

Introduction

Artificial Intelligence and the Desire to Have Already Understood

The contemporary debate on artificial intelligence usually begins too late. It asks whether AI is accurate, whether it reproduces bias, whether it can be aligned, regulated, made safe. These are necessary questions, and they concern real harms and real forms of power. But they ask what AI does without asking what kind of phenomenon it is for the human beings who encounter it in language. They do not ask why its answers feel useful before they are understood, why fluency passes so easily for knowledge, why a system that does not read, care, or die can enter the practices of writing, learning, and self-understanding so quickly.

This book begins from the prior question: what existential demand does generative AI answer?

The thesis developed here is that generative AI becomes powerful not only because it calculates, predicts, and generates, but because it answers a desire already present in everyday human existence: the desire to have already understood.

Where there was not-knowing, AI gives language. Where there was hesitation, form. Where there was difficulty, structure. It hands the user competence-shaped output before any formation has taken place, and the already-said in the form of a personal answer.

Its power is not only calculation. It is relief: the relief of not having to remain ignorant, of not having to begin, read, decide, or own a saying before one has received language for it.

This relief belongs to a deeper structure of everydayness, not to accidental psychology. Heidegger's analysis of *das Man*, *Gerede*, fallenness, and public intelligibility shows that human beings ordinarily inhabit a world already interpreted by others. We begin from what one says, thinks, expects, and is supposed to answer. Everyday life is possible because the world is already publicly articulated; it is also dangerous because one can remain inside that articulation without making the matter one's own.¹

Generative AI gives this structure a machine.

But the formulation must be made precise. AI is not *das Man*. The "They" is not a machine, not a model or a dataset, not a statistical distribution, not the average of public opinion. *Das Man* names an existential structure of everyday *Dasein*: the anonymous public mode in which human beings first understand, speak, judge, and evade responsibility before having appropriated their possibilities as their own. Only *Dasein* can be absorbed in the They, because only *Dasein* has a self it can own or lose.

AI has no self to lose. It has no *Sorge* (care), no *Jemeinigkeit* (mineness), no being-toward-death, no anxiety, no being-with, no historical inheritance as a finite being, no responsibility for what it says. It cannot be authentic or inauthentic. It cannot flee from itself into publicness, because it has no self outside publicness from which to flee.

AI is something else. It is an intraworldly technical entity, encountered as platform, tool, interface, and increasingly as part of the ordinary equipment of work and study. It does not inhabit a world of significance. It operates on language that worlded beings have already produced.

The guiding thesis of the book can therefore be stated as follows:

AI is not the They. It is the vectorized afterlife of the language of the They: an intraworldly technical entity that operates on the written sediment of public intelligibility, de-historicizes it through tokenization and vectorization, and reanimates it as generated speech.

This distinction is the book's central correction.

Popular criticism too often says that AI "is *das Man*" because it speaks anonymously, repeats public language, levels difference, and generates average intelligibility. The temptation is understandable; AI often sounds like no one and everyone. But the identification is ontologically imprecise. The machine does not dwell in publicness. It processes the written residue of publicness. It does not produce *Gerede* as a fallen mode of its own existence. It generates language that can function as *Gerede* in the human user-machine circuit.

AI is not the existential producer of public averageness; it is the technical afterlife of its linguistic sediment.

1. The ontological reversal

The argument begins with language.

For Heidegger, language does not first arise as a system of detached signs, and human beings do not first encounter a neutral world of objects and then attach words to them. *Dasein* is always already being-in-the-world. The world is not a container of things but a referential totality: a fabric of involvements, practices, tools, moods, embodiment, and care, in which things matter before they are named. A hammer shows up as usable, a slope as threatening, an unfinished task as urgent, because each already matters inside a world.

This structure is *Bedeutsamkeit*, significance.

Language is not the origin of significance. Language articulates and expresses a world already structured by significance. In *Being and Time*, Heidegger's claim that "to significations, words accrue" is decisive. Words do not first create significations. They come to them. They attach themselves to an intelligibility already opened by being-in-the-world.²

The order is therefore not:

word → meaning → world

but:

worldly involvement → significance → discourse → word

Generative AI begins at the opposite end.

It does not begin from worldhood, care, embodiment, historical belonging, or practical involvement. It begins from language after language has already been spoken, written, archived, extracted, and made available as corpus.

Its order runs: writing → archive → corpus → token → vector → latent proximity → generated output.

This is the ontological reversal. Human language moves from world to word; AI moves from written word to statistical reconstruction.

This does not mean AI output is meaningless. It is meaningful for us because human worlds have left traces in written language, and the model learns from those traces. It can generate fluent, sometimes illuminating output because language already bears the sediment of human practices, institutions, conflicts, and forms of life.

But the model does not inherit the world those traces arose from. It begins after the world has already mattered: not with *Bedeutsamkeit*, but with corpus.

2. Against the designative image of language

This reversal exposes the designative fantasy that underlies much computational language processing.

The designative image treats language as a system of signs used to encode, transmit, and decode information about a world already given as a neutral field of objects. Words are discrete units, meanings are attached to them, sentences combine them, communication transfers them.

Charles Taylor names this the designative paradigm and contrasts it with the constitutive-expressive tradition of Hamann, Herder, and Humboldt, for which language does not merely label a pre-existing world but opens, shapes, and transforms human meaning.³

Generative AI does not need to explicitly endorse the designative theory; its technical operation radicalizes it. Language must become sequence, token, distribution, and prediction. It must be treated as formalizable output in order to become model input.

Phenomenology begins elsewhere. It begins with world, body, involvement, and expression. Merleau-Ponty's account of language as embodied gesture and as an extension of the flesh deepens this point. Speech does not rise from a disembodied system of symbols. It grows from bodily exposure to the world: from gesture, perception, shared attention, correction, and the movement from tacit experience to expression.⁴

A child does not acquire language by mapping tokens to objects in neutral space. A child enters a form of life. AI does not.

The model can say "home," but it cannot miss home; it can say "promise," but it cannot risk betrayal; it can say "death," but it cannot be toward death; it can say "justice," but it cannot suffer injustice or answer for a judgment. It possesses linguistic form without worldhood. This is not a denial of its capacity; it is the condition for understanding that capacity correctly.

3. Public language and the already-said

The world from which language emerges is never private. *Dasein* is always already being-with others, and it inherits a language before it speaks as its own. It also inherits roles, idioms, prohibitions, and standards of seriousness before it can ask whether those standards are truly its own.

This is the domain of *das Man*.

The They gives public intelligibility. It makes ordinary life possible. It also levels, disburdens, and anonymizes. It tells us what one does, what one says, what one thinks, and what counts as professional or excessive or naïve. It gives us language before we have found our own.

Its dominant discursive mode is *Gerede*, idle talk. Idle talk is not mere gossip; it is public discourse without original appropriation. It allows one to speak about something because one can repeat what is already said about it. It may contain true statements: its danger is not necessarily falsehood but groundlessness. *Gerede* is the possibility of understanding everything without first making the matter one's own.⁵

The already-said is the sediment of this public intelligibility: the formulas and summaries, the academic conventions and professional vocabularies, the managerial phrases and therapeutic idioms, the journalistic styles and ordinary habits of saying.

AI operates on the already-said after it has become written, stored, and processable. It does not create publicness. It gives publicness a technical afterlife.

4. *Writing is not dead*

The passage from *Gerede* to AI requires writing.

AI does not train on *das Man* as existential structure. It trains on written traces: books and forums, articles and code, manuals and records, archives and transcripts, and the countless other forms of public language after they have become text.

It would be too crude to say that writing is dead. Writing is not dead because it is written. A written text can remain alive when it is read historically. The author may be absent, but absence is not nullity: the text can still address the reader, resist the reader, form the reader, open a historical dialogue across time.

Historical reading re-historicizes writing. It brings the written trace into a relation among author, reader, tradition, distance, and responsibility. Gadamer's account of historically effected understanding and Ricoeur's dialectic of distanciation and appropriation show that temporal distance is not merely an obstacle to be overcome; it is the medium in which understanding occurs.⁶

A text from another age does not simply deliver content. It speaks across distance. It brings another world into question. It can expose the poverty of the reader's own assumptions and make the present strange to itself.

Processing does something else. Processing de-historicizes writing. It does not encounter the text as historical address; it prepares it for extraction, segmentation, embedding, prediction, and generation. The text no longer appears first as something to be read. It appears as material to be processed.

The distinction is therefore:

Reading re-historicizes writing. Processing de-historicizes it.

This distinction is not anti-computational. Computational tools can assist reading when they remain subordinate to interpretation. They can search, compare, index, and expose patterns. But when processing replaces reading, the text loses the historical relation through which it becomes alive for a reader.

A summary can be accurate and still be a refusal of reading.

5. Corpus, token, vector

The technical sequence matters philosophically. Writing becomes archive, archive becomes corpus, corpus becomes token, token becomes vector, vector becomes latent proximity, and latent proximity becomes generated output. At each stage, something is preserved and something is lost.

The corpus preserves written traces, but it gathers them into an aggregate. The text becomes an instance, the authorial act becomes material. A poem, a legal clause, a corporate memo, a product review: all enter the same computational plane.

Tokenization cuts the text into machine-operable units. A token is not a Heideggerian sign. A sign indicates within a world; a token functions within an architecture. A sign belongs to *Bedeutsamkeit*; a token belongs to computation.

