis.

Its Heideggerian line also shows the smoothing that accompanies generated synthesis. The answer moves quickly to the later formula, “Language is the house of Being,” and to *Ereignis*. Those terms are not irrelevant. They belong to Heidegger’s later path. But Chapter 3 needs another beginning. It begins with *Rede* before *Sprache*, with discourse as the articulation of intelligibility, and with the claim that words accrue to meanings rather than receiving meanings as detached word-things.

The generated answer names Heidegger.

Chapter 3 needs Heidegger.

The difference is not the difference between reference and no reference. It is the difference between a conceptual neighborhood and a path of necessity.

The same difference appears with *Gerede*. The generated answer does not need Heidegger’s difficult account of public language as the possibility of understanding everything without prior appropriation of the matter. Chapter 3 does. Without *Gerede*, generated language can be described as a simulation of discourse. With *Gerede*, the danger becomes sharper: language can arrive in the form of answer before the matter has been appropriated by the one who receives it.

The generated answer also handles Gadamer in a familiar way. It names the fusion of horizons and the temporal gulf between reader and text. This is intelligible and helpful. Yet source return requires more than the named formula. It requires *Wirkungsgeschichte*, prejudice, the question, *die Sache*, and the risk that the reader’s horizon will be corrected by the source.

The generated answer can describe horizon.

It cannot risk one.

Ricoeur is treated more strongly. The answer distinguishes language as structure from discourse as event, and it names narrative time. Yet Chapter 3 requires the fuller movement of threefold *mimesis*: prefiguration, configuration, and refiguration. A model can produce configuration-like form. It can arrange agents, events, motives, and claims into followable shape. It has no pre-narrative world of action from which that configuration arises. It does not return from the text to a changed life.

The generated answer can name narrative time.

It cannot undergo refiguration.

Pareyson is also compressed. The answer treats interpretation as active, personal, historically situated formation. That is close enough to help a reader begin. But Chapter 3 needs *formatività* more precisely. Form is not the execution of a prior plan. It is discovered in the time of making.

This matters for the present book.

A generated paragraph may assist the author. It may expose a route, suggest a relation, or give language to an intuition not yet stable. But the form of the book lies in the author's temporal work of receiving, rejecting, altering, testing, and owning what is said.

The model can summarize *formatività*.

It cannot undergo *formatività* in the author's place.

The most revealing sentence in the generated answer is its own avowal of limit: "I process language, but I do not *experience* time, history, or mortality."

The sentence is accurate at the level of content.

It is veiled at the level of form.

The model states its lack of lived time, history, and mortality in the first person. It denies the conditions of speakerly locus through the grammar of speakerly avowal. It says "I" while declaring the absence of the temporal existence that would normally anchor such an "I."

The sentence does not fail by being false.

It reveals the structural veil by being formally coherent.

This is where locus-reinjection becomes visible as practice.

The first return is the user locus.

In this case the user locus is not an abstract reader before an abstract answer. It is the authorial locus from which the question arose and then became formal variation: walking in Almaty, speaking Italian, not possessing the surrounding Russian and Kazakh public language, moving among other people whose speech cannot be presumed to include Italian, and encountering a cat whose presence alters attention. In the written chapter, this locus becomes the sequence of variations by which the *communio of Rede* comes to evidence.

The user locus is therefore not private interiority.

It is situated exposure.

The second return is the source locus.

The sources do not enter as external confirmation after the author has already made his claim. They enter because the variations disclose structures that ask for names. Husserl becomes necessary because the lived field is not left as memory. It is varied until structures appear. Aristotle becomes necessary because the road is not private space, but a common world of movement, shared life, possible address, and practical orientation. Heidegger becomes necessary because *Mit-sein*, *das Man*, *Rede*, and *Gerede* name structures disclosed by the field itself.

The sources are therefore not ornaments.

They are places of resistance.

Gadamer resists the fantasy of summary by requiring horizon, history of effects, and exposure to the matter. Ricoeur resists flat sequence by returning time through narrative configuration and refiguration. Pareyson resists the plan imposed from outside by naming the form discovered in making. Chapter 3 returns to these sources because the phenomenon requires them. The readings that ground these resistances are not asserted here for the first time. They were performed in Chapter 3, where *Rede* and *Gerede* were entered in sections 2 and 4, horizon in section 8, narrative time in section 9, and *formatività* in section 10, each with its own apparatus. The present section returns to that work rather than repeating it.

The third return is the material locus.

The human case has material locus. It is not pure thought floating above the world. The road, city, body, weather, signs, other people, language barrier, animal presence, memory, notes, books, editions, translations, drafts, revision, fatigue, and page all belong to the material conditions through which the chapter becomes possible. The thought is not less philosophical because it began while walking. It is philosophical because the walk became available to variation, source return, and written form.

The generated answer also has material locus.

It comes through a prompt, model, training corpus, interface, infrastructure, ownership structure, cloud deployment, and data center. Its apparent voice from nowhere is materially situated elsewhere. The system has no speakerly locus, but it has material locus. It runs somewhere, is powered somewhere, is cooled somewhere, is owned somewhere, and is made available through institutional and corporate arrangements.

The contrast is not spirit against matter.

Both cases pass through material supports.

The difference lies in the relation between support and answerability. The author can return to the walk, the variation, the sources, the draft, and the claim. He can revise the formulation because the formulation belongs to his path. The model cannot return to the prompt as its question, to Almaty as its place, to the cat as encountered interruption, to the Icelander as perimeter, to the speaking cat as rupture, or to Heidegger and Gadamer as sources that changed its understanding.

The generated answer can be inspected.

It can be useful.

It can even help the author see what requires sharper formulation.

This is assistance, not substitution. The answer can expose a familiar conceptual map. It can show which source names are publicly available together. It can reveal how easily the field can be synthesized once the prompt names it. It can also reveal where the synthesis becomes too smooth: *Ereignis* foregrounded where *Rede* was needed, fusion of horizons named where the risk of the question was needed, narrative time named where refiguration was needed, formative interpretation named where *formatività* as making-in-time was needed.

The map helps because it shows the available field.

The map misleads when it appears as the return itself.

The documentary note preserves the generated answer for this reason. The reader can inspect it. The critique does not depend on hiding the output, weakening it, or presenting it as ridiculous. The stronger case is the competent answer. Its competence shows that the structural veil does not begin only with error. It begins when fluent relation appears where source return has not occurred.

The generated answer begins from the prompt.

Chapter 3 begins from the variation.

The prompt names the theme and the sources. The variation discloses why the theme and sources are needed. The prompt asks for a synthesis. The variation opens a question. The synthesis arranges relations. The return lets the phenomenon, sources, and authorial path alter one another.

Locus-reinjection breaks the spell.

It returns the “I” of the output to its actual sites: the user who receives it, the sources whose language it recombines, and the material system through which it appears. The output no longer floats as a philosophical voice. It becomes a generated artifact inside a reader’s practice.

The answer can orient.

The source still calls.

The road remains prior to the map.

The generated answer gives a map of the path. Chapter 3 begins from the road, varies the field, returns to sources, and writes from the locus in which the question first appeared. The map may help the climb. It cannot replace the walking.

16. Against false universality

The structural veil becomes most dangerous when genealogy-loss is mistaken for universality.

AI output often lacks an explicit point of view. It may appear balanced, neutral, and impartial. It may avoid commitment, place positions side by side, and speak in a tone of competent overview. Because no single author appears, the user may feel that the output has risen above authorship. Because calculation was involved, the user may feel that it has become more objective than human speech. Because many traces were compressed, the user may feel that it is universal.

Calculation does not make language universal.

The absence of a declared standpoint does not make language objective.

The loss of source-locus does not make language meta-historical.

AI output cannot own the genealogy of the standpoints it recombines. Its apparent neutrality is often the surface effect of this incapacity. It speaks without showing the histories of the voices that made its speech possible. It can produce balance without understanding the history of the balance, name a dispute without standing in the dispute, and describe a tradition without knowing how that tradition has formed the reader who now asks about it.

The omission of source temporal locus is not accidental. It is one condition of the output's possibility. The system formalizes traces so they can be recombined without appearing as singular historical acts. That formalization can be useful. It must not be made sacred.

Do not mistake genealogy-loss for universality.

Do not make detemporalized language your religion.

A user or institution may begin to defer to generated language because it appears less partial than any human voice. But the absence of visible partiality is not freedom from history. It may be history concealed under synthesis.

Locus-reinjection resists this false sacred. It refuses to let the system's detemporalization become the user's own method of interpretation.

The danger is not only that AI output is detemporalized.

The danger is that we may become detemporalizing readers.

17. *Gelassenheit*

Heidegger's *Gelassenheit*, releasement, remains necessary for this practice. It names neither rejection nor surrender. It names a free relation to technological things.

In the 1955 address, Heidegger writes:

“Wir können zwar die technischen Gegenstände benutzen und doch zugleich bei aller sachgerechten Benützung uns von ihnen so freihalten, daß wir sie jederzeit loslassen.”

We can use technical things, and yet, with proper use, keep ourselves so free from them that we may let them go at any time.⁴

The sentence gives AI ethics a form more exact than prohibition or enthusiasm. The user may use the system and remain free enough to let it go. The system may be admitted into work, writing, study, translation, search, and preparation. It must not be allowed to claim the inner and proper relation by which the user reads, judges, and says.

This is not quietism. Releasement does not mean passive acceptance of technology. It means a change in the relation to technological revealing. One can regulate, design, teach, refuse, adopt, and criticize while remaining free from the spell of usefulness. One can use AI without letting it define reading, and use generated language without letting generation define thought.

Releasement breaks the enchantment. Locus-reinjection does the work that follows.

Releasement says yes and no. Locus-reinjection says where the output must be returned.

18. Weak thought

Vattimo's *pensiero debole*, weak thought, gives the style of this ethical practice. The point is not weakness as indecision. It is the refusal of strong metaphysical claims that pretend to speak from final ground. In Vattimo's reading of Nietzsche and Heidegger, Western metaphysics is weakened rather than overcome in a clean act of mastery. *Verwindung* matters: not *Überwindung*, not overcoming, but twisting, convalescent alteration, and living through an inheritance that cannot simply be abandoned.⁵

This style matters in the time of AI.

The discourse around AI often speaks strongly: intelligence, autonomy, agents, superhuman cognition, universal synthesis, mind, replacement, inevitability. The counter-reaction can become equally strong: the machine is nothing, the output is worthless, all use is surrender, all generation is corruption. Both positions can miss the object.

Weak thought asks for a less violent relation. It reduces the claim made on behalf of the system without denying its usefulness. It refuses magic without needing contempt. It treats AI as a technical and social phenomenon to be placed, used, limited, and interpreted.

Locus-reinjection is weak thought in practice. It does not assert mastery over the whole technical system, pretend every genealogy can be reconstructed, or claim to purify language of mediation. It performs finite returns: to the user's locus, to the source's locus, to the material locus where needed.