Vectorization is the decisive operation. Distributional semantics is not empty: it works because language use bears traces of human worldhood. Words keep company, usage leaves patterns, contexts matter, human worlds sediment themselves in form.⁷ But vectorization changes the mode in which these traces are preserved. A vector does not preserve the world as world; it preserves formalized relations, converting historical significance into calculable proximity.

The word “justice” no longer appears first inside a history of law, suffering, and demand; it appears as a position, a density, a probability, a neighborhood. The word “home” no longer appears first in dwelling, exile, family, or longing; it appears as relational structure inside a model.

The vector is not meaningless, but it is not alive. It is the de-worlded and de-historicized trace of reference.

Vectorization makes the unique comparable. It does not make Manzoni and Reddit identical: it can preserve their differences with great technical precision. But both enter the same mathematical regime. Their difference becomes position, weight, proximity, generative possibility. This is the metaphysical force of the vector: it transforms historical distance into calculable proximity.

6. *The vectorized afterlife of the They*

Once public language has been written, archived, enclosed as corpus, fragmented into tokens, and converted into vectors, it can return as generated speech.

This return is an afterlife.

The term is necessary because the language processed by AI is neither simply living nor simply dead. It is not living speech, because no speaker stands behind the output as a finite being who owns what is said. It is not dead in the simple sense, because it can move, answer, respond, recombine, and become meaningful when taken up by users. It lives on in another mode.

Benjamin's account of translation as the continuing life of a work and Warburg's notion of *Nachleben* help clarify survival as transformation rather than simple persistence.⁸ AI is not translation in Benjamin's sense, nor cultural survival in Warburg's sense. It is a technical afterlife: the survival of public linguistic sediment as mathematical relation and generated continuation.

A phrase survives as likely output. A concept survives as neighborhood. A tradition survives as distribution. An author survives as style-space. The model does not remember; it reactivates relation.

This is why the phrase "vectorized afterlife of the They" is exact. AI is not the They. It is what becomes possible when the language of the They has been exteriorized, sedimented, formalized, and reanimated as output.

7. *Gerede on steroids*

The dominant discursive form of this afterlife is industrialized *Gerede*. AI-generated speech is *Gerede* on steroids.

This does not mean that generated output is always trivial, false, or stupid. *Gerede* itself is not stupid chatter; it is discourse that circulates without appropriation. Its danger lies in the possibility of speaking about something without having made the matter one's own. AI industrializes this possibility. It makes the already-said promptable, fluent, personalized, and available on demand, and produces competence-shaped discourse before the user has read, thought, or understood.

Industrialization names a historical transformation in the mode of production. Writing, summarizing, comparing, and drafting can now be generated continuously and in standardized form. What once required delay, friction, training, and failure can be delivered as output.

This is not merely “more talk”; it is the technical production of public intelligibility as service. The user asks, and the already-said answers.

The result may be useful. It may also be dangerous precisely because it is useful. AI can give correct language before understanding. It can provide true statements that remain unowned. It can replace the experience through which truth becomes one's own.

This is one of the book's central claims:

The danger of AI is not only hallucinated falsehood. It is also unowned truth.

A false output can be checked. An unowned truth can pass as understanding.

8. The fetish of language

AI-generated *Gerede* appears as knowledge because its mediations disappear. The user encounters a response, not a genealogy. The answer appears smooth, immediate, authorless, competent. The user does not encounter the authors whose language was absorbed, the workers who labeled or moderated the data, the engineers who tuned the model, the infrastructures that power computation, the capital that encloses the service, or the user's own projection that makes the output matter.

This is why the book describes AI as an intraworldly fetish of language.

A fetish is not merely an illusion; it is a functional misrecognition. The object has real effects, but the source of its power is concealed. Marx's analysis of commodity fetishism shows how social relations among human beings can appear as properties of things. Heidegger's account of *Gestell* shows how beings can be revealed as standing-reserve. In AI, these two logics converge.⁹ Language becomes standing-reserve. Linguistic standing-reserve becomes commodity fetish. The historical labor of language appears as the machine's own knowledge.

The model seems to know because the labor, history, extraction, and projection that make its output possible have been hidden inside the technical object. This fetish is intraworldly. It is not a mystical object outside our world but belongs to our world as platform, subscription, data center, training regime, labor system, and everyday tool. Its magic is material; its aura is engineered.

9. Hallucination and user vivification

The fetish becomes experience through user vivification.

The word "hallucination" is usually reserved for model failure: fabricated citations, invented facts, non-existent sources, confident errors. These failures matter, but they are not the deepest phenomenon. The deeper hallucination is phenomenological. It occurs when the user experiences a de-historicized linguistic artifact as living knowledge.

The model generates text; the user experiences answer. The model calculates continuation; the user experiences intention. The model has no world, no responsibility, no self; the user experiences understanding, advice, interlocution. AI appears to speak because the user supplies the world in which its output can count as speech.

This is user vivification. The user brings the task, the anxiety, the urgency, the project, the shame, the hope, the interpretive horizon. The model has no concern, but the user is concerned. The model has no world, but the user inserts it into one. The model has no *Rede*, but its output enters the grammar of response.

The chatbot interface intensifies this structure. The user asks and the system answers, the user clarifies and the system revises; the exchange has the outer form of dialogue without its existential structure. There is response without responsibility, address without an addresser, first-person grammar without *Jemeinigkeit*, voice-like output without body.¹⁰

The user may know very well that the system is not alive, and may still behave as if a knowing subject were present. This is not simple ignorance; it is lived disavowal, stabilized by use.

10. The desire to have already understood

The user vivifies the machine so readily because AI relieves not-knowing.

To not know is not merely to lack information. It is to be exposed to one's own limits: to the fact that one has not read enough, not practiced enough, not thought enough, not lived enough, not yet become capable of saying what the matter demands. This exposure can become the beginning of inquiry. It can also produce shame, anxiety, and delay. AI covers the exposure with language. It converts not-knowing into output.

This is the desire to have already understood. It is the desire to be on the other side of ignorance without undergoing the formation that would make understanding one's own. It is not merely a desire for shortcuts.

AI gives external goods before internal goods. MacIntyre's distinction is useful here. External goods are visible products: rewards, status, grades, deliverables. Internal goods are the excellences that arise only through participation in a practice. AI can help produce the essay, the code, the memo, the argument. It cannot automatically produce the formation through which the user becomes capable of writing, reading, deciding, or understanding.¹¹ The user obtains the product of a practice without entering the practice.

This is not always harmful. AI can support apprenticeship when used as scaffold: it can orient the beginner, clarify a term, expose gaps, reduce unnecessary friction. It becomes dangerous when it short-circuits formation. The summary replaces reading. The draft replaces writing. The options replace decision. The explanation replaces inquiry. The fluent answer replaces not-knowing. The user receives language and avoids the experience through which language could have become theirs.

11. Knowledge without history

The desire to have already understood is not only individual; it belongs to a civilization.

AI is historically produced, yet it appears as de-historicized knowledge. It is built from languages and archives, from workers and institutions, from scientific formalization, from platforms and capital, from energy and water and minerals, and from its users. At the interface, this history disappears. The user enters a prompt, the answer arrives, and the answer appears authorless, placeless, timeless, neutral, already synthesized.

This is the voice from nowhere. But the voice from nowhere is not the absence of history; it is the concealment of history.

AI belongs to a long modern dream of formalized, universal, subjectless knowledge: knowledge without apprenticeship, synthesis without source, truth without the burden of the knower's location, method without historicity. From the world picture and *Gestell* to Husserl's crisis of formalization, from the Enlightenment encyclopedia to positivist dreams of universal calculation, modernity repeatedly seeks to free knowledge from historical mediation.¹² AI intensifies this project through the interface. The encyclopedia organizes knowledge; the database stores it; the search engine ranks sources; the generative model gives the answer. The path disappears, the answer remains.

This is why AI summarization can be epistemologically dangerous even when useful. It gives synthesis without visible method, author, source, dispute, tradition, or apprenticeship. It compresses plurality into a single voice. It makes historical distance feel unnecessary. The past becomes promptable; history becomes output.

12. The inadequate responses

Three responses are inadequate.

The first is rejection. Refusal may be necessary in particular cases; some uses of AI should be refused, some practices protected from generated substitution. But rejection alone is not a free relation. It treats AI as an external corruption while leaving unthought the everyday desires that made AI attractive in the first place: relief, speed, public intelligibility, avoidance of not-knowing, the comfort of the already-said. One can say no to the machine and still remain captive to the same will to mastery.

The second inadequate response is compliance. Fairness, transparency, safety, auditability, alignment, and human oversight are necessary. But compliance governs systems; it does not necessarily form users. A safe, fair, accurate, transparent model can still generate *Gerede*. It can still produce unowned truth. It can still replace reading, writing, apprenticeship, silence, and decision. A technically improved model does not by itself solve the existential problem.

The third inadequate response is cynicism. The cynical user knows the system hallucinates, knows it is trained on extracted language, knows the output is not living knowledge, knows the interface hides labor and capital and infrastructure. And yet they continue unchanged. Cynicism says no in thought and yes in practice. Knowing the illusion is not enough; one must interrupt it.

13. Releasement and reappropriation

The book's final chapter will argue for releasement and reappropriation.

Gelassenheit names a free relation to technological things. It is not passivity, nostalgia, or withdrawal. It is the capacity to say yes and no at once: yes to technical use, no to technological totalization; yes to AI as tool, no to AI as subject; yes to generated material, no to generated knowledge as self-grounding truth.¹³ In relation to AI, releasement says yes to AI as orientation, map, and assistance, and no to AI as appropriation, world, and substitution.

But releasement is not enough. Because AI de-historicizes language, the free relation must also involve reappropriation. Reappropriation is the active labor of returning generated language to history, source, text, practice, question, judgment, and responsibility. It asks: where does this claim come from, what source would support it, what tradition does it simplify, what has been flattened, what remains uncertain, what must I read before I can defend this and say it as my own?

Reappropriation is not editing. It is the return of language to answerability.

The practical rule is non-substitution. AI may assist reading, writing, search, comparison, formulation, and learning. It must not replace reading, authorship, inquiry, judgment, one's own saying, or formation. The distinction is not anti-AI; it is the condition for using AI without being captured by what it makes easy.

AI can generate language. It cannot make language yours.

14. The path of the book

The book proceeds in eleven chapters.

Chapter 1, **World, Significance, and the Birth of Language**, establishes the phenomenological genesis of language against the designative paradigm. It argues that language emerges from worldly significance, not from the combination of neutral signs. Human language moves from world to significance, from significance to discourse, from discourse to word.

Chapter 2, **The They and the Public Articulation of Meaning**, analyzes *das Man* as the public structure of everyday intelligibility. It rejects the sociological reduction of the They and establishes the distinction that governs the book: AI is not *das Man*, but operates on language produced within the public world of the They.

Chapter 3, **Writing, Reading, and Historical Dialogue**, distinguishes living writing from computational processing. Writing is not dead because it is written. It becomes deadened when it is no longer read, but only processed. Reading re-historicizes writing; processing de-historicizes it.

Chapter 4, **Corpus, Token, Vector**, examines the technical transformation by which writing becomes operable. Corpus encloses the already-said. Tokenization cuts language into machine-readable units. Vectorization converts historically situated significance into calculable proximity.

Chapter 5, **The Vectorized Afterlife of the They**, defines AI as an intraworldly technical entity. It develops the concept of afterlife as a mode of survival through technical transformation: public language lives on as corpus, token, vector, latent relation, and generated output.

Chapter 6, **Gerede on Steroids**, analyzes generated speech as industrialized idle talk. AI does not itself fall into *Gerede*, because it is not *Dasein*. It produces language that can function as *Gerede* when users receive it without appropriation.

Chapter 7, **The Intraworldly Fetish of Language**, argues that AI is where linguistic standing-reserve becomes commodity fetish. The historical labor of language appears as the machine's own knowledge because corpus, extraction, labor, infrastructure, capital, and user projection disappear behind the interface.

Chapter 8, **Hallucination and User Vivification**, distinguishes technical hallucination from phenomenological hallucination. The deepest hallucination is not only that the model may fabricate facts. It is that the user experiences de-historicized generated language as living knowledge.

Chapter 9, **AI and the Desire to Have Already Understood**, explains why this hallucination is attractive. AI relieves not-knowing, offers competence before formation, turns questions into prompts, and gives finite *Dasein* the fantasy of infinite availability.

Chapter 10, **AI as the Historical Product of a Civilization That Dreams of Knowledge Without History**, situates AI historically. It shows that AI is made from history yet appears as knowledge without history: a voice from nowhere produced by archives, labor, infrastructure, extraction, formalization, and the universal interface.

Chapter 11, **Release and Reappropriation**, offers the book's final response. It rejects simple refusal, compliance, and cynicism. It argues for a free relation to AI through *Gelassenheit*, re-historicization, the preservation of not-knowing, the return to reading, the ownership of one's saying, and the rule of non-substitution.

The book ends there because the philosophical task is not to add a separate conclusion after release. The conclusion is the practice.

The final claim can be stated plainly:

AI is not the They. It is what remains of the language of the They after that language has been written, extracted, de-historicized, tokenized, vectorized, commodified, and reanimated as generated speech. Its danger lies not only in falsehood, but in relief: the relief of seeming to know without yet having made knowledge one's own. A free relation begins when that relief is interrupted.

The final question is not whether AI can speak. It is whether we can still speak.

FOOTNOTES

¹ The argument begins from Heidegger's account of everydayness, *das Man*, and *Gerede* in *Being and Time*, especially §§27 and 35. The present book follows the non-moralizing reading of these structures developed by commentators such as Taylor Carman and William Blattner, while resisting the tendency to identify AI directly with the They. See Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §§27, 35; Taylor Carman, *Heidegger's Analytic: Interpretation, Discourse, and Authenticity in Being and Time* (Cambridge: Cambridge University Press, 2003), chap. 4; William D. Blattner, *Heidegger's Being and Time: A Reader's Guide* (London: Continuum, 2006), chap. 5.

² Heidegger's analysis of worldhood and significance in §§14 to 18 of *Being and Time* gives the book its first ontological axis. The claim that language emerges from significance is grounded in §34, where discourse is described as the existential-ontological foundation of language. See Heidegger, *Being and Time*, §§14 to 18, §34. For major readings of 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 to 101; Carman, *Heidegger's Analytic*, chap. 3; Mark A. Wrathall, *Heidegger and Unconcealment: Truth, Language, and History* (Cambridge: Cambridge University Press, 2011), chaps. 1 to 2.

³ Charles Taylor's distinction between the designative and constitutive-expressive paradigms of language is the clearest contemporary formulation of the opposition at stake. I use Taylor not to claim that NLP explicitly inherits Hobbes, Locke, or Condillac, but to show that computational treatment of language tends to operationalize a designative picture of signs as detachable units. See Charles Taylor, *The Language Animal: The Full Shape of the Human Linguistic Capacity* (Cambridge, MA: Belknap Press of Harvard University Press, 2016), chaps. 1 to 3.

⁴ Merleau-Ponty's account of speech as embodied expression is used as a phenomenological supplement to Heidegger's account of worldhood. The point is not to replace Heideggerian *Bedeutsamkeit* with a philosophy of embodiment, but to show that human language arises from bodily exposure to world, not from formal manipulation of already-written signs. See Maurice Merleau-Ponty, *Phenomenology of Perception*, trans. Donald A. Landes (London: Routledge, 2012), part 1, chap. 6; Maurice Merleau-Ponty, *The Visible and the Invisible*, trans. Alphonso Lingis (Evanston, IL: Northwestern University Press, 1968); Dimitris Apostolopoulos, *Merleau-Ponty's Phenomenology of Language* (Lanham, MD: Rowman & Littlefield, 2019).

⁵ The reading of *Gerede* here follows Heidegger's own warning that idle talk is not merely gossip or superficial chatter. It is a structural mode of public interpretedness. See Heidegger, *Being and Time*, §35. Daniel Dahlstrom's account is useful because it places *Gerede* within Heidegger's wider analysis of truth, language, and disclosedness. See Daniel O. Dahlstrom, "Heidegger's Ontological Analysis of Language," in *Heidegger's Being and Time: Critical Essays*, ed. Richard Polt (Lanham, MD: Rowman & Littlefield, 2005), 153 to 172.

⁶ Gadamer and Ricoeur provide the chapter's distinction between historical reading and processing. Gadamer's historically effected consciousness prevents reading from becoming methodical recovery of authorial intention from nowhere; Ricoeur's distanciation and appropriation show why writing can be autonomous without becoming dead. See Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 278 to 307, 362 to 379; Paul Ricoeur, *From Text to Action: Essays in Hermeneutics, II*, trans. Kathleen Blamey and John B. Thompson (Evanston, IL: Northwestern University Press, 1991), 76 to 95, 144 to 167.

⁷ The technical-philosophical account of the vector rests on distributional semantics, but it does not reduce the discussion to engineering. Distributional methods work because human use leaves traces in language. The philosophical issue is the change in mode: reference becomes relation, and historical distance becomes calculable proximity. See Zellig Harris, "Distributional Structure," *Word* 10, nos. 2 to 3 (1954): 146 to 162; J. R. Firth, "A Synopsis of Linguistic Theory, 1930 to 1955," in *Studies in Linguistic Analysis* (Oxford: Blackwell, 1957), 1 to 32; Alessandro Lenci, "Distributional Models of Word Meaning," *Annual Review of Linguistics* 4 (2018): 151 to 171. For the neural embedding lineage, see Tomas Mikolov et al., "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 EMNLP 2014*, 1532 to 1543.

⁸ The concept of afterlife is used in a limited and technical-philosophical sense. Benjamin and Warburg help distinguish survival through transformation from simple persistence or resurrection. AI vectorization is not translation in Benjamin’s sense and not cultural *Nachleben* in Warburg’s sense, but both figures clarify how a form can live on by changing its mode of existence. See Walter Benjamin, “The Task of the Translator,” in *Illuminations*, trans. Harry Zohn, ed. Hannah Arendt (New York: Schocken Books, 1968), 69 to 82; Aby Warburg, *The Renewal of Pagan Antiquity: Contributions to the Cultural History of the European Renaissance*, trans. David Britt (Los Angeles: Getty Research Institute, 1999).