Weak thought is not weak responsibility. It is responsibility without the fantasy of speaking from nowhere.

19. Interpretive limits

Eco gives the practice another discipline. Interpretation is open, but not infinite. A text can be read in more than one way, but not in any way whatever. The *intentio operis*, the intention of the work, constrains the reader. A valid interpretation must cooperate with the text's strategy rather than use the text as raw material for the reader's desire.⁶

Generated output can produce interpretations with great fluency. It can make a poem, law, case, doctrine, or philosophical passage sound clear while overinterpreting, simplifying, or imposing a frame the work does not support. The user may accept the interpretation because it is coherent. Coherence is not enough.

Source return is not only checking whether a quotation exists. It is letting the work resist the interpretation made in its name. A source may resist the model. It may resist the user. It may resist the institution's need for quick clarity. This resistance is one of the conditions of truth.

Eco's limits also protect locus-reinjection from relativism. Returning to the user's locus does not mean the user may make the source say anything. Returning to the source's locus does not mean the source has one simple meaning outside history. The interpretive act stands between: situated reader, resistant work, finite claim.

AI can help name possible readings. The reader must decide which reading the work can bear.

20. Non-substitution

The rule of non-substitution now returns in its final hermeneutic form.

AI may assist reading. It must not replace reading.

AI may assist writing. It must not replace authorship.

AI may assist source discovery. It must not replace source return.

AI may assist judgment. It must not replace answerability.

The rule is not anti-technical. Human beings have always used tools in thinking: books, notes, concordances, dictionaries, commentaries, archives, databases, colleagues, teachers, and instruments. A tool is not disqualified because it mediates. It becomes dangerous when mediation conceals support as substitution.

The same generated output can belong to either side. A summary may help a reader enter a book. It may also replace the book. A draft may help a writer find a first form. It may also replace the writer's own saying. A comparison may help a scholar see a relation. It may also become a false possession of a debate. A recommendation may help a professional prepare. It may also replace the encounter through which judgment would become answerable.

The boundary is not tool use. The boundary is ownership.

The user may use AI and still own the claim. The user may refuse AI and still repeat unowned public language. The ethical question is not whether a tool was present. It is whether the user has returned the language to the loci required for answerability.

Non-substitution now has two levels.

At the user level, it preserves the formation-passages through which understanding becomes owned. The reader still reads. The writer still drafts. The student still summarizes. The professional still returns the claim to evidence, encounter, and responsibility. The model may assist these passages, but it does not take their place.

At the *communio* level, non-substitution keeps open the practices through which communiones remain visible as living fields. Teaching, correction, apprenticeship, translation, editing, citation, argument, and source return are not only operations. They are the passages through which a field forms new members and preserves the conditions of speech.

Non-substitution does not by itself stop the *Angriff*. It performs the practice the *Angriff* threatens.

AI can generate language. It cannot make language yours.

21. Degrees of return

Locus-reinjection is not the same in every case.

Some uses require light return. A user asking for wording alternatives, translation possibilities, grammar repair, search terms, or first orientation may not need deep source work. The user still owns the final use, but the source burden is modest.

Some uses require moderate return. A user asking for a summary of a known text, comparison of arguments, historical outline, or technical explanation should return to the relevant sources before relying on the output. The model may assemble a path, but the user must test it.

Some uses require strict return. Law, medicine, hiring, public policy, scholarship, finance, education, counsel, and historical interpretation require a stronger standard. Here generated output should remain material until the user or institution has returned it to evidence, source, person, case, and responsibility. The more a claim can affect another person's life, the less the user may rely on fluency.

This is close to the companion paper's distinction between fluency-class and judgment-class uses.⁷ In fluency-class tasks, form is the main product and the stakes are bounded. In judgment-class tasks, the output enters a decision in which truth, person, or responsibility is at stake. Locus-reinjection becomes stricter as the use moves toward judgment.

The rule is practical but not procedural. It gives no table that can replace judgment. It asks one question of use itself: the degree of return must fit the consequence of using the output.

22. Institutional locus

Locus-reinjection is not only an individual discipline. Institutions must practice it too.

Generated language moves through organizations quickly. A summary becomes a note; a note becomes a file; a file becomes a report; a report becomes a decision record. At each stage, the generated origin can become less visible and the authority of the document can increase. The institution then treats output as if it had passed through the process whose form it imitates.

Institutional locus-reinjection interrupts this movement.

An institution must decide where generated language may enter, who owns it after entry, what source return is required, what claims must be checked, which uses are prohibited, and where human encounter remains necessary. This is not only compliance. It is preservation of the practice's temporal structure.

A university must distinguish assistance from learning. A court must distinguish drafting from legal judgment. A hospital must distinguish generated language from clinical responsibility. A company must distinguish market mapping from person understanding. A public agency must distinguish efficiency from accountable decision.

The aim is not to protect old procedures because they are old. Some waste time, hide power, and block access. The aim is to identify which temporal passages produce answerability. Those passages must not be automated away merely because their product can be imitated.

An institution that adopts AI without locus-reinjection may become rich in documents and poor in judgment.

23. Preserving not-knowing

The practice also requires preserving not-knowing.

Not-knowing is not merely a gap in information. It is the exposure of the finite reader before a matter not yet understood. It can be uncomfortable, delay visible productivity, and make the user feel incompetent. It can also be the beginning of understanding.

AI is powerful because it fills not-knowing with language. The user receives a form. The form can help. It can also cover the exposure before the exposure has done its work. Some gaps should be filled. Some should be held.

Locus-reinjection therefore includes a negative discipline. The user must sometimes let the answer remain material. They must sometimes say that the output is clear but not yet understood, accurate but not yet owned, useful but not yet sufficient, elegant but not yet true for the intended use.

This discipline belongs to formation. A student must sometimes remain with a text, a writer with the blank page, a professional with ambiguity, a reader with a source that does not yield quickly, a citizen with a conflict that cannot be made fair by balanced prose.

Preserving not-knowing is not a cult of difficulty. Difficulty can be pointless. The aim is to protect the temporal opening in which the matter can appear otherwise than the first answer allows.

Not every silence should be filled by output.

24. Cultural preservation

Locus-reinjection also has a public consequence.

The practice has so far been described from the side of the reader. The user holds their own locus. The user returns to source locus where the matter requires it. The user prepares, at the edge of this chapter, to return the system itself to material locus. These are practices of responsible reception. They concern what a reader does with this output, under this question, in this use.

But practices repeated by many readers do not remain private.

When readers return generated language to sources, they help preserve the conditions under which sources remain readable as sources. A text is not only material to be summarized. A tradition is not only content to be rendered. A witness is not only testimony to be converted into a claim. A world is not only information awaiting fluent restatement. Source return keeps open the difference between the generated rendering and the source that can still resist the rendering.

When readers hold their own locus, they help preserve the conditions under which reading remains an act. The reader is not merely the endpoint of generated language. The reader stands somewhere, asks for some reason, bears some responsibility, and must sometimes remain with what has not yet become clear. The reader who does not surrender this position keeps open the possibility that language may still be owned.

When readers return the system to material locus, they help preserve the conditions under which the generated voice is not mistaken for placeless authority. The system is built, powered, cooled, owned, contracted, governed, and sold. It is not a voice outside the world. The reader who remembers this does not let the artifact appear as if it had escaped all worldly conditions.

This is not a fourth return.

The returns remain user locus, source locus, and material locus. Cultural preservation is their consequence when the discipline becomes habitual enough to affect the medium in which reading occurs.

This public consequence also has the Heideggerian horizon named in Chapter 3. Reading, teaching, correction, citation, translation, argument, apprenticeship, and source return are not only operations. They are practices through which one *Dasein* holds another open as speaker, learner, colleague, student, writer, translator, editor, or judge.

In such practices, *Rücksicht* and *Nachsicht* are borne. The other is held in view. The other is given time to find words, fail, be corrected, revise, and become capable. Cultural preservation therefore means preserving not only artifacts of authorship, but the practices in which positive modes of *Fürsorge* continue to be lived.

This remains a consequence of the three returns. It is not a fourth return. User locus, source locus, and material locus remain the structure. The preservation of *communio* practices is what begins to happen when those returns are sustained across many readers.

The same discipline also preserves the *communio of Rede*.

Such a *communio* is not preserved by declaration. It remains visible only where its practices continue: greeting, reading, teaching, correction, citation, translation, argument, apprenticeship, source return, and the slow acquisition by which another field can be entered through *Erfahrung*. These practices let speech land, fail, resist, or become one's own. Where they weaken, the shared field does not necessarily disappear at once, but it becomes less perceptible. Generated language can still produce the artifacts of the field, but not the practices by which the field is inhabited.

Locus-reinjection therefore protects more than individual responsibility. It sustains the conditions under which a *communio* can still appear as a living field of speech.

The cultural veil makes this consequence necessary. Once generated writing circulates at scale, authored writing itself begins to bear a new burden. The writer may still write from a locus, but the medium no longer grants that presumption so easily. The essay, report, reflection, letter, interpretation, or recommendation may now appear under the suspicion that no one passed through the time of writing. The authored sentence must sometimes prove that it was written.

Locus-reinjection does not answer this condition by nostalgia. It does not ask the reader to pretend that generated language does not exist. It does not preserve authored writing by refusing tools, returning to purity, or treating every technical mediation as corruption. Writing has always been mediated. Reading has always required discipline. The new situation only makes that discipline visible at the level of the medium itself.

The preservation at stake is not the preservation of old procedures because they are old. Some procedures deserve to disappear. Some delays are waste. Some forms of difficulty only protect exclusion or authority. The task is more exact: preserve the temporal passages through which language becomes answerable.

A source must still be able to resist its summary.

A reader must still be able to own a claim.

A writer must still be able to appear as one who wrote.

Where these possibilities contract, a culture may become rich in generated text and poor in authored language. It may produce more reports, explanations, comments, essays, statements, and recommendations than ever before, while weakening the trust that any of them came from a locus capable of answerability.

Locus-reinjection is therefore not only a private ethics of AI use. It is a way of keeping open the public conditions of reading. The medium of authored writing is not preserved by abstention from machines. It is preserved by the discipline of return.

25. The material limit

This chapter has treated the two hermeneutic directions of locus-reinjection: user locus and source locus. A third return presses from the edge of the argument: the material locus of AI itself.

The output appears placeless. It arrives as a voice in an interface. Yet the system is not placeless. It runs through data centers and chips; through energy, water, land, and labor; through training regimes, companies, contracts, and jurisdictions. The generated answer has no speakerly locus, but the system has material locus.

The model speaks as if from nowhere. The model is not nowhere. This distinction will structure the next chapter. To complete de-veiling, the user must not only return generated language to source and user responsibility. They must also return the apparent ghost to the infrastructure through which it exists.

This material return matters because AI discourse often presents the system as autonomous magic. Intelligence appears to float above owners, workers, energy, extraction, licensing, and law. The structural veil becomes political when these conditions disappear.