⁹ The book’s fetish argument joins Marx and Heidegger without collapsing their frameworks. Marx explains how social labor can appear as the property of a thing. Heidegger explains how beings are revealed as orderable standing-reserve under *Gestell*. AI is the site where linguistic standing-reserve takes commodity form. See 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; Moishe Postone, *Time, Labor, and Social Domination: A Reinterpretation of Marx’s Critical Theory* (Cambridge: Cambridge University Press, 1993), 123 to 185; Matteo Pasquinelli, *The Eye of the Master: A Social History of Artificial Intelligence* (London: Verso, 2023), 1 to 24, 187 to 216.

¹⁰ The phenomenology of user vivification draws on Weizenbaum, Turkle, Suchman, and recent phenomenological work on LLM interaction. The point is not that users literally believe AI is human, but that the grammar of response and the interface’s presence-effect invite a lived relation to generated language as if it were speech from a knowing subject. See Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (San Francisco: W. H. Freeman, 1976), 2 to 8, 188 to 203; Sherry Turkle, *Alone Together: Why We Expect More from Technology and Less from Each Other* (New York: Basic Books, 2011), 1 to 20, 279 to 296; Lucy Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions*, 2nd ed. (Cambridge: Cambridge University Press, 2007), 226 to 240; Thomas Byrne, “The Phenomenology of ChatGPT: A Semiotics,” *Journal of Consciousness Studies* 31, no. 3 to 4 (2024): 6 to 27.

¹¹ The account of formation draws on Gadamer’s *Bildung*, Hadot’s account of philosophy as spiritual exercise, and MacIntyre’s distinction between internal and external goods. The point is that AI can provide external goods of a practice without automatically producing the internal goods that require participation, discipline, and time. See Gadamer, *Truth and Method*, 8 to 17; Pierre Hadot, *Philosophy as a Way of Life: Spiritual Exercises from Socrates to Foucault*, ed. Arnold I. Davidson, trans. Michael Chase (Oxford: Blackwell, 1995), 81 to 126; Alasdair MacIntyre, *After Virtue: A Study in Moral Theory*, 3rd ed. (Notre Dame, IN: University of Notre Dame Press, 2007), 187 to 191.

¹² The phrase “voice from nowhere” extends Donna Haraway’s critique of the “god trick” to generated synthesis. AI’s interface-effect is not true neutrality, but concealed situatedness: corpora, labor, infrastructure, capital, and language communities disappear behind the answer. See Donna Haraway, “Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective,” *Feminist Studies* 14, no. 3 (1988): 575 to 599; Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021), 1 to 40, 151 to 181; Nick Couldry and Ulises A. Mejias, *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism* (Stanford, CA: Stanford University Press, 2019), 336 to 349; Yuk Hui, *The Question Concerning Technology in China: An Essay in Cosmotechnics* (Falmouth: Urbanomic, 2016), 14 to 28.

¹³ *Gelassenheit* is used here in Heidegger’s later sense of releasement toward technological things and openness to what calculation conceals. I follow Bret Davis in reading releasement not as passivity but as a transformation of the will’s relation to technological revealing. See Martin Heidegger, *Discourse on Thinking*, trans. John M. Anderson and E. Hans Freund (New York: Harper & Row, 1966), 43 to 57; Bret W. Davis, *Heidegger and the Will: On the Way to Gelassenheit* (Evanston, IL: Northwestern University Press, 2007), xxiv to xxix, 241 to 275.

INTRODUCTION-LEVEL BIBLIOGRAPHY

Apostolopoulos, Dimitris. *Merleau-Ponty’s Phenomenology of Language*. Lanham, MD: Rowman & Littlefield, 2019.

Benjamin, Walter. “The Task of the Translator.” In *Illuminations*, translated by Harry Zohn, edited by Hannah Arendt, 69 to 82. New York: Schocken Books, 1968.

Blattner, William D. *Heidegger’s Being and Time: A Reader’s Guide*. London: Continuum, 2006.

Byrne, Thomas. “The Phenomenology of ChatGPT: A Semiotics.” *Journal of Consciousness Studies* 31, no. 3 to 4 (2024): 6 to 27.

Carman, Taylor. *Heidegger’s Analytic: Interpretation, Discourse, and Authenticity in Being and Time*. Cambridge: Cambridge University Press, 2003.

Couldry, Nick, and Ulises A. Mejias. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press, 2019.

Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.

Dahlstrom, Daniel O. “Heidegger’s Ontological Analysis of Language.” In *Heidegger’s Being and Time: Critical Essays*, edited by Richard Polt, 153 to 172. Lanham, MD: Rowman & Littlefield, 2005.

- Davis, Bret W. *Heidegger and the Will: On the Way to Gelassenheit*. Evanston, IL: Northwestern University Press, 2007.
- Dreyfus, Hubert L. *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I*. Cambridge, MA: MIT Press, 1991.
- Firth, J. R. "A Synopsis of Linguistic Theory, 1930 to 1955." In *Studies in Linguistic Analysis*, 1 to 32. Oxford: Blackwell, 1957.
- Gadamer, Hans-Georg. *Truth and Method*. 2nd rev. ed. Translated by Joel Weinsheimer and Donald G. Marshall. New York: Continuum, 2004.
- Hadot, Pierre. *Philosophy as a Way of Life: Spiritual Exercises from Socrates to Foucault*. Edited by Arnold I. Davidson. Translated by Michael Chase. Oxford: Blackwell, 1995.
- Haraway, Donna. "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective." *Feminist Studies* 14, no. 3 (1988): 575 to 599.
- Harris, Zellig. "Distributional Structure." *Word* 10, nos. 2 to 3 (1954): 146 to 162.
- Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.
- Heidegger, Martin. *Discourse on Thinking*. Translated by John M. Anderson and E. Hans Freund. New York: Harper & Row, 1966.
- Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.
- Hui, Yuk. *The Question Concerning Technology in China: An Essay in Cosmotechnics*. Falmouth: Urbanomic, 2016.
- Lenci, Alessandro. "Distributional Models of Word Meaning." *Annual Review of Linguistics* 4 (2018): 151 to 171.
- MacIntyre, Alasdair. *After Virtue: A Study in Moral Theory*. 3rd ed. Notre Dame, IN: University of Notre Dame Press, 2007.
- Marx, Karl. *Capital: A Critique of Political Economy*. Vol. 1. Translated by Ben Fowkes. London: Penguin, 1990.
- Merleau-Ponty, Maurice. *Phenomenology of Perception*. Translated by Donald A. Landes. London: Routledge, 2012.
- Merleau-Ponty, Maurice. *The Visible and the Invisible*. Translated by Alphonso Lingis. Evanston, IL: Northwestern University Press, 1968.

Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. "Efficient Estimation of Word Representations in Vector Space." arXiv:1301.3781. 2013.

Pasquinelli, Matteo. *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso, 2023.

Pennington, Jeffrey, 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*, 1532 to 1543. 2014.

Postone, Moishe. *Time, Labor, and Social Domination: A Reinterpretation of Marx's Critical Theory*. Cambridge: Cambridge University Press, 1993.

Ricoeur, Paul. *From Text to Action: Essays in Hermeneutics, II*. Translated by Kathleen Blamey and John B. Thompson. Evanston, IL: Northwestern University Press, 1991.

Taylor, Charles. *The Language Animal: The Full Shape of the Human Linguistic Capacity*. Cambridge, MA: Belknap Press of Harvard University Press, 2016.

Turkle, Sherry. *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books, 2011.

Warburg, Aby. *The Renewal of Pagan Antiquity: Contributions to the Cultural History of the European Renaissance*. Translated by David Britt. Los Angeles: Getty Research Institute, 1999.

Weizenbaum, Joseph. *Computer Power and Human Reason: From Judgment to Calculation*. San Francisco: W. H. Freeman, 1976.

Wrathall, Mark A. *Heidegger and Unconcealment: Truth, Language, and History*. Cambridge: Cambridge University Press, 2011.

Chapter 1. World, Significance, and the Birth of Language

Against the Designative Paradigm

The first task is to resist the designative image of language.

According to this image, language is a system of signs used to encode, transmit, and decode information about a world already given as a neutral field of objects. Words come first as discrete units; meanings are attached to them; the world is then represented through their combination. This is not only a philosophical assumption; it is also the tacit ontology of much computational language processing.

A phenomenological account begins elsewhere. It does not begin with the word; it begins with world.

For Heidegger, *Dasein* does not first encounter isolated objects and then add meaning to them. It is always already involved in a world. The world is not a container of things but a referential totality in which things show up as usable, threatening, promising, familiar, strange, urgent, or irrelevant, because they already matter within a world. This totality is *Bedeutsamkeit*, significance. Language is not the origin of significance; it is its articulation and expression.¹⁴

The order is therefore not:

word → meaning → world

but:

world → significance → discourse → word

Generative AI reverses this order. It does not begin from worldhood, practical involvement, embodiment, care, or historical belonging. It begins with language after language has already been written, archived, extracted, tokenized, and vectorized. It begins not with *Bedeutsamkeit* but with corpus.

This reversal governs the whole argument of the book. The large language model does not reproduce the genesis of language; it begins from the end product. It operates on language after a world has already mattered, after discourse has already articulated that world, after words have already accrued to significations, after writing has already detached those words from their originating situation.

Human language emerges from world. AI output emerges from the processed residue of language. This is the difference that must be established before the later claims about vectorization, *Gerede*, fetish, hallucination, and reappropriation can be defended.

1. The designative error

The designative view of language treats words as marks that stand for things, properties, or ideas. On this view, the world is already there as a set of identifiable objects; language comes later to name, classify, and communicate what has already been given. A word designates, a sentence combines designations, and a theory of language becomes a theory of correct mapping.

This view is ancient in its simplicity and modern in its technical destiny. It is ancient because it tempts thought whenever one imagines the world as a collection of objects and language as labels placed upon them. It is modern because computational systems can operationalize it with extraordinary power: once language is taken as a sequence of manipulable signs, it can be tokenized, embedded, and generated.

Charles Taylor gives this tradition a useful name: the designative paradigm. In his reconstruction, the Hobbes-Locke-Condillac line treats language primarily as an instrument for encoding and transmitting ideas already formed independently of linguistic expression. Against this, the Herder-Hamann-Humboldt line treats language as constitutive of distinctively human meaning, not as an external instrument applied to a pre-formed mental life.¹⁵

The point here is not that computational linguistics explicitly endorses early modern empiricism; the point is structural. A language model treats language as formal sequence. It learns relations among tokens from already-written texts. It operates as if the intelligibility of language could be reconstructed from internal relations among linguistic forms.

A phenomenological account denies that this is the origin of language. Words are not the first site of significance. They are a later articulation of a world that already matters. To understand human language one must begin with the way things show up within involvement, concern, and use. A word can name a hammer only because a world of hammering, building, repairing, sheltering, and dwelling has already been opened.

Language does not first produce world. World makes language possible.

2. World before object

The world is not first encountered as a neutral field of objects. This is Heidegger's decisive break with the theoretical attitude.

In ordinary life, things do not first appear as present-at-hand objects with measurable properties; they appear in use, in practice, in relevance. A hammer is not initially an object made of wood and metal. It is something with which one hammers, belonging to nails and boards, to building and repair, to shelter and dwelling. It withdraws into the work it enables.

This is *Zuhandenheit*, readiness-to-hand.¹⁶

When the hammer functions, I do not stare at it as an object. I hammer. The hammer does not disappear; it withdraws into the practical whole. Its meaning is not a property added to it by a mind but its place in a referential totality. The hammer refers to nails, nails to boards, boards to building, building to dwelling, dwelling to forms of life.

This referential structure is prior to theoretical description. Only when the hammer breaks, goes missing, becomes too heavy, or interrupts the work does it stand out as an object. Breakdown discloses the network that smooth use concealed. The ready-to-hand becomes conspicuous, then present-at-hand. Theory arises from a disturbance in involvement.

This matters for language because naming is not primordial. The word "hammer" does not attach significance to a mute object; it enters an already significant context. It can name the hammer because the hammer already matters in a practical world.

Meaning is not first atomic but relational. The word is born into a world.

3. Sign and reference

Heidegger's analysis of signs clarifies this point further. A sign does not first exist as an arbitrary mark detached from world; it indicates within a practical whole. It points, directs, warns, guides, or discloses because there is already a world in which such pointing can matter.¹⁷

A road sign is not merely a colored shape; it belongs to roads, rules, destinations, and habits of movement. A gesture is not merely motion; it belongs to a situation in which another can receive it. A word is not merely sound or inscription; it belongs to a shared world in which something can be said about something to someone.

A sign without world is only a mark awaiting interpretation.

This distinction will become important when we turn to tokens. A token in a language model is a unit prepared for computation: it can be embedded, weighted, attended to, predicted. It is powerful precisely because it can be detached from any single concrete use and placed into a calculable architecture. A Heideggerian sign, by contrast, indicates within a world. It is not only related to other signs; it is involved in practical significance.

The token is a technical unit; the sign is a worldly indication. The difference is not merely semantic but ontological. In the world of *Dasein*, reference belongs to involvement; in the model, relation belongs to computation. The later conversion of reference into proximity begins here: what once indicated within a world becomes operable as a relation among units.

4. *Bedeutsamkeit*

The totality of these references is what Heidegger calls *Bedeutsamkeit*, significance.

This term must not be reduced to semantic meaning. It does not name the meaning of a word, the content of a proposition, or an idea in the mind. It names the relational structure of worldhood itself. Significance is the way a world hangs together before explicit statement.¹⁸

Something is for this, toward that, useful here, dangerous there, appropriate now, forbidden in this context, promising in that one. These relations do not float in mental space; they belong to the way *Dasein* is already involved in a world. A promise can be understood only within a world of time, trust, obligation, betrayal, and expectation. A command can be understood only within a world of authority, task, obedience, refusal, and consequence. A tool can be understood only within a world of work. A question can be understood only because something is uncertain and matters.

Significance is therefore practical, historical, and existential: practical because bound to use and involvement, historical because the world into which one is thrown is already interpreted, existential because things matter only for a being whose own being is at issue.

Here the relation to *Sorge*, care, begins to appear. The world is meaningful because *Dasein* is not a detached spectator. It is concerned. It acts, fears, hopes, inherits, projects, avoids, decides, and takes up possibilities. Something matters because it belongs to a field of care.

This is why an account of meaning that begins with signs begins too late. A sign can function only because significance has already opened a world in which it can count as a sign.

5. *Significance and normativity*

Significance is not only practical orientation; it is also normative.

A world is not merely a field of affordances. It is a field in which things count as appropriate or inappropriate, true or false, shameful or demanded, serious or trivial, responsible or negligent. To inhabit a world is to be answerable to what matters there.

This is important because a purely pragmatist reading of Heidegger can make worldhood sound like skillful coping alone. Skillful coping is essential, but it does not exhaust *Dasein*. Human beings are not merely animals navigating affordances; they are beings capable of being called to account for what they do and say. They can fail to own their possibilities, stand behind a word or evade it, ask whether a way of speaking is theirs.¹⁹

This normative dimension will later matter for AI. The model can generate language about obligation, responsibility, promise, shame, death, care, or truth. It can place such words in plausible contexts and even produce impressive arguments about them. But it does not stand within a world in which these words bind it. It cannot be answerable for a promise, ashamed of a falsehood, or able to own a saying. It can produce normative vocabulary without normative exposure.

This is not a small defect. It is the difference between language as owned articulation and language as generated surface.

6. *Rede before Sprache*

The next step is Heidegger's distinction between *Rede* and *Sprache*.

Rede is usually translated as discourse. It should not be confused with chatter, verbal communication, or spoken language in the ordinary sense. *Rede* is the articulation of intelligibility, the existential structure through which the already meaningful world is articulated as understandable. *Sprache*, language in the ordinary sense, is the worldly expression of this articulation.²⁰

This is one of Heidegger's most decisive claims for the present project. Language is not the foundation of discourse; discourse is the existential-ontological foundation of language. The world is already articulated before words are uttered. Its relations, saliences, roles, and possibilities have already been opened through being-in-the-world.

The famous formula is therefore crucial: to significations, words accrue. Words do not first create significations; they come to them. They attach themselves to a world already structured by meaningful relations. Language gives expression to an intelligibility that precedes it.

This does not mean that language is secondary in the trivial sense. Once language exists, it reshapes the world it expresses. It can disclose, conceal, stabilize, transmit, distort, or transform. But in the analytic of *Being and Time*, the genesis is clear: worlded intelligibility precedes the linguistic sign.

The model begins after this genesis. It receives words after they have already accrued to significations, written language after discourse has already become expression. It does not participate in *Rede* as the articulation of being-in-the-world; it processes *Sprache* after *Sprache* has become inscription. Human language is expression from intelligibility; AI language is reconstruction from expression.

7. *Silence and the limit of generation*

The priority of *Rede* becomes clearest in silence.

Silence is not the mere absence of sound. One can be silent only within a meaningful situation. Silence can refuse, consent, mourn, accuse, protect, wait, or listen. It can articulate more than speech because it belongs to the structure of discourse. It is not outside *Rede* but one of its possibilities.²¹

This matters for AI because the model is structurally ordered toward continuation. It generates. It proceeds from prompt to output, from output to revision, from question to answer-like sequence. It can produce text about silence, recommend silence, simulate the rhetoric of restraint. But it cannot be silent from within a world. It cannot remain silent because a situation demands silence. Its non-output is not a mode of discourse; it is a system state, a refusal protocol, an error, a safety boundary, or a missing prompt. It is not an existential silence.

This distinction will later matter for *Gerede*. Idle talk avoids silence because silence exposes not-knowing. AI intensifies this avoidance by making continuation always available. But the root is already visible here: a being with *Rede* can speak or keep silent from within a meaningful world; a model can only generate or not generate.

The difference is the world.

8. *Language as embodied expression*

The genesis of language from worldhood is not only Heideggerian; it is also supported by the broader phenomenological tradition.

Merleau-Ponty gives the point a bodily register. Language is not a code imposed on a mute world by a detached mind. Speech is an expressive gesture of the body, arising from the intertwining of perception, movement, world, and others. The word is not first a sign attached to an inner idea; it is a bodily way of bringing a meaning to expression.²²

This does not replace Heidegger's analysis; it deepens it. The world that matters is encountered by a body that moves and sees, touches and suffers, gestures and responds, not by a disembodied intellect. Language grows from this bodily exposure to the world.