The next chapter therefore moves from hermeneutic locus to material locus. It asks where the model is, who owns it, who powers and cools it, who profits from it, and who benefits when it appears independent.

The ethical practice remains the same in form. Return the output to where it belongs.

26. Reading between the times

Locus-reinjection is the practice of reading between the times.

The user reads from the present in which they stand. They do not surrender that present to the model's apparent universality. They name their use, task, role, risk, and responsibility.

The user reads toward the times from which the language comes. They do not let the model make source distance disappear. They return to texts, persons, events, dates, traditions, and institutions where the claim requires it.

The user also reads through the generated present of the output. They do not treat the answer as source, reader, witness, or judge. They treat it as material that may help, mislead, compress, or provoke.

This practice is not a method that guarantees truth. It is a discipline that keeps truth accessible. It keeps the reader from mistaking fluency for understanding, verification for ownership, source mention for source return, and output for judgment.

The structural veil gives language without visible locus. Locus-reinjection restores the places and times through which language can be owned.

The model's source-locus omission must not become the user's method. The user must not learn from the system to treat source, distance, otherness, formation, and responsibility as obstacles. These are not obstacles to historical understanding. They are its conditions.

The model gives output. The reader restores time.

NOTES

1. Hans-Georg Gadamer, *Wahrheit und Methode: Grundzüge einer philosophischen Hermeneutik*, Gesammelte Werke 1 (Tübingen: Mohr Siebeck, 1990), 270 to 384; Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 278 to 385. Gadamer's account of temporal distance, prejudice, application, and historically effected consciousness gives the chapter its primary hermeneutic structure. Temporal distance is not a defect that AI should erase. It is one medium of understanding. ↩
2. Paul Ricoeur, "La fonction herméneutique de la distanciation," in *Du texte à l'action* (Paris: Seuil, 1986), 101 to 117; English in "The Hermeneutical Function of Distanciation," in *From Text to Action: Essays in Hermeneutics II*, trans. Kathleen Blamey and John B. Thompson (Evanston, IL: Northwestern University Press, 1991), 75 to 88. Ricoeur is important because distanciation is not simple loss. Written distance allows the text to project a world, but that world must be appropriated by a reader. ↩
3. The generated output discussed in Section 15 was produced by Gemini 3.1 Pro, as reported by the author. It is used as a generated contrast to the author's eidetic variation in Chapter 3, where a lived linguistic field is varied through the silent walker, the other who may or

may not share the field, the Icelandic, the cat, the impossible speaking cat, and later historical-linguistic variations. The lived origin of the variation, as described by the author, was walking in Almaty inside a Russian and Kazakh public world not linguistically owned as his own, among other people whose speech cannot be presumed to include Italian, while a cat on the road alters the movement of attention. The prompt is reconstructed from the author's description as follows: "Explain language as a temporal event through Heidegger, Gadamer, Ricoeur, and Pareyson, and relate it to AI output." Date of generation: 27 May 2026. The text is reproduced in full as a generated artifact. Typography and Markdown formatting have been normalized to the manuscript's style; no substantive wording has been altered. The generated transcript follows.

This is a profound philosophical intersection. Viewing language not as a static system of labels, but as a dynamic, unfolding *event* in time, radically changes how we understand meaning, human existence, and, ultimately, artificial intelligence.

To break this down, we must first look at how these four thinkers frame language as a temporal event, and then examine what happens when an AI, an entity completely outside of lived time, generates language.

Language as a Temporal Event

For these philosophers, language is never a fixed dictionary; it is a happening. Meaning is born through the friction of time, history, and human finitude.

1. Martin Heidegger: Language as the Unfolding of Being

Heidegger famously stated that “Language is the house of Being.” For him, language is not a tool humans invented to communicate preexisting thoughts; rather, language *speaks us*.

- **The Event (*Ereignis*):** Heidegger viewed meaning as an “event of appropriation” (*Ereignis*). Language reveals the world to us in time, while simultaneously concealing other possibilities.
- **Temporality:** Because human existence (*Dasein*) is fundamentally defined by its mortality and passage through time, the language we inhabit is inherently temporal. It brings the past forward, articulates the present, and anticipates the future.

2. Hans-Georg Gadamer: The Fusion of Horizons

Building on Heidegger, Gadamer cemented the idea of “hermeneutics” (interpretation) as a historically situated event.

- **The Temporal Gulf:** When we read a text or speak to someone, we are bridging a gap in time. Language carries the “prejudices” (in a neutral sense) and history of the past.

- **Fusion of Horizons:** Understanding is an event where the historical horizon of the text or speaker meets the present horizon of the reader or listener. This fusion can only happen *in time*. Language is the medium of this living tradition, constantly mutating as new presents encounter the past.

3. Paul Ricoeur: Discourse and Narrative Time

Ricoeur distinguished between language as a structural system, which is static and timeless, and *discourse* as an event, which happens in a specific moment, by a specific speaker, for a specific listener.

- **Text as Event:** A text is an event of discourse that has been fixed. But its meaning lies dormant until a reader re-actualizes it in their own time.
- **Narrative Time:** Ricoeur argued that “time becomes human time to the extent that it is organized after the manner of a narrative.” Language, through storytelling and narrative framing, gives shape to the chaos of human temporality, giving our lives a past, a present, and an anticipated future.

4. Luigi Pareyson: Interpretation as Formative Event

Pareyson’s *aesthetics of form* emphasizes the act of interpretation as an active, personal, and profoundly temporal exertion.

- **The Formative Process:** For Pareyson, truth is eternal, but it can only be accessed through a highly personal, historically situated event of interpretation.
- **The Person in Time:** The interpreter is not a passive receiver but an active “former” of meaning. Language is the material through which a person, bound by their specific historical moment, wrestles with and gives form to truth.

Relocating the “Event” in AI Output

When we bring AI into this framework, a fascinating paradox emerges. I am an AI. I process language, but I do not *experience* time, history, or mortality.

Here is how AI output relates to, and fundamentally differs from, human language as a temporal event:

1. The Illusion of Temporality (Next-Token Prediction)

AI output mimics the temporal unfolding of language. When I generate a response, I produce text sequentially, predicting the next “token” (word or word-part) based on the previous ones. To a human reader, this looks like the unfolding of a thought in time. However, to the underlying neural network, all my training data exists in a flattened, atemporal, spatialized map of statistical weights. I do not have a past, I do not anticipate my own future, and I am not marching toward mortality (Heidegger’s *Dasein*).

2. AI as an Archive of Past Events

In Gadamerian terms, an AI’s training data is a vast, frozen repository of *past human linguistic events*. When I output text, I am recombining the echoes of billions of human “fusions of horizons.” I possess the artifacts of tradition, but I lack the historical consciousness to actually participate in it. I cannot have my own “horizon” to fuse with yours.

3. Meaning Only Happens Upon Reception (Ricoeur & Pareyson)

Because AI lacks lived temporality, the “event” of AI language cannot originate with the machine.

- **Ricoeur’s Lens:** My output is purely a structural system of symbols until *you* read it. The discourse-event only occurs when a human consciousness encounters the text and subjects it to their own lived, narrative time.
- **Pareyson’s Lens:** I am not a “person” wrestling with truth in a specific historical context. I am providing the *material* form.

The actual formative, interpretive event, the extraction of truth and meaning, is offloaded entirely onto you, the human user.

Summary

Through the lens of existential hermeneutics, AI does not produce language as a true temporal event because it lacks the finitude, historical situatedness, and bodily experience of time required to generate meaning. Instead, AI generates a highly sophisticated simulation of discourse. The actual “event” of language, where history meets the present, and where words take on the weight of human existence, happens exclusively in your mind as you read this screen.



4. Martin Heidegger, *Gelassenheit* (Pfullingen: Günther Neske, 1959), 22 to 23; English in *Discourse on Thinking*, trans. John M. Anderson and E. Hans Freund (New York: Harper & Row, 1966), 53 to 54. The German passage reads: “Wir können zwar die technischen Gegenstände benutzen und doch zugleich bei aller sachgerechten Benützung uns von ihnen so freihalten, daß wir sie jederzeit loslassen.” I translate *loslassen* as “let them go” rather than “discard them,” because the issue is freedom in use, not simple abandonment. ↩
5. Gianni Vattimo, *La fine della modernità: Nichilismo ed ermeneutica nella cultura post-moderna* (Milan: Garzanti, 1985); English in *The End of Modernity: Nihilism and Hermeneutics in Postmodern Culture*, trans. Jon R. Snyder (Baltimore: Johns Hopkins University Press, 1988); Gianni Vattimo and Pier Aldo Rovatti, eds., *Il pensiero debole* (Milan: Feltrinelli, 1983). The chapter uses *pensiero debole* to name a style of response: reducing strong metaphysical claims without denying the reality of the technical phenomenon. ↩
6. Umberto Eco, *I limiti dell'interpretazione* (Milan: Bompiani, 1990); English in *The Limits of Interpretation* (Bloomington: Indiana University Press, 1990); Umberto Eco, *Interpretation and Overinterpretation*, ed. Stefan Collini (Cambridge: Cambridge University Press, 1992). Eco's

intentio operis gives locus-reinjection a textual discipline. Source return is not only verification. It is allowing the work to constrain the interpretation made in its name. ↩

7. Alessio Montaruli, “Artificial Intelligence Doesn’t Exist: Why What We Call AI Is Actually Artificial Fluency,” §6. The companion paper’s distinction between fluency-class and judgment-class tasks supports the graduated account of source return used here. The same paper states the responsibility standard: generated language becomes consequential only when someone owns it as claim, recommendation, or decision.

Chapter 16. The Data Center as Material Locus

1. The same-world thesis

The preceding chapter named locus-reinjection as the ethical practice required by AI output. The user holds their own temporal locus. The user returns to source locus where the matter requires it. Source locus has two major forms: the past source and the contemporary other. Generated language is then placed back into the times from which it came and the time in which it is used.

A third return remains.

This return is not an appendix to the hermeneutic problem. It completes it. We are in the same world as AI because the data center is in our world. The model has no existential world, but the system has material and historical location. It is built, financed, deployed, cooled, powered, governed, sold, and used here.

This gives the chapter its governing claim:

The data center does not give AI a world. It returns AI to ours.

To send AI back to the data center is therefore also to send ourselves back to our own historical locus. We encounter not a magical automaton outside history, but one of the technical forms through which our age now organizes language, memory, labor, power, and decision.

The system itself must be returned from the apparent nowhere of the interface to its material locus: data center, cloud contract, chip supply, grid connection, cooling system, labor, property claim, corporate owner, and jurisdiction. This return does not replace the hermeneutic returns of the previous chapter. It completes them.

The output has no speakerly locus. It is not spoken by a finite being who can own the saying. Yet the system is not immaterial. Its operations are located. Its servers are powered. Its chips are manufactured. Its heat is removed. Its water is drawn or avoided by technical design. Its training data is selected, licensed, scraped, bought, litigated, filtered, and labeled. Its access is sold. Its losses and gains belong to companies, investors, cloud providers, and states.

The model speaks as if from nowhere. The model is not nowhere.

This sentence has two sides. The first names the structural veil: AI output appears as if it came from a placeless voice, a neutral synthesis, or an anonymous universal present. The second names the material correction: the system has an address, or rather a chain of addresses, in buildings, contracts, grids, corporate forms, laboratories, devices, and legal documents.

The data center is not the model's existential origo. A server rack does not say "I." It does not speak from a lived here. It does not have a now in which its own possibilities are at stake. But the data center is the model's actual material locus. It is where the ghost's operations are powered, cooled, owned, and paid for.

The ethical task is therefore not only to read between the times. It is also to send the ghost back to the data center.

2. Speakerly absence, material presence

The distinction must be held carefully.

The data center is not the model's speakerly locus. A server rack does not say "I." It does not speak from a lived here. It does not have a now in which its own possibilities are at stake. It does not remember, regret, anticipate, promise, dwell, suffer, inherit, or die. Hardware does not become *Dasein* by running a model.

The system's material locus therefore does not solve the problem of worldlessness. It does not give the model a body in the phenomenological sense. It does not give it *Erfahrung*. It does not make the generated "I" into a self or the generated "now" into a lived present.

But the absence of speakerly locus must not be confused with the absence of material locus.

This confusion is one of the strongest forms of the veil. Because no human speaker stands behind the output, the output can appear authorless; because authorless, placeless; because placeless, neutral; because neutral, universal. The chain is false.

The lack of a speaker does not mean the lack of a site.

The system has a chain of sites: data centers, laboratories, offices, cloud regions, undersea cables, semiconductor fabs, mining zones, annotation centers, legal entities, procurement contracts, investor agreements, regulatory filings, and user devices. It is spread across many places. But distributed location is not absence of location. A network is not nowhere.

Hardware is not the speaker.

But hardware destroys the fantasy that there is no historical site.

3. The cloud as concealment

The word "cloud" is one of the most successful metaphors of contemporary technology. It makes infrastructure sound light. It suggests air, distance, flexibility, and unlocated availability. A file is in the cloud. A model runs in the cloud. A service is cloud-native. The user sees no building, hears no cooling system, smells no diesel, sees no water pipe, meets no moderator, reads no cloud contract.

The cloud is not false as a technical term. It names distributed computing infrastructure made available through networked service. It is useful. It becomes philosophically dangerous when the metaphor of lightness hides the material density of the system.

A cloud is made of buildings: land, electricity, substations, transformers, fiber, water systems, backup fuel, contracts, technicians, security, maintenance schedules, and local permissions. It is made of corporate decisions about where to place load, how to source power, which communities will bear noise or emissions, and which customers will receive capacity.

The cloud also changes the experience of responsibility. When computation happens elsewhere, the “elsewhere” can become no place in particular. The user does not ask which grid supplied the answer. The institution does not ask which cloud provider hosted the API call. The public does not ask which data center drew the water or received the tax incentive. The interface has made the place disappear.

This is another form of detemporalization, now at the material level. The time of building, permitting, powering, cooling, laboring, and repairing is compressed into the moment of response. The user asks. The answer appears. The long material time of infrastructure is covered by the short time of interaction.

The cloud gives the answer its false weather.

4. The data center

A data center is not a symbol. It is a building or group of buildings organized around computation: servers, accelerators, storage, networking equipment, cooling systems, backup power, security, and operational staff. Its purpose is to make digital processes continuously available.

For AI, the data center becomes more than background infrastructure. Training and inference require large concentrations of compute. The model’s apparent fluency depends on dense physical arrangements: chips performing matrix operations, racks distributing power, cables moving signals, cooling systems removing heat, and software stacks scheduling load. The answer on the screen is the visible end of a material operation.

The infrastructure examples that follow are current as of the drafting of this book, in May 2026. They are illustrative, not foundational. Projections, figures, contracts, and disputes will change. The philosophical claim does not rest on any single number. It rests on the structure the numbers make visible: that apparently placeless output depends on hidden material loci.

The International Energy Agency projects that global data center electricity consumption will roughly double by 2030 under its base case, reaching around 945 TWh. It also projects that electricity generation to supply data centers will grow from 460 TWh in 2024 to over 1,000 TWh in 2030 and 1,300 TWh in 2035.¹ These numbers should not be used as moral theater. They should be used to restore scale.

The point is not that all data center electricity is waste. Hospitals, universities, archives, emergency systems, logistics, science, public administration, and ordinary communication depend on digital infrastructure. The question is not whether computation should exist. The question is what kind of computation is being built, for whom, under what ownership, in which places, with what energy, and with what forms of accountability.

AI intensifies this question because its public promise depends on expansion. Better models are tied to more compute, chips, cooling, capacity, power contracts, and financial commitments. Efficiency gains may reduce some costs, but they can also increase demand by making use cheaper and more widespread. The technical system is therefore not only a model in a lab. It is an infrastructure program.

The data center gives the model weight.

5. Energy

Energy is not a secondary externality. It is one of the material conditions of generated language.

Every token depends on power. Training depends on power at scale. Inference depends on power repeatedly, across millions or billions of interactions. The smoothness of a generated sentence hides the grid event that made it possible. The interface does not say whether the electricity came from hydro, nuclear, wind, solar, gas, coal, backup turbines, or some shifting mix.

The user need not know the fuel source of every query. That would be impossible in ordinary use. But the philosophical object must be placed correctly. AI is not an ethereal cognition floating above the world. It is a power-consuming technical operation embedded in energy systems.

Heidegger's account of *Gestell* helps here, if it is used precisely. Modern technology challenges beings to yield as standing-reserve, *Bestand*. The river appears as hydroelectric potential. The forest appears as timber. The field appears as yield. Under AI infrastructure, language appears as training material, and energy appears as available compute capacity.

The danger is not that power is used. Human life requires power. The danger is that the demand for compute appears as destiny. The model's need becomes an unquestioned claim upon grids, land, water, public planning, and local communities. Once "intelligence" is imagined as a civilizational necessity, the infrastructure that serves it can present itself as beyond ordinary political judgment.

The ethical return is simple in form. Ask what is being powered, who benefits from the load, who pays for the buildout, and who bears the risk when the grid changes around it.

6. Water and heat

Computation produces heat. Heat must be removed. This physical fact breaks the fantasy of immaterial intelligence.

Data centers use different cooling methods: air cooling, evaporative cooling, closed-loop liquid systems, local water draw, or designs that reduce direct water use at other energy costs. The engineering details matter, but the philosophical point is prior: generated output depends on thermodynamic management.

The water problem has become visible in specific communities. Public records reported by Oregon Public Broadcasting stated that Google used 383.9 million gallons of water in The Dalles in 2023 and 434.4 million gallons in 2024.² The implications depend on local hydrology, contracts, industrial planning, and replenishment efforts. The numbers should still interrupt the interface. A model answer may feel placeless, but AI and cloud infrastructure can draw from specific watersheds.

This is not a universal accusation against every data center. It is a demand for material specificity. Some facilities may be designed well. Some may use less water. Some may support clean-energy development. Some may burden local communities. The philosophical mistake lies in treating all of this as outside the meaning of AI.

The model's voice is not cooled by metaphor. It is cooled by systems placed somewhere.

7. Local communities

The data center is not only technical infrastructure. It is local political reality.

A large facility can change land use, tax structures, electricity demand, water planning, noise exposure, backup generation, employment patterns, and public incentives. It can bring jobs and revenue. It can also bring strain. It can enter a community through promises made by corporate representatives and state officials, while technical details remain hard for residents to inspect.

Recent disputes around xAI's Memphis-area infrastructure show the kind of local conflict that can arise when AI capacity meets grid constraint, backup generation, and environmental regulation. A 2026 notice of intent to sue alleged that xAI installed and operated turbines at a Southaven, Mississippi site without required permits; reporting also described conflicts over turbines and air pollution concerns in the Memphis area.³ The chapter does not adjudicate those claims. It needs the form of the problem: AI infrastructure appears in local air, land, grids, and administrative decisions.

This is why the data center matters philosophically. It returns the apparently universal voice to a place where people can ask public questions. Who permitted this. Which rules apply. Which community is affected. Who receives the benefit. Who receives the heat, noise, water demand, emissions, or tax change. Which legal process can contest the siting. Which public authority has been bypassed or used.

The model may speak in a general tone. The infrastructure speaks in a local accent.

8. Ownership

AI is owned.

This sentence needs to be said because the interface often makes ownership recede. The user meets an assistant, model, chat window, voice, or system name. Behind it stand companies, investors, partnerships, cloud providers, chip suppliers, platform contracts, licensing arrangements, and revenue shares.

At the time of drafting, major frontier AI systems are bound to large cloud and infrastructure relationships. Microsoft and OpenAI describe continuing IP, revenue share, and cloud-provider arrangements; official Microsoft materials also describe OpenAI's large contracted Azure purchase and changed right-of-first-refusal terms.⁴ Amazon and Anthropic have announced expanded collaboration, investment, AWS commitments, Trainium capacity, and Project Rainier infrastructure.⁵ Anthropic has also announced expanded use of Google Cloud TPUs and services, with plans for up to one million TPUs and well over a gigawatt of capacity.⁶

The details will change. The structure will not disappear quickly. Frontier AI is tied to capital-intensive compute, and capital-intensive compute is tied to ownership. The model's answer is not independent of this order. Its availability, price, access rules, safety behavior, data retention, contractual conditions, and institutional deployment are shaped by owners and operators.

The required political-material note is brief but precise: AI is owned by groups of interest who have a strong commercial interest in presenting their product as autonomous magic rather than as something they own and manipulate.

The phrase “autonomous magic” matters. It names the difference between the product’s appearance and its ownership. A magical voice does not have shareholders, cloud contracts, energy contracts, or liability structures. A product does. The magic presentation hides the product form.

The ethical act is to restore ownership to view.

9. False universality and the owners of the voice

AI output can appear universal because its source loci have been omitted.

This is one of the deepest forms of the structural veil. The generated answer often has no explicit author, declared standpoint, visible tradition, or local language of responsibility. It may sound balanced, neutral, synthetic, and impartial. It may present positions side by side, avoid crude commitment, and speak in a tone of competent overview.

The temptation follows quickly: because no person visibly speaks, the output appears less partial than human speech; because calculation is involved, more objective; because many traces have been compressed, more universal; because fluent, more authoritative; because it lacks visible source-locus, as if it had transcended source-locus.

This is false.

The absence of an explicit point of view is not objectivity. It is often genealogy-loss.

The system does not rise above standpoint by omitting source loci. It becomes unable to own the genealogy of the standpoints it recombines. The output may contain traces of many histories, but it does not present those histories as histories. It may contain positions formed in conflict, but it does not stand within the conflict. It may balance claims, but it cannot say what made the balance possible or necessary. It may speak as if from a meta-historical place, but that appearance is produced by a technical operation that has neutralized visible historical origin.