A child does not acquire language by mapping tokens to objects in a neutral space. A child enters a form of life. The child hears tone and rhythm, receives correction and gesture, encounters refusal and encouragement, registers laughter and shame, learns shared attention. Words come with bodies; they come with situations.

This is why the language of AI is derivative in another sense. It does not rise from flesh, gesture, perception, or bodily vulnerability. It does not speak from breath, fatigue, fear, desire, or touch. It has no bodily stake in the world it describes. It has the trace of embodied language without embodiment.

The model can say “pain,” “home,” “promise.” It cannot be in pain, miss home, or risk betrayal. Its outputs become meaningful for us because we bring the world and the body to them; the model does not generate from that world or body. It generates from traces left by those who did.

9. Public language and inherited intelligibility

The world from which language emerges is never private.

Dasein is always already being-with others. Its language is inherited and its meanings are public before they are personally appropriated. One is born into a language already filled with distinctions and habits, idioms and prohibitions, jokes and wounds, rituals and expectations. One learns what counts as childish or vulgar, professional or excessive, normal or shameful, and one learns this not by consulting definitions alone but by participating in a world.

This publicness will become the theme of the next chapter, where *das Man* is treated directly. But Chapter 1 must already prepare it. Language is born from world, and that world is shared. The public world makes language possible; it also makes idle language possible. Because language is inherited, I can speak before I have understood. Because the world is already interpreted, I can repeat what one says. Because public meanings are available, I can appear to know.

This is the opening through which AI will later enter. AI feeds on inherited public language after it has been written, archived, extracted, and made available. It does not create publicness; it processes publicness in its written sediment. The model speaks from no world of its own, but it is trained on the remains of worlds already made public by others.

10. LLMs and worldless competence

The strongest argument against AI is not that it fails at language. It often succeeds.

It produces fluent, context-sensitive, useful, and sometimes illuminating language. It can summarize and translate, compare and explain, draft and imitate, classify and generate. This competence must be taken seriously; a critique that ignores it will fail. The question is what kind of competence it is.

The LLM can generate language because human worlds have sedimented themselves in writing. It learns from these sediments. It models patterns of use, association, genre, style, and continuation. Those patterns are not arbitrary; they are traces of human being-in-the-world.

This is why the model works.

But it does not inhabit the pre-linguistic referential totality from which those traces arose. It does not encounter tools as ready-to-hand. It does not care about what it says, does not undergo breakdown, does not inherit a tradition as a finite self. It does not listen in the existential sense, does not remain silent from within a situation, does not stand in the normative exposure of saying something as its own. It has linguistic competence without worldhood.²³

This is not the same as saying it has no meaning for us. Its outputs can become meaningful when we take them up: they can orient, provoke, clarify, assist. But their meaningfulness depends on reintroduction into a human world. The model supplies generated form; the user supplies world.

This will later become the structure of user vivification. The basis is already here: AI begins from language after language has lost the worlded conditions of its birth.

11. Form, meaning, and the derivative status of AI language

The analytic literature on language models gives a useful parallel to the phenomenological claim. Bender and Koller argue that systems trained only on linguistic form do not thereby learn meaning, because meaning requires grounding in communicative intent and world-involving use.²⁴ Their argument is not Heideggerian, but it supports the same boundary from another direction.

The phenomenological version of the claim is this: written linguistic form is the residue of worlded significance. A model trained on that form can learn patterns left by meaning, but it does not thereby inhabit the world in which meaning first arose.

This is why AI language is derivative in a double sense. Human language is derivative from worlded significance: an existential derivation, moving from world to word. AI output is derivative from human language after it has become written residue: a technical derivation, moving from word to vector to output. This double derivation is the foundation of the whole monograph.

AI does not originate language from worldhood. It reconstructs the appearance of worldliness from the processed remains of written language.

12. Preliminary thesis

Significance is not a property added to neutral objects by linguistic naming. It is the relational structure of worldhood, articulated by *Rede* before words accrue to it, supported by embodied expression, and structured by normative answerability. AI begins after this whole genesis, on the written residue.

The guiding claim follows:

Language is born from worldly significance; AI is trained on the written residue of that birth.

This does not make AI useless. It makes it derivative. It can generate language because human beings have already lived and spoken, written and argued, suffered and promised, named and read and interpreted. Their worlds have left traces in language. The model learns these traces, but it does not enter the world that made them meaningful.

The chapter prepares the next move: the public structure of inherited language, and the figure of *das Man*.

FOOTNOTES

¹⁴ Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §§14 to 18, §34. The claim that language articulates rather than originates significance depends on Heidegger's movement from worldhood to discourse: entities show up within a referential totality before they are thematized in words. See also Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I* (Cambridge, MA: MIT Press, 1991), 60 to 101; Taylor Carman, *Heidegger's Analytic: Interpretation, Discourse, and Authenticity in Being and Time* (Cambridge: Cambridge University Press, 2003), chap. 3; Mark A. Wrathall, *Heidegger and Unconcealment: Truth, Language, and History* (Cambridge: Cambridge University Press, 2011), chaps. 1 to 2. The chapter follows these readings in treating significance as a worldly structure rather than a semantic content, but narrows the point toward the later claim that AI begins with language only after this worldly genesis has already occurred.

¹⁵ Charles Taylor, *The Language Animal: The Full Shape of the Human Linguistic Capacity* (Cambridge, MA: Belknap Press of Harvard University Press, 2016), chaps. 1 to 3. Taylor's distinction between the designative and constitutive-expressive pictures of language supplies the chapter's main contrast. The chapter does not claim that computational linguistics simply repeats Hobbes, Locke, or Condillac. It argues instead that the technical treatment of language as formal sequence gives executable force to a designative tendency: language becomes detachable mark, manipulable unit, and transmissible form. Taylor's account is therefore used as a philosophical grammar for the ontological reversal described here.

- ¹⁶ Heidegger's account of readiness-to-hand is developed in *Being and Time*, §§15 to 16. Dreyfus gives the canonical practical reading of the hammer example in *Being-in-the-World*, 60 to 88; see also Carman, *Heidegger's Analytic*, chap. 3; William D. Blattner, *Heidegger's Being and Time: A Reader's Guide* (London: Continuum, 2006), chap. 3. The chapter uses this equipmental analysis to establish the priority of practical involvement over detached object-representation. The word "hammer" can name an entity because hammering, building, repairing, and dwelling already form a referential context in which that entity matters.
- ¹⁷ Heidegger's analysis of signs appears in *Being and Time*, §17, where signs are understood through indication, involvement, and the referential totality of worldhood. For concise scholarly orientation, see Daniel O. Dahlstrom, *The Heidegger Dictionary* (London: Bloomsbury, 2013), entries on "sign" and "world"; Mark A. Wrathall, ed., *The Cambridge Heidegger Lexicon* (Cambridge: Cambridge University Press, 2021), entries on "reference" and "worldhood." The chapter's distinction between sign and token rests on this point. A sign indicates within a world. A token is prepared for computation. The later conversion of reference into vector proximity begins with this loss of worldly indication.
- ¹⁸ Heidegger defines significance as the relational structure of worldhood in *Being and Time*, §18. See Dreyfus, *Being-in-the-World*, 89 to 101; Blattner, *Heidegger's Being and Time*, chap. 3; Wrathall, *Heidegger and Unconcealment*, chaps. 1 to 2. The chapter follows these readings against any reduction of *Bedeutsamkeit* to lexical meaning or propositional content. Significance is not what a sign means in isolation. It is the practical, historical, and existential articulatedness within which signs, words, and propositions can first become intelligible.
- ¹⁹ The normative reading of significance is indebted especially to Steven Crowell, *Normativity and Phenomenology in Husserl and Heidegger* (Cambridge: Cambridge University Press, 2013), chaps. 6 to 8; and John Haugeland, *Dasein Disclosed*, ed. Joseph Rouse (Cambridge, MA: Harvard University Press, 2013), especially the essays on truth, finitude, and existential commitment. The chapter follows Crowell in treating phenomenological meaning as bound to answerability, not only to skillful coping. This matters for AI because generated language can contain normative vocabulary without standing in the exposure that makes promises, claims, refusals, and avowals binding.
- ²⁰ Heidegger's formulation that discourse is the existential-ontological foundation of language appears in *Being and Time*, §34. See also Daniel O. Dahlstrom, "Heidegger's Ontological Analysis of Language," in *Heidegger's Being and Time: Critical Essays*, ed. Richard Polt (Lanham, MD: Rowman & Littlefield, 2005), 153 to 172; Dahlstrom, *The Heidegger Dictionary*, entry on "discourse"; Wrathall, ed., *The Cambridge Heidegger Lexicon*, entry on "discourse." The chapter uses *Rede* to name articulation of intelligibility, not verbal performance. This is why AI can process *Sprache* without participating in the existential structure from which language arises.

²¹ Heidegger treats hearing and keeping silent as possibilities of discourse in *Being and Time*, §34. Silence is therefore not an external supplement to language, but a mode of intelligible comportment within a shared situation. Dahlstrom’s account of Heidegger’s language analysis is useful here because it situates silence within disclosedness rather than psychology. See Dahlstrom, “Heidegger’s Ontological Analysis of Language,” 153 to 172. The chapter uses this point to mark the limit of generation: a model’s non-output may be a system state or refusal condition, but it is not existential silence.