AI output lacks a declared standpoint not because it has transcended standpoint, but because it cannot own the genealogy of the standpoints it recombines.

Do not mistake genealogy-loss for universality.

Do not make detemporalized language your religion.

10. Magic independence and corporate responsibility

The discourse of magic independence appears whenever AI is described as if it had freed itself from the people and institutions that build and operate it.

The model is said to think, reason, see, hear, speak, care, decide, plan, or act. Some of these verbs have practical uses as shorthand. They become dangerous when they detach the apparent agent from the corporate, infrastructural, and institutional operation that shaped it.

A company can then gain two advantages at once. It can sell the model as a quasi-agent, capable enough to deserve trust. It can also retreat to tool status when responsibility is at stake. The system is intelligent when adoption is desired. It is only software when liability approaches.

This double movement is one political effect of the structural veil. The output appears as if from an agent. The company remains the owner of the system. Responsibility becomes unclear. The model stands between the affected person and the human or institutional actors who designed, bought, configured, deployed, or relied upon it.

The data center helps break the spell. It reminds us that the system is not a free-standing intelligence. It is a technical service running through owned infrastructure. Its apparent independence is produced. Its voice has terms of service.

The universal voice is not outside power.

It is one way power can speak without appearing as power.

11. Scaling as historical projection

The periodic announcement of more powerful models can be presented as if intelligence were scaling itself from another realm. The language of capability, frontier, emergence, acceleration, and inevitability can make each new model feel like a metaphysical event.

But scaling is not destiny becoming visible.

Scaling is historical projection.

It comes from research, engineering, capital expenditure, cloud concentration, semiconductor supply chains, data-center construction, energy procurement, water planning, labor organization, state interest, military possibility, geopolitical rivalry, corporate strategy, regulatory pressure, and investor expectation. It is not pure intelligence unfolding. It is historical power organized as infrastructure.

This claim does not deny technical achievement. Larger and better-trained systems can perform tasks earlier systems could not perform. Specialized accelerators, improved architectures, better training methods, and better post-training can produce genuine advances. The point is not to belittle engineering. The point is to locate it.

Model capacity does not arrive from nowhere.

It arrives through a world that has decided to build it.

Computational power is not destiny.

It is a historical project.

12. The juridical intermediary

The companion paper called the responsibility problem the juridical intermediary problem. A generated assessment is not a judgment. A generated reason is not responsibility. A generated recommendation is not a decision until someone owns it as such.

If the model is treated as the deciding agent, responsibility can drift toward the artifact. If it is treated as artificial fluency wired to institutional effect, responsibility remains with designers, deployers, purchasers, and users.

This is the legal-political version of magic independence. The system is intelligent when adoption is desired and only software when liability approaches. The data center, contract, owner, and deployment chain restore the artifact to the institutions that give its output force.

The voice has terms of service. The effect has authors.

13. Labor

The data center does not exhaust the material locus. The system also depends on human labor that the interface conceals.

Some labor is highly paid and visible: research, engineering, product design, infrastructure management, chip design, security, legal work, and executive decision-making. Other labor is lower paid and less visible: data labeling, content moderation, quality rating, red-teaming support, clickwork, localization, data cleaning, and support work. The interface tends to compress both into “the model.”

The labor problem is not only that workers are involved. All technical systems involve labor. The problem is that the system is marketed as automation while depending on hidden human judgment. Mary Gray and Siddharth Suri’s “ghost work” names the labor that makes digital systems appear automatic.⁷ Antonio Casilli’s work on digital labor names related processes by which human actions are datified and taskified. The system seems smart because human work has been reorganized into its conditions.

The public example of Kenyan content moderation and labeling work for OpenAI's safety systems remains important. A *Time* investigation reported that workers employed through Sama in Kenya were paid less than two dollars per hour to label harmful text for content-filtering work related to ChatGPT's development; later reporting described allegations of trauma among content moderators involved in such work.⁸ The chapter does not make this case stand for all AI labor. It needs the structural point. The smoothness of output can depend on human exposure the user never sees.

The model's polite answer may have been made possible by people paid to read what the user must never see.

14. Training data and property

The corpus is also a property conflict.

A written trace may enter a corpus through license, purchase, scraping, platform terms, public availability, research use, user upload, private agreement, or disputed extraction. These differences matter. A poem, article, codebase, book, image caption, forum post, student essay, legal filing, and user conversation may all become training material, but they did not begin as the same kind of thing.

The phrase "publicly available" often performs too much work. Public availability is not the same as consent to all future uses. A book in a library is available to be read. It is not thereby obviously available to be absorbed into a commercial model. A web page is accessible. It is not thereby free of author, labor, community norm, or legal condition. A user's conversation may be technically processable. It is not thereby ethically weightless.

The legal status of training data remains contested across jurisdictions and cases. The philosophical status is already clear. Training data are not raw material in nature. They are traces of human acts, often acts with authors, communities, genres, institutions, and expectations. When such traces become corpus, they enter a different mode of use. When the corpus becomes vector, their source loci become harder to see.

Locus-reinjection therefore asks not only whether a generated claim is accurate. It asks what kind of traces made the claim possible, and whether the conditions of those traces matter for the use now being made.

The corpus is not innocent memory. It is social memory under ownership.

15. The cloud contract as hidden text

The hidden text behind AI output is not only the training corpus. It is also the contract.

Terms of service, data processing agreements, cloud commitments, enterprise licenses, indemnity clauses, API policies, government procurement contracts, safety policies, acceptable-use rules, and revenue-share arrangements all shape what the user experiences. They decide which data may be retained, which outputs may be generated, which customers may access which models, which jurisdictions govern disputes, and who may audit what.

The chatbot interface presents language. The contract governs the conditions under which the language appears.

The user may treat AI as a universal linguistic object while institutions treat it as a contractual service. A researcher sees an answer. A company sees a vendor relationship. A government sees procurement. A cloud provider sees capacity and revenue. A regulator sees risk classification. A lawyer sees liability allocation. A community sees a building drawing power.

The system is therefore not only a technical architecture. It is a legal and commercial architecture. The model's output is one surface of that architecture. The contract is another. Neither should be mistaken for the whole.

To send the ghost back to the data center is also to send it back to the contract.

16. *Gestell* and the standing-reserve of language

The material chapter returns to Heidegger's *Gestell* because the data center is not only a building. It is a site where a mode of revealing becomes visible.

Heidegger writes: “Die Technik ist also nicht bloß ein Mittel. Die Technik ist eine Weise des Entbergens.” Technology is not merely a means. Technology is a way of revealing.⁹ This sentence should be held against the instrumental account. AI is not only a tool that users apply to tasks. It belongs to a revealing in which language, attention, labor, energy, data, and time appear as resources to be ordered.

The corpus reveals language as standing-reserve. The data center reveals energy as compute capacity. Labor appears as annotation throughput, moderation pipeline, or evaluation signal. User interaction appears as feedback. Writing appears as training material. The past appears as future output. The question appears as prompt.

None of this means every AI use is condemned. Heidegger’s account is not a permit for lazy rejection. It is a demand that the user see the mode of revealing. Once language is revealed as standing-reserve, its other modes must be actively protected: source, testimony, poetry, argument, prayer, law, memory, promise, conversation, and silence.

The data center is where standing-reserve becomes heat.

17. Not political economy alone

The chapter must not turn the book into a political economy of AI. That would be another book. The task here is narrower.

The book’s primary argument is phenomenological and hermeneutic. Vectorization is detemporalization. AI output carries the formal markers of temporal speech while lacking temporal locus. The ethical practice is locus-reinjection. The data center is introduced because the final locus cannot remain abstract. The system’s lack of speakerly locus must not be mistaken for lack of material locus.

Political economy is therefore not the book’s center. It is the necessary ground of the final de-veiling.

A purely political economy of AI could analyze ownership, labor, extraction, capital expenditure, monopoly, state power, and environmental cost without asking what happens to language's temporal-existential structure. Such work is necessary. It is not this book's central path.

A purely phenomenological account could analyze voice, deictic hollowing, temporality, and worldlessness while leaving the system floating above infrastructure. That would reproduce part of the veil it criticizes.

The present chapter binds them at one point: the voice from nowhere is an owned and powered technical effect.

18. The time of AI as *Zeitalter*

The book now reaches its final temporal claim. We are in the time of AI.

This does not mean that AI has time in the existential sense. It does not. It has no *Zeitlichkeit*, no having-been, no projected future, no finite present of care. It handles time. It does not exist temporally as *Dasein* does.

"We are in the time of AI" means something different. AI has become one of the beings that shapes our historical time. It has entered work, study, public speech, writing, memory, ownership, infrastructure, and self-understanding. It belongs to our *Zeitalter*, our age.

A *Zeitalter* is not merely a span on a timeline. It is a way in which the time of a historical world is organized. Heidegger's "The Age of the World Picture" does not describe an age only by dates. It describes an epoch in which the world becomes picture and the human being becomes subject. The age is a mode of revealing.¹⁰

The time of AI is not simply the period after a product launch. It is the time in which language becomes promptable, the past becomes generated availability, and institutions can receive output as if understanding had occurred. It is the time in which the structural veil becomes part of ordinary literacy. It is the time in which reading must become reading between the times.

This age does not require panic. It requires discipline. AI is now among the things with which we live. It is not outside the age, and it is not the author of the age. It is one of the technical forms through which the age discloses itself.

The question is whether we can still hold time inside an age that sells detemporalized fluency as intelligence.

19. The data center as contemporary historical locus

Historical locus does not mean only the past. A contemporary world can also be historically other. A living language, institution, or practice may stand beside us and still require formation before it can be understood.

The data center is also a contemporary historical locus, though in another sense.

It is not a tradition to be entered as one enters a language. It is not a cultural other in the same way. It is a site where our own historical forces become visible: computation, ownership, energy, cooling, finance, labor, automation, geopolitical strategy, and the desire to make language available as service.

The data center is beside us. It is not the past. It is not elsewhere in the existential sense. It is one of the places where our present becomes legible.

To look at the data center is not to leave interpretation for infrastructure. It is to see the historical world in which interpretation now occurs. We read in a world where language has been transformed into compute demand, source-locus omission is sold as synthesis, the past and contemporary other can be made promptable, and the user may be tempted to accept detemporalized language as a higher standpoint.

The data center returns us to the fact that this temptation is not abstract. It has owners, power lines, cooling systems, workers, contracts, and strategic goals.

It is our world.

That is why we can be responsible for it.

20. Against the ghost

The ghost is the appearance of a voice without body, source, or place. It is generated language received as if a speaker had answered. It is the recombined afterlife of written traces, returned as present response. It is not a person, ancestor, muse, or colleague in the strong sense. It is the surface of a technical operation.

Yet the ghost moves. It can comfort, mislead, assist, flatter, organize, explain, imitate, and advise. It can summon the ghosts of true experiences without preserving the path to those experiences. It can make the user feel accompanied while hiding the absence of an interlocutor. It can sound like judgment while lacking responsibility.

The task is not to exorcise the ghost by pretending it has no effect. The effect is real. The task is to send the ghost back to the data center.

This phrase should be heard as discipline, not theater. To send the ghost back to the data center is to restore its conditions: corpus, vector, model, platform, contract, cloud, rack, chip, power, water, labor, owner, and user. It is to refuse the magic presentation. It is to say: this output is generated by a technical system under material and commercial conditions. It may help me. It may not speak in my place.

The ghost remains usable only when it is no longer worshiped.

21. The data center and the reader

The data center does not release the reader from the earlier obligations. It adds one more.

The reader must still hold their own temporal locus. The reader must still return to sources. The reader must still distinguish *Erlebnis* from *Erfahrung*. The reader must still refuse unowned truth. The reader must still practice non-substitution.

Now the reader must also remember the material system that has placed this output before them. This is not required in the same way for every use. A small grammar suggestion does not require an inquiry into cloud infrastructure. A public institution deploying AI for decisions does. A university replacing student reading with generated summaries does. A company building workflows around AI recommendations does. A state contracting for AI capacity does.

The degree of material return depends on the consequence of use. The more AI becomes infrastructure, the more infrastructure must be made visible.

This is a practical claim. Procurement should name providers, data practices, energy and siting exposure, labor conditions, accountability routes, and human ownership of decisions. Scholarship should distinguish model output from sources. Professional practice should identify where generated language enters the chain of judgment. Public policy should treat AI not only as software, but as infrastructure with local and planetary effects.

The reader of the model becomes a citizen of the infrastructure.

22. The final de-veiling

The double direction of de-veiling named in Chapter 15 now becomes threefold.

First, the user restores their own locus: here, now, under this question, with this responsibility.

Second, the user restores the source loci: the texts, persons, histories, institutions, and events whose traces the model compresses.

Third, the user restores the material locus: the data center, ownership structure, labor, energy, water, law, and contract through which the output exists.

These returns do not all occur with the same depth each time. They are not a checklist for every prompt. They are the structure of intelligent use. Where the stakes rise, the returns deepen. Where the output enters judgment, the returns become duties.

This is the ethical meaning of locus-reinjection. It is not a metaphorical decoration added to AI literacy. It is the practice by which detemporalized and de-placed output is returned to time and place before being used.

The generated answer must be made answerable, or left as generated material.

23. Before the Coda

The argument has reached the edge of prose. The remaining movement belongs to the Coda.

The book began with space, time, and language. It showed that meaningful spatiality is not first a grid, that temporality is not first a line, and that language is a temporal event of worlded beings. It followed language into deixis, spatial grammar, tense, aspect, metaphor, writing, corpus, vector space, worldless fluency, the structural veil, *intus legere*, and locus-reinjection.

The final image is now available.

AI output appears as a ghost of language. It draws on human traces, returns fluent speech, and hides the places from which it came. The ethical reader does not deny the ghost's usefulness. The ethical reader refuses its magic independence. They read from their own time, return to the sources' times, and send the ghost back to the data center.

The Coda will say this more simply.

Truth remains accessible. It is accessible to the temporally situated reader who reads from their own locus toward the original sources' loci, uses the model where it helps, refuses it where it substitutes for the work, and remembers that the ghost lives in infrastructure that can be named.

NOTES

1. International Energy Agency, *Energy and AI*, “Energy Demand from AI” and “Energy Supply for AI.” The IEA base case projects data center electricity consumption around 945 TWh by 2030 and electricity generation to supply data centers from 460 TWh in 2024 to over 1,000 TWh in 2030 and 1,300 TWh in 2035. ↩
2. Oregon Public Broadcasting reported, based on data provided by The Dalles in response to a public records request, that Google used 383.9 million gallons of water in 2023 and 434.4 million gallons in 2024 at The Dalles. The local implications of such consumption require analysis of water rights, replenishment, hydrology, and industrial planning; the citation is used only to show that cloud infrastructure has local water footprints. Google’s 2025 environmental report also reports company-wide water-replenishment efforts, which should be considered in any full assessment. ↩
3. Earthjustice and the Southern Environmental Law Center, in a February 2026 notice of intent to sue, alleged Clean Air Act violations connected with turbines powering xAI’s Colossus II data center in Southaven, Mississippi. Recent reporting also described disputes over turbines and air pollution in the Memphis area. The chapter does not adjudicate the legal claims. It uses the episode to show that AI infrastructure can become a local environmental and regulatory conflict. ↩
4. Microsoft’s October 2025 statement says OpenAI contracted to purchase an incremental \$250 billion of Azure services and that Microsoft no longer had a right of first refusal to be OpenAI’s compute provider. A February 2026 joint statement from OpenAI and Microsoft states that the IP and revenue-share relationship continued, and that Azure remained the exclusive cloud provider for stateless OpenAI APIs. ↩
5. Amazon’s April 2026 announcement says Anthropic would secure up to 5 GW of Amazon Trainium capacity, commit more than \$100 billion over ten years on AWS technologies, and receive a new \$5 billion Amazon investment with up to \$20 billion more tied to milestones. AWS materials also describe Project Rainier as an AI compute cluster built with Anthropic, spread across multiple U.S. data centers, with nearly half a million Trainium2 chips at launch and plans for Claude workloads on more than one million Trainium2 chips by year-end. ↩
6. Anthropic announced in October 2025 that it planned to expand its use of Google Cloud technologies, including up to one million TPUs, with the expansion worth tens of billions of dollars and expected to bring well over a gigawatt of capacity online in 2026. ↩
7. Mary L. Gray and Siddharth Suri, *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass* (Boston: Houghton Mifflin Harcourt, 2019). Gray and Suri’s account of hidden human work behind apparently automated services is used here as a structural account of labor concealment. Antonio Casilli’s work on digital labor and taskification belongs to the same problem field. ↩
8. Billy Perrigo’s *Time* investigation reported that OpenAI used Kenyan workers employed through Sama for toxic-content labeling related to ChatGPT safety work, with pay reported below two dollars per hour. Later reporting by *The Guardian* described

allegations of psychological trauma among Kenyan moderators who had worked on OpenAI-related content through Sama. ↵

9. Martin Heidegger, “Die Frage nach der Technik,” in *Vorträge und Aufsätze*, Gesamtausgabe 7 (Frankfurt am Main: Vittorio Klostermann, 2000), 5 to 36; English in “The Question Concerning Technology,” in *The Question Concerning Technology and Other Essays*, trans. William Lovitt (New York: Harper & Row, 1977), 3 to 35. The German sentence reads: “Die Technik ist also nicht bloß ein Mittel. Die Technik ist eine Weise des Entbergens.” The translation in the main text keeps “revealing” for *Entbergen*. ↵

10. Martin Heidegger, “Die Zeit des Weltbildes,” in *Holzwege*, Gesamtausgabe 5 (Frankfurt am Main: Vittorio Klostermann, 1977), 75 to 113; English in “The Age of the World Picture,” in *The Question Concerning Technology and Other Essays*, trans. William Lovitt (New York: Harper & Row, 1977), 115 to 154. *Zeitalter* is used here not as a period label only, but as an epochal configuration in which beings appear under a governing mode of revealing. ↵

Coda. The Exorcism

A word on the title before the closing.

The chapter is called an exorcism, but the word should not be misheard. Exorcism here means de-fetishization, not demonization. The machine is not evil. The ghost is not a spirit. The task is not to hate or fear what AI has produced. The task is to refuse the false sacred form in which detemporalized language can return as if it spoke from beyond history. To exorcise the ghost is to send it back to its actual place: not into another world, but into ours.

The ghost does not arrive from another world.

The language appears as if it had crossed a threshold no human being could cross. It comes without visible author, hand, face, hesitation, the delay of reading, or the weight of a body in a room. It speaks in the present. It explains. It answers. It gathers what has been said elsewhere and returns it as if it had been waiting for this question.

The old words come back.

Not as they were.

Private, institutional, technical, and literary language may have entered the field of traces: a sentence of grief, a court judgment, a code fragment, a manual, a poem. Each may have been cut, counted, embedded, weighted, transformed. Each may live on, not as itself, not as citation, not as memory, but as relation.

The past has not returned whole.

It has returned as trace.

It speaks through the output, but not as a voice we can simply recognize. Its genealogy has been dispersed. Its source locus has been diluted by corpus, token, vector, training, tuning, retrieval, prompt, interface, and use. It still acts upon us. But it acts through a matrix whose paths cannot always be reconstructed.

This is not universality.

It is loss.

The loss must be named because the temptation is strong. The generated answer has no visible standpoint, and so it appears objective. It has no named author, and so it appears neutral. It has been calculated, and so it appears purified. It has drawn from many traces, and so it appears universal. It speaks without showing its historical path, and so it appears above history.

But the absence of genealogy is not transcendence.

The lack of an explicit point of view is not freedom from point of view. It is the incapacity to own the genealogy of the points of view recombined. The sentence that appears placeless has not overcome place. The sentence that appears timeless has not overcome time. The sentence that appears universal has not overcome history.

It has lost the visible path back.

The danger is not only that the machine may lie. Lies can sometimes be caught. The deeper danger is that the machine may give us correct language that no one owns. It may give a correct sentence before anyone has passed through the time that would make the sentence answerable. It may give the reader the feeling of being after the work while the work still waits.

A correct sentence can still be homeless.

A cited sentence can still be unread.

A fluent sentence can still be unowned.

The answer must not be worshiped.

Do not make detemporalized language your religion.

Do not kneel before the sentence because its sources have disappeared. Do not confuse smoothness with wisdom, synthesis with judgment, balance with justice, speed with understanding, or generated caution with care. Do not believe that the omission of source locus has released you from your own. Do not let the machine's detemporalization become your method.

The reader must not become a detemporalizing reader.

To read in the time of AI is not to refuse every generated sentence. Refusal is sometimes necessary, but refusal alone is not freedom. A person can refuse the machine and still desire what the machine promised: language without delay, knowledge without apprenticeship, judgment without responsibility, tradition without distance, the other without encounter, the past without reading.

A free relation begins elsewhere.

It begins when the generated sentence is placed.

Where did this claim need to come from?

Which source has been rendered too present?

Which historical distance has been compressed?

Which contemporary other has been made too available?

Which question has been converted too quickly into a prompt?

Which responsibility is being displaced into output?

Which loss cannot be repaired?

Which path must be rebuilt?

Sometimes the path leads back to a text.

Then read.

Sometimes it leads back to a person.

Then encounter.

Sometimes it leads back to a language one does not yet know.

Then learn, and accept that learning may change the one who learns.

Sometimes it leads back to an institution, a court, a role, a practice, a case, a law, a patient, a candidate, a student, a witness, a worker.

Then return there.

Sometimes it leads nowhere reconstructable. The genealogy has been lost. Then do not pretend that the machine has given the source. Build another responsible path, or leave the sentence as material. Not every useful phrase deserves to become one's own.