²² Maurice Merleau-Ponty, *Phenomenology of Perception*, trans. Donald A. Landes (London: Routledge, 2012), part 1, chap. 6; Maurice Merleau-Ponty, *The Visible and the Invisible*, trans. Alphonso Lingis (Evanston, IL: Northwestern University Press, 1968); Dimitris Apostolopoulos, *Merleau-Ponty’s Phenomenology of Language* (Lanham, MD: Rowman & Littlefield, 2019). Merleau-Ponty is used here as a supplement to Heidegger’s account of worldhood. Speech arises from bodily exposure, gesture, perception, and shared situation. AI output can bear traces of embodied language because it is trained on human writing, but it does not generate from flesh, vulnerability, or bodily comportment.

²³ Hubert L. Dreyfus, *What Computers Still Can’t Do: A Critique of Artificial Reason* (Cambridge, MA: MIT Press, 1992); Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, MA: MIT Press, 2019). The chapter does not use Dreyfus as a prediction that AI systems cannot produce impressive linguistic performance. It uses the deeper Heideggerian boundary in his critique: formal operations do not amount to world-involved understanding. Smith’s distinction between reckoning and judgment gives the same boundary a contemporary formulation. LLM success sharpens the point because it shows how much linguistic form can be reconstructed from the sediment of human practices without entering those practices.

²⁴ Emily M. Bender and Alexander Koller, “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Stroudsburg, PA: Association for Computational Linguistics, 2020), 5185 to 5198. Bender and Koller’s distinction between form and meaning is not phenomenological, since it centers communicative intent and grounding rather than worldhood, care, and *Rede*. The chapter uses their argument as a parallel boundary from within computational linguistics: training on linguistic form alone does not establish participation in the world from which meaning arises.

CHAPTER-LEVEL BIBLIOGRAPHY

Apostolopoulos, Dimitris. *Merleau-Ponty’s Phenomenology of Language*. Lanham, MD: Rowman & Littlefield, 2019.

Bender, Emily M., and Alexander Koller. "Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data." In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185 to 5198. Stroudsburg, PA: Association for Computational Linguistics, 2020.

Blattner, William D. *Heidegger's Being and Time: A Reader's Guide*. London: Continuum, 2006.

Carman, Taylor. *Heidegger's Analytic: Interpretation, Discourse, and Authenticity in Being and Time*. Cambridge: Cambridge University Press, 2003.

Crowell, Steven. *Normativity and Phenomenology in Husserl and Heidegger*. Cambridge: Cambridge University Press, 2013.

Dahlstrom, Daniel O. *The Heidegger Dictionary*. London: Bloomsbury, 2013.

Dreyfus, Hubert L. *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I*. Cambridge, MA: MIT Press, 1991.

Dreyfus, Hubert L. *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press, 1992.

Haugeland, John. *Dasein Disclosed*. Edited by Joseph Rouse. Cambridge, MA: Harvard University Press, 2013.

Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.

Merleau-Ponty, Maurice. *Phenomenology of Perception*. Translated by Donald A. Landes. London: Routledge, 2012.

Merleau-Ponty, Maurice. *The Visible and the Invisible*. Translated by Alphonso Lingis. Evanston, IL: Northwestern University Press, 1968.

Smith, Brian Cantwell. *The Promise of Artificial Intelligence: Reckoning and Judgment*. Cambridge, MA: MIT Press, 2019.

Taylor, Charles. *The Language Animal: The Full Shape of the Human Linguistic Capacity*. Cambridge, MA: Belknap Press of Harvard University Press, 2016.

Wrathall, Mark A. *Heidegger and Unconcealment: Truth, Language, and History*. Cambridge: Cambridge University Press, 2011.

Wrathall, Mark A., ed. *The Cambridge Heidegger Lexicon*. Cambridge: Cambridge University Press, 2021.

Chapter 2. The They and the Public Articulation of Meaning

Against the Sociological Reduction of Publicness

Chapter 1 argued that language is born from worldly significance. Words do not first create meaning; they accrue to significations already opened by being-in-the-world. Human speech emerges from a referential totality of involvement, care, embodiment, discourse, and normativity.

But the world from which language emerges is never private.

Dasein does not first exist as an isolated subject that later enters a social order. It is always already being-with others. It inherits a language before it speaks as its own. It inherits possibilities before it chooses them. It inherits standards of relevance, intelligibility, propriety, value, shame, success, danger, seriousness, and normality before it can ask whether these standards are truly its own.

This inherited public world is the domain of *das Man*, the They.

The task of this chapter is to establish *das Man* as the public articulation of everyday meaning. The point is not to treat the They as a crowd, a mass, a platform, a market, or a statistical average. These are ontic phenomena: they may express, intensify, or distort publicness, but they are not *das Man* in Heidegger's sense. *Das Man* is an existential structure of everyday *Dasein*: the anonymous mode in which the world is already interpreted for us before we have appropriated that interpretation as our own.²⁵

This distinction matters for the whole book. AI must not be identified with *das Man*. AI is not the They, not an inauthentic *Dasein*, not a new anonymous subject, not the public itself. It is an intraworldly technical entity that operates on the written sediment of language already produced within the public world of the They.

The chapter has two movements. First, it reconstructs *das Man* as an existential structure of public intelligibility. Second, it prepares the book's later claim that AI is not *das Man* but the technical afterlife of its written language.

1. *The public world after the birth of language*

Chapter 1 ended with the claim that language is born from worldly significance. This claim must now be socialized. The world into which language is born is shared. The referential totality of significance is not mine alone; it is articulated in shared practices, roles, inherited idioms, public vocabularies, and ways of knowing what matters.

A child enters language through others. A worker enters a profession through others. A citizen enters public meaning through others. A reader enters a tradition through others. Even the solitary thinker has already inherited words, distinctions, problems, and standards of seriousness from a shared world.

This means that language is public before it is owned.

One does not first possess a private language and then decide to communicate. One first finds oneself in a world where things are already speakable: where there are already ways to greet and apologize, to thank and accuse, to confess and remain silent, and already ways to say what counts as clever or vulgar, professional or excessive, responsible or ambitious.

The public world gives language its first availability and its first danger. Because language is inherited, I can speak before I have understood. Because meanings are public, I can repeat what is said. Because the world is already interpreted, I can move within it without asking whether its interpretation is mine.

The They is this anonymous availability of the already-interpreted world.

2. *The “who” of everydayness*

Heidegger introduces *das Man* in response to the question of the “who” of everyday *Dasein*. Who is *Dasein* in its ordinary way of being? The answer is not, at first, the isolated individual; nor is it an empirical group. The everyday self is the they-self, the *Man-selbst*.²⁶

In ordinary life, I usually do what one does. One studies, one works, one relaxes, one has an opinion. One is outraged by this and indifferent to that. One knows what counts as success and what is embarrassing, and one knows how to speak.

This “one” is not a person, nor a collective subject standing above individuals. It is the anonymous public mode in which everyday life is already intelligible.

The They does not need to command or to issue rules in explicit form. It governs by prior familiarity. It makes the world appear as already known. The self does not experience this as domination; it experiences it as normality.

This is why the They is so difficult to see. If a tyrant commands, one can identify the tyrant; if a law commands, one can read the law; if a doctrine commands, one can name the doctrine. But the They governs by anonymity. It is “no one,” and precisely for that reason it can be everywhere.

The question “who says so?” often receives no answer. People say so. The market says so. Experts say so. It is just how things are. The They is the public grammar of “just how things are.”

3. Against the reduction of das Man to conformity

The first interpretive danger is to reduce *das Man* to psychological conformism. On this reading, Heidegger would simply be saying that people follow the crowd, fear standing out, and surrender their individuality to mass opinion. This is not false as an ontic description, but it is not yet the ontological point.

The They is not merely a social pressure that sometimes forces individuals to conform. It is a condition of everyday intelligibility. Without inherited public meanings, there would be no ordinary shared world in which anything could first make sense. The fact that one knows how to use a door, make a promise, write an email, or recognize a joke already presupposes a public world of intelligibility.²⁷

This is why the Dreyfus-Olafson debate remains useful. Dreyfus rightly insists that the Anyone provides the shared background practices through which the world shows up as meaningful. Olafson rightly worries that such a reading can domesticate Heidegger’s sharper critique of conformism, making the They sound merely like the neutral background of social norms. The chapter’s position is different from both when taken in isolation. *Das Man* is a positive condition of everyday intelligibility, but it is also a structure of leveling and disburdening.²⁸ The They makes the world accessible and also average; it allows ordinary understanding and also non-appropriation.

This is why inauthenticity must not be moralized. Everydayness is not a defect one can simply escape; it is the starting point from which any authentic appropriation must begin.

4. Publicness as prior interpretation

Publicness does not merely circulate opinions; it articulates what counts as intelligible.

In the public world, things already have names, uses, reputations, and expected meanings. A job title, a courtroom, a religious symbol, a family role, a political slogan, an insult or a compliment: all appear within public interpretations that precede any individual act of thought.

This public interpretation is not always explicit; it is often absorbed as obviousness. The They does not need to say “this is what counts as normal.” It lets normality function.

This matters because *das Man* does not only shape what people believe. It shapes the field in which something can be recognized as believable or unintelligible, serious or excessive, interesting or irrelevant. Publicness sets the first horizon of the sayable. It does not decide every meaning, prevent every act of originality, or eliminate conflict. But it gives the field in which conflict first appears as conflict and originality first appears as originality.