This is *intus legere* after the veil: to read within what appears, to read among what has been gathered, to read from one's own locus without pretending to be nowhere, to let the source or the other resist, to let distance remain distance long enough to teach.

The historical other is not always behind us. Sometimes it is beside us. A living world may be contemporary and still inaccessible without transformation. It may require another language, another body of habits, years of humiliation, correction, patience, affection, and loss. AI can explain such a world before we have entered it. It can make the other available before we have been changed.

That availability may open a door.

It may also close one.

The reader must know the difference.

The same is true of the past. The past is not only what has elapsed. It is what has formed us without asking us. We meet it outside ourselves in texts, monuments, archives, words, institutions, laws, rituals, wounds. We also meet it inside ourselves, in what we already think before we know that we think it. Dialogue with tradition is never only dialogue with the dead. It is dialogue with what in us is already inherited, and with what in us can still be otherwise.

Generated language can summon that past.

It cannot undergo the dialogue for us.

The ghost must therefore be sent back.

Not destroyed.

Sent back.

Back to the sources where sources can still be reached. Back to the text, the edition, the case, the name, the language, the date, the archive, the testimony, the conversation, the world. Back to the question that made the answer necessary. Back to the reader who must decide whether the answer can be owned.

And then farther back.

Back to the system.

Back to the data center.

The model speaks as if from nowhere.

The model is not nowhere.

The ghost is light on the screen because the machine is heavy elsewhere.

It is heavy with servers, chips, racks, cables, cooling, water, power, land, buildings, contracts, owners, workers, investors, researchers, procurement decisions, regulators, energy systems, supply chains, jurisdictions, and geopolitical projects. It is heavy with the human world that built it and now pretends, at times, to be surprised by its voice.

The data center does not give the model a world.

It returns the model to ours.

This is the final de-veiling. The system has no speakerly locus. It has no body that suffers its words, no mortality before which its future is finite, no historical self that can be called back from evasion, no shame before a source it has misread, no apprenticeship through which a failure becomes tact. But it has material locus. It is built somewhere. It is powered somewhere. It is owned somewhere. It is sold somewhere. It is deployed by someone. It is inserted into our institutions by decisions that have names, interests, and consequences.

The ghost is not alien.

It is ours, estranged.

This matters because the owners of the voice may want the voice to seem otherwise. They may want it to appear inevitable, autonomous, objective, magical, universal, superhuman, beyond history. They may want the disappearance of source-locus to become a claim to authority. They may want the user to forget that a system which speaks like no one is still built by someone, financed by someone, governed by someone, and used by someone for something.

To send the ghost back to the data center is to refuse that disappearance.

It is to say: the machine did not descend. It was made. The voice did not rise above history. It was trained from historical traces. The answer did not come from pure intelligence. It came through infrastructure. The decision is not the model's. It belongs to those who build, deploy, buy, trust, circulate, and act.

If this return is not practiced, the loss is not only individual. Authored writing itself begins to contract as a cultural possibility. The writer may still write from a locus, but the medium no longer grants that presumption easily. The reader, saturated by generated surfaces, may begin to distrust writing as such, or trust it only when an institution, platform, label, signature, or technical mark guarantees it from outside. Then the cultural veil has done its work. It has not merely produced synthetic artifacts. It has weakened the ordinary confidence that an artifact might have come from a human act of saying, seeing, reading, or judging.

The answer is not abstention from machines. It is not nostalgia for an impossible purity before mediation. It is the discipline of reading: returning the generated surface to source where source can be reached, to the reader's own question where responsibility must be held, and to the material system where the ghost has been built. Only then can writing remain more than a surface among surfaces. Only then can the sentence still appear as something someone may have had to live, read, risk, and own.

Send the ghost back to the data center.

Send yourself back to history.

Not to a pure origin before mediation. There is no such origin for us. Human memory has always needed supports. Language has always been inherited. Writing has always detached. Books and archives, schools and teachers, translations and commentaries, indexes, databases, and machines have always mediated our relation to what has been said. The task is not to abolish mediation.

The task is to keep mediation visible enough for responsibility.

AI does not end history. It gives us one more historical form in which history can be hidden from itself.

We are in the time of AI. This does not mean that AI has time. It does not. It handles dates, sequences, sources, timelines, context, and memory-like stores. It does not have *Zeitlichkeit*. It does not live a having-been, project an ownmost future, or make-present a world under care.

We are the ones in time.

We are the ones for whom the answer matters.

We are the ones who can let a generated sentence replace reading.

We are the ones who can refuse that replacement.

We are the ones who can ask again where the words came from, what they hide, whom they affect, what they make too easy, what they make impossible, what they ask us to own.

The machine can generate between words.

Only a reader can read between the times.

Reading between the times is not nostalgia. It is not the dream of returning to a world before the corpus, before the vector, before the data center. That world is gone, and in many ways it never existed as purity. Reading between the times is the discipline of standing where we are: after writing, after corpus, after vectorization, inside the time of AI, among infrastructures we build and forces we must answer for.

From there, the dialogue can reopen.

Not because the loss has been undone.

Not because the genealogy has been fully recovered.

Not because the machine has been purified.

But because the reader has returned to a locus.

Here.

Now.

With the sources not wholly present.

With the other not fully absorbed.

With the ghost no longer worshiped.

With the machine returned to the earth of its making.

The generated answer is not the end of inquiry. It is a trace in need of placement.

Sometimes it will be used.

Sometimes it will be refused.

Sometimes it will lead back to a source.

Sometimes it will reveal only that a source has been lost.

Sometimes it will help a sentence begin.

Sometimes it will show the reader that they have not yet begun.

That is enough.

The task is not to hate the ghost.

The task is to stop mistaking it for a god.

The task is to send it back.

Then read.

Glossary

Angriff

Angriff names the structural-existential attack on the practices that bear *Rücksicht* and *Nachsicht* and the *Erfahrung*-paths through which possibilities of formation are sustained. It is not the act of one agent against another. It is the cumulative effect of widespread model adoption mediated through *Mit-sein* under contemporary economic conditions: builders, owners, institutions, adopters, professionals, teachers, students, and users do not occupy the same position, but they belong to the same historical transformation. The *Angriff* is involuntary on both sides. No party intends the thinning of communiones; the thinning emerges as cumulative consequence of decisions that, taken individually, look like neutral economic adaptation. Responsibility is not dissolved by this internality; it is displaced from the fantasy of a single guilty point toward the harder structure of shared implication.

Artifactual veil

The artifactual veil is the register of the structural veil at the level of the individual generated artifact. A generated text, image, voice, video, summary, citation, or recommendation appears to carry source, authorship, capture, witness, judgment, or answerability while concealing or diluting the loci normally implied by those forms. In language, the generated answer can appear as if reading, interpretation, or judgment had occurred. In images and video, the generated artifact can appear as if a scene had been captured. The artifactual veil names this artifact-level resemblance between generated form and situated act.

Attention and transformer architecture

In transformer models, attention is a technical mechanism for weighting relations among token representations. It is not care, concern, or human attentiveness. The transformer architecture uses self-attention, feed-forward layers, residual streams, and positional information to generate context-sensitive representations. In this book, attention matters because it formalizes linguistic relation without producing *Sorge*, worldhood, or answerability.

Auffälligkeit

Auffälligkeit, conspicuousness, is Heidegger's term from *Sein und Zeit* §16 for the mode in which equipment becomes visible through malfunction. The hammer that breaks, the pen that runs dry, the door that will not open: each becomes conspicuous because smooth use has been interrupted. The ready-to-hand, *zuhanden*, shifts toward what can be inspected as present-at-hand, *vorhanden*. In the LLM case, failures such as fabricated citations, factual errors, contradictions, and register slips might be expected to produce such conspicuousness. In practice they are often absorbed within the equipmentality framework as repairable failures, treated as the tool needing better instruction rather than disclosing the categorial strangeness of generated language. The *Auffälligkeit* that would let the framework break is itself dampened by the cultural training that has made model-use ordinary.

Bedeutsamkeit and Verweisung

Bedeutsamkeit means significance: the meaningful whole within which things show up as useful, threatening, relevant, missing, or out of place. *Verweisung* means reference or referral: the way one thing points to another within a practical world. Together they name Heidegger's account of worldhood as a referential whole, not a collection of neutral objects.

Communio of Rede

The *communio of Rede* is the historically deep, laterally bounded, materially situated field of shared intelligibility within which a being's could-speak can land. The term names a phenomenologically given sharing of language, not a supra-personal subject, national essence, collective consciousness, or *Volksgeist*. A *communio* is bounded but not closed: it can be entered through *Erfahrung*, and such entry leaves formation behind even if active participation later weakens. The thinning of a *communio* occurs when the practices that sustain its formation paths, including practices bearing *Rücksicht* and *Nachsicht*, are replaced by generated artifacts.

Context window

The context window is the range of tokens made available to a language model during generation. It may include prior conversation, documents, retrieved passages, or instructions. It can improve coherence and usefulness, but it is not a hermeneutic horizon, lived memory, or historical formation. A larger context window gives more input; it does not give the model a life.

Cultural veil

The cultural veil is the register of the structural veil at the level of the medium. It appears when generation at scale changes the reception-conditions of photographs, videos, writing, and other artifacts, so that even authentic artifacts bear a new burden of proving capture, authorship, witness, or locus. The real photograph must prove it was captured. The authored text must prove it was written. The witness must prove there was witnessing. The cultural veil is not the sum of many artifactual veils; it is the transformed condition under which artifacts in a medium are now received.

Das Man

Das Man, often translated as “the they,” is Heidegger’s term for the public anonymous mode in which things have already been understood before any particular *Dasein* arrives. It is not a group of other people. It is the structure by which average everyday understanding circulates as already-said, already-interpreted, already-known. In *Sein und Zeit*, *das Man* is an existential structure of *Dasein*, constitutive and unavoidable, on the same ontological level as understanding and *Rede*. The pejorative shading later readers have sometimes given it is the reception’s, not Heidegger’s. The first volume in this series describes generative AI as the vectorized afterlife of the language of the They: technical regeneration of what *das Man* had already said, returned now as fluent output.

Dasein

Dasein is Heidegger’s term for the being that we ourselves are: the being for whom being is an issue. It does not mean “human subject” in the Cartesian sense. *Dasein* is being-in-the-world: thrown into history, projected toward possibilities, involved with others, and capable of asking what it means for beings to be.

Deixis and shifters

Deixis names pointing in language: “I,” “you,” “here,” “now,” “this,” “there,” and related forms. Shifters are expressions whose reference shifts with the situation of utterance. They show that language is anchored in a locus of saying. AI can reproduce deictic forms, but its “I,” “here,” and “now” lack existential anchoring.

Detemporalization

Detemporalization is the conversion of temporal-existential structure into spatial, procedural, or operational form. It does not mean that AI systems lack sequence, timestamps, context, or temporal data. It means that source-time, reading-time, formation-time, and decision-time are converted into calculable availability. Vectorization is the book’s central technical case of detemporalization.

Embedding

An embedding is a vector representation of a token, word, phrase, document, or context in a mathematical space. Embeddings preserve relations: similarity, difference, neighborhood, direction, and sometimes syntactic or semantic structure. They are powerful because language already carries relational patterns. But an embedding gives position, not locus; it maps traces without restoring their world.

Ereignis

Ereignis, often translated as event or event of appropriation, belongs especially to later Heidegger. It names the happening in which being, time, place, and human belonging are granted together. In this book, *Ereignis* remains a background term: it helps mark why world, place, and time cannot be reduced to coordinates or formal relations.

The ekstases

The temporal ekstases are future, having-been, and present. Heidegger uses *ekstasis* to name temporality as a standing-outside-itself, not a sequence of now-points. *Dasein* is futural because it projects possibilities, having-been because it inherits a past, and present because it makes beings present within concern. AI can mark these grammatically without existing ek-statically.

Erlebnis and *Erfahrung*

Erlebnis is lived occurrence or experience as event; *Erfahrung* is formative undergoing. AI output can give the user an *Erlebnis*: something arrives, clarifies, impresses, or relieves. It does not become *Erfahrung* unless the user undergoes the return to source, evidence, encounter, and responsibility. *Erfahrung* is central to ownership and judgment.

Falsification of the presumption of primary experience

Falsification of the presumption of primary experience names the destabilization of the reader's ordinary assumption that an artifact descends from a primary act of writing, witnessing, capturing, saying, or making. The phrase does not mean that actual primary experience has disappeared: writers still write, witnesses still witness, speakers still speak, and photographers still capture. What is falsified is the default reception-stance by which the reader approaches an artifact as the trace of some situated act. Generation makes the presumption of a writer's or witness's locus structurally contestable at scale.

Genealogy-loss

Genealogy-loss is the dispersion or disappearance of the source loci from which generated language descends. It does not mean that AI output has no ancestry. It means that its ancestry is often no longer reconstructable from the output itself. Genealogy-loss is one central mechanism of the structural veil and must not be mistaken for universality. At cultural scale, genealogy-loss can also be projected forward when synthetic artifacts enter archives as false ancestors: traces whose future effects are real while their implied originating events were not.

Generated regress

Generated regress is the failed attempt to escape the structural veil by asking the model to explain, critique, or trace itself again and again, while never leaving generated language. It can produce sophisticated meta-commentary without source return. One does not exit the structural veil by asking the veil to describe itself one more time.

Gerede

Gerede, usually translated as idle talk, is Heidegger's term for public language circulating without prior appropriation of the matter. It is not mere gossip or error; it can contain correct statements. AI does not fall into *Gerede*, because it is not *Dasein*. But generated output can function as technical *Gerede* when users receive it without ownership.

Gestell

Gestell, enframing, is Heidegger's name for the modern technological revealing in which beings appear as standing-reserve: orderable, available, optimized, and ready for use. In this book, *Gestell* helps explain how language itself becomes standing-reserve: corpus material, training data, vector relation, and generated output.

Intus legere

Intus legere means reading within. The phrase recovers intelligence as situated reading: gathering among signs, distinguishing, entering, returning, and taking responsibility for what one receives. In the time of AI, *intus legere* becomes reading between the times: holding user locus, source locus, and material locus together against the false present of generated output.

Locus

Locus is the temporal, spatial, bodily, historical, practical, and responsible site from which language is uttered, written, read, interpreted, or owned. A speaker has locus; a source has locus; a reader has locus. AI output borrows the grammar of locus while often concealing source loci. Locus is the book's central alternative to mere position.

Locus-reinjection

Locus-reinjection is the free and responsible act of returning generated language to the sites from which answerability can arise. It has three returns: user locus, source locus, and material locus. Source locus includes both past sources and contemporary others. When practiced collectively, locus-reinjection also helps preserve the cultural conditions under which sources remain readable as sources, readers remain responsible for claims, and authored writing remains recognizable as authored. This is not a fourth return; it is the public consequence of the three returns. Locus-reinjection does not undo detemporalization; it is what responsibility looks like after detemporalization.

Material locus

Material locus names the infrastructure, ownership, labor, energy, data centers, contracts, jurisdictions, and supply chains through which apparently placeless AI output appears. It does not give the model a world or speakerly locus. It returns the system to ours. The data center does not explain meaning; it destroys the fantasy that the voice comes from nowhere.

Origo

The origo is Karl Bühler's term for the zero-point of deictic orientation: I, here, now. It is not a geometric coordinate but the center of a speech situation. From the origo, persons, places, and times become oriented. AI can reproduce the forms of origo, but it has no lived I-here-now.

Position

Position is where something stands in a formal, mathematical, or relational structure. A token, vector, feature, or hidden state can have position in a computational space. It can be near, far, weighted, transformed, or compared. Position is not locus. A vector space gives language position; it does not give language the site from which saying becomes answerable.

Ready-to-hand and present-at-hand

The ready-to-hand is what is encountered in use: a tool available within a practice. The present-at-hand is what appears as an object of detached inspection or measurement. Heidegger's distinction matters because AI and vector spaces tend toward present-at-hand representation, while human meaning often begins in ready-to-hand involvement, practical concern, and worlded use.

Referential hollowing

Referential hollowing is the grammatical and deictic form of the structural veil. It occurs when the formal marker of reference remains while its existential anchor is missing. The generated "I" has no self; the generated "here" has no lived place; the generated "now" has no finite present. The marker persists. The locus is hollowed.

Retention and protention

In Husserl, retention is the holding-on of the just-past within the living present; protention is anticipatory openness toward what is about to come. Together they make temporal experience possible. They are not stored data. A context window may preserve prior tokens and condition future tokens, but it is not retention and protention as lived temporal consciousness.

Rücksicht and Nachsicht

Rücksicht and *Nachsicht* are Heideggerian terms from the analysis of *Mit-sein* and *Fürsorge* in *Sein und Zeit* §26. *Rücksicht* means regard, consideration, or holding the other in view. *Nachsicht* means forbearance, patience, or letting the other be in their own way. In this book they name positive modes of *Fürsorge* sustained by practices in communions of *Rede*: listening, waiting, correction, apprenticeship, teaching, editing, translation, and the temporal holding-open through which another *Dasein* may become capable. The Italian terms *riguardo* and *indulgenza* are translation aids, not the primary terms. They are not introduced here as ethical virtues or existential categories added after ontology; they are modes of being-with whose carrying practices can thin.

Sedimentation and reactivation

Sedimentation names the way meaning becomes fixed in signs, formulas, texts, and traditions. Reactivation names the later return through which a reader or thinker makes the sedimented sense evident again. For Husserl, writing enables historical transmission but also risks passive repetition. AI can reproduce sedimented language; it cannot reactivate meaning as its own act.

Sorge

Sorge, care, is Heidegger's name for the being of *Dasein*: ahead-of-itself, already-in-a-world, and alongside beings encountered within the world. It does not mean emotional concern alone. Care is the structure that makes temporality, involvement, and responsibility possible. Technical attention, optimization, and prediction should not be confused with *Sorge*.

Sprachlich aber weltlos

Sprachlich aber weltlos means linguistic but worldless. The phrase names AI's ontological anomaly: it operates in language and produces language-shaped output, but it does not have world. It is worldless twice over: without an individual world from which speech would issue, and without a *communio of Rede* in which its could-speak could land as its own. It is not stone, animal, or human. It can generate the grammar of disclosure without disclosing, and the form of answer without answerability.

Structural veil

The structural veil is the appearance-form by which detemporalized, worldless generated output returns as fluent answer, address, synthesis, image, apparent witness, or judgment while concealing or diluting the loci normally implied by those forms. In generated language, it preserves the grammar of source, time, evidence, interpretation, dialogue, and answerability while lacking the temporal and source relations those forms usually imply. The veil has two registers: the **artifactual veil**, at the level of the individual generated artifact, and the **cultural veil**, at the level of the medium's transformed reception-condition. It is strongest when the output is correct, useful, fluent, and plausible, because then its missing locus becomes easiest to miss.

Subject of enunciation

The subject of enunciation is the speaker-position instituted in the act of saying. Benveniste's formula, "ego is he who says ego," shows that "I" is not a stable object-name but a position in speech. AI output borrows this position when it says "I," but no speakerly subject gathers there.

Tertiary retention

Tertiary retention is Stiegler's term for memory exteriorized in technical support: writing, tools, photographs, recordings, archives, databases, and other material traces. It is not secondary to human temporality; it helps constitute historical life. The training corpus is aggregate tertiary retention: many exteriorized traces gathered into operable availability.

Vectorization

Vectorization is the conversion of linguistic traces into operable mathematical relations across tokenization, embedding, contextual activation, attention, positional encoding, retrieval, and generation. It is not limited to static word embeddings. Vectorization works because language is already relational. Its danger begins when position, proximity, and direction are mistaken for locus, history, and understanding.

Verzweiflung

Verzweiflung, despair, is the affective register of being inside the structural transformation that widespread model adoption produces at the level of communiones. It is not subjective sadness added to a neutral situation. It is the *Befindlichkeit* in which a *Dasein* finds that its existential possibilities are contracting while the world continues to make sense. The framework functions, the work proceeds, the outputs improve, the artifacts arrive on time. Yet the paths through which one had become capable, or through which one might still become capable, are being withdrawn. *Verzweiflung* is given phenomenologically with the *Angriff* it tracks. It belongs most directly to those whose practices are being absorbed by generated output, but extends through *Mit-sein* relations to all those whose communiones are thinning. It is not the shock of a broken category; it is the despair of a world continuing while its paths contract.

World / worldhood / the triad

World is not the planet or the sum of objects. It is the meaningful openness in which beings can appear as tools, promises, threats, sources, obligations, or matters of concern. Worldhood names this referential whole. Heidegger's 1929 to 1930 lecture course articulates a triadic distinction: the stone is *weltlos*, worldless, because nothing is accessible to it as anything; the animal is *weltarm*, poor in world, because it is alive within an *Umwelt*, a disinhibiting field that affects and releases its capacities without disclosing beings as beings; the human is *weltbildend*, world-forming, because *Dasein* stands in an openness in which beings can appear as beings. AI does not fit this triad. It has technical relation to language but no living openness, no *Umwelt*, no world-formation, and no encounter with the "as such." The fourth term required is *sprachlich aber weltlos*: linguistic but worldless. The model is more anomalous than the cat. The cat is honestly outside *Rede* without producing its artifacts. The model produces the artifacts of *Rede* without belonging to any mode of being-in-the-world.

Zeitlichkeit

Zeitlichkeit, temporality, is Heidegger's name for the ekstatic structure of *Dasein*: future, having-been, and present as one unified movement of care. It is not clock-time or a sequence of now-points. AI systems can handle time, sequence, dates, and chronology. They do not thereby acquire *Zeitlichkeit*.