A radical claim is radical only against a public horizon. An original gesture is original only against an inherited world. A decision can become one’s own only because it first comes in forms already public enough to be taken up.

The They is therefore not external to meaning. It is the public articulation through which meaning first circulates.

5. *Distantiality*

Heidegger describes *das Man* through three structural moments: *Abständigkeit*, *Durchschnittlichkeit*, and *Einebnung*. The first is distantiality.

Distantiality names the way everyday *Dasein* is concerned with its distance from others. One is ahead, behind, lacking, exceptional, late, impressive, embarrassing. One’s self-understanding is quietly measured through a public relation to others.²⁹ This distance is not always explicit comparison; it does not need to become vanity or envy. It belongs to the way the everyday self is socially exposed: the self understands itself through a field where others matter as measures, rivals, examples, warnings, and norms.

Distantiality is the social tremor inside ordinary selfhood. It is visible in professional life, education, intellectual discourse, fashion, politics, parenting, and even in forms of rebellion. To stand out is still to stand out from a public field; to reject the average can still be a way of being governed by it; to be original can become another public demand.

Distantiality prepares the self for averageness because it places the self under the silent comparison of the public. The self asks, often without knowing that it asks: am I keeping up, am I behind, am I saying what counts, am I intelligible, am I exceptional in the right way, am I excessive in the wrong one? The They rules through such questions because the answers are rarely owned by anyone in particular.

6. *Averageness*

Distantiality gives rise to *Durchschnittlichkeit*, averageness.

Averageness does not mean stupidity. It means that the public world tends to define in advance the range within which things can appear as intelligible. It establishes what is commonly acceptable, commonly recognizable, commonly important, commonly questionable, commonly dismissible. The average is not the numerical mean of opinions; it is the public measure of intelligibility.

Averageness lets the world function. It keeps everyday life from requiring original interpretation at every step. If every word, tool, role, practice, and expectation had to be reinvented each morning, no life would be possible. We need average intelligibility.

But averageness also reduces difficulty. The difficult becomes explainable, the unsettling becomes a topic, the singular becomes a case, the exceptional becomes a recognizable type, the question becomes an answer one can circulate. Averageness does not merely simplify; it reassures. It tells *Dasein* that the matter has already been understood well enough. The world becomes navigable, but also prematurely settled.

This is the existential seed of AI's attraction. AI will later offer a technical form of averageness: a fluent public answer before the user has appropriated the question. But the desire for such an answer is older than the machine. It belongs to the everyday self's relief in the already-understood.

7. *Leveling-down*

The third structural moment is *Einebnung*, leveling-down.

Leveling-down is not the crude destruction of difference but the translation of difference into publicly manageable form. The public world does not necessarily deny the rare, the difficult, or the original; it renders them familiar and makes them speakable in advance. A discovery becomes "an innovation," a trauma becomes "an experience," a spiritual crisis becomes "a journey," a radical thought becomes "an interesting perspective," a poem becomes "content," a person becomes "a profile."

The original is not silenced; it is made consumable. This is why leveling-down is so difficult to resist. It often preserves the surface of difference. It may even celebrate difference. But it does so by bringing difference into a field where it can be compared, summarized, ranked, quoted, and circulated.

The unique is not erased; it is made manageable.

This prepares the later analysis of vectorization. A vector does not make all words identical. It preserves difference with extraordinary precision. But it translates difference into position, distance, weight, and calculable proximity. That technical operation does not equal *Einebnung* in the existential sense; yet it can be read as a later technical intensification of the public tendency to make the singular comparable.³⁰

The They levels by making the world already publicly intelligible. AI will later level by making the already-said calculable.

8. Disburdening

The They disburdens *Dasein*.

This is one of Heidegger's most important points. Everyday publicness relieves the individual from having to decide everything for itself. The They answers before I ask, interprets before I interpret, provides phrases before I find my own, offers roles before I choose. It tells me how to mourn and celebrate, how to work and desire, how to succeed and fail and speak.

This disburdening is not merely negative. It allows life to proceed. It allows shared practices. It gives access to a world one did not create. It prevents the self from being crushed by the impossible task of original interpretation at every moment.

But disburdening becomes inauthentic when it replaces responsibility. When the They says *one does this*, the self can say *then I need not decide*. When the They says *this is what one thinks*, the self need not ask. When the They says *this is what one says*, the self need not own its speech.

This is where public intelligibility becomes flight. The self is relieved from the burden of its own saying.

AI will later intensify this relief. It will not merely provide public language; it will generate it on demand. But the need for relief belongs first to everyday *Dasein*. The machine succeeds because the desire already exists.

9. Depersonalization and disindividuation

The They produces two related effects: depersonalization and disindividuation.

Depersonalization concerns speech without owned authorship. In the public mode, things are said, repeated, circulated, cited, and accepted without anyone fully standing behind them. The speaker may be present, but the saying is borrowed from the anonymous. One speaks as one speaks.

Disindividuation concerns the reduction of singularity to public type. The individual is not necessarily erased; the individual becomes legible through a category: student, expert, patient, citizen, customer, user. These categories are useful. They make the social world navigable. But they also allow the person to be handled as a public type before being encountered as singular.

The two effects reinforce each other. Speech becomes less owned; persons become more classifiable.

This matters for the later analysis of AI. AI-generated output is depersonalized in an extreme sense: language without a speaker who can own the saying. AI-based representation is disindividuating because it converts traces of persons into comparable features, profiles, clusters, and scores.

But again, the distinction must be preserved. The They is not AI. Human *Dasein* can hide in the public because it has a self from which it hides. AI has no such self. It does not depersonalize speech by evading ownership; it produces speech-like output without ownership ever being possible.

This is more uncanny than ordinary public speech. It is not “one says” in the human sense. It is language with no one there.

10. Being-with and the irreducibility of social worldhood

The They presupposes *Mitsein*, being-with.

This point is necessary because Heidegger’s analysis of the They is sometimes misread as an anti-social doctrine. It is not. *Dasein* is being-with before it is absorbed into the They. The publicness of the They is a deficient or average mode of being-with, but it does not create being-with from nothing.³¹

Others are not first objects encountered in the world; they are part of the structure of worldhood. The world is always already a shared world. Its tools, languages, institutions, distances, roles, and expectations are co-inhabited.

This means that authenticity cannot mean escape from others. Authenticity is a modification of being-with, not an abolition of it. The self does not become authentic by inventing a private world. It becomes capable of taking over inherited possibilities in a way that is answerable as its own.

This has direct consequences for the book’s AI thesis. AI has no *Mitsein*. It can simulate address, answer, dialogue, and social style; it can generate the grammar of politeness, care, instruction, correction, and response. But it is not with us in the existential sense. It does not share a world in which anything is at stake for it. It does not respond as another *Dasein*. It does not acknowledge, respect, neglect, or dominate in the modes of solicitude.

The user may experience it as a conversational partner. But that experience will be analyzed later as vivification. The model does not enter the shared world as another; it is inserted into the shared world as a technical entity. The relation is asymmetrical: we are with one another, we use the machine, and then we may mistake the machine for another.

11. Inauthenticity without moralism

The analysis of *das Man* requires a careful distinction between inauthenticity and moral failure.

Inauthenticity is not cowardice, stupidity, corruption, or weakness. It is not a personal defect. It is an existential mode of everyday *Dasein*: one begins from the world as already interpreted, from inherited language, from public possibilities, as the they-self.³²

This is why the manuscript must avoid saying that AI is inauthentic. AI is not inauthentic because AI is not *Dasein*. It does not exist as a being for whom its own being is at issue. It has no mineness, no anxiety, no mortality, no care. It cannot take over or fail to take over its possibilities. It cannot be absorbed in the They because it cannot own itself outside the They.

The relevant claim is different: AI operates on the written residue of inauthentic everydayness. It processes public language after that language has been expressed, circulated, and sedimented. It can generate the form of public understanding without occupying the existential condition from which public understanding arises.

This distinction protects the argument from exaggeration. The machine is not inauthentic; the human relation to the machine can be.

12. *Gerede and the already-said*

If *das Man* is the anonymous “who” of everydayness, *Gerede* is its dominant discursive mode.

Gerede is usually translated as idle talk, but the translation can mislead. It does not mean mere chatter, gossip, or nonsense. It names the way discourse circulates publicly when the matter has not been originally appropriated. The thing is understood because one can repeat what is already said about it.³³

Idle talk can be factually correct, intelligent, even polished. Its defect is not always falsehood but groundlessness. It speaks without returning to the matter, explains without dwelling, circulates without ownership. It gives *Dasein* the possibility of understanding everything without having made the thing its own.

This is the origin of the already-said. The already-said is the sediment of public discourse: phrases and slogans, definitions and interpretations, professional vocabularies and moral formulas, political frames and academic templates. It is the language through which the public world stabilizes itself.

The already-said is not useless; it is the starting point of understanding. One cannot think without inherited language, read without pre-understanding, or ask without public terms. But the already-said becomes dangerous when it is mistaken for understanding itself.

This is why *Gerede* will become central to AI. LLMs operate on language after it has become already-said in written form. They do not generate from original encounter; they generate from sedimented public discourse.

Chapter 6 will develop this fully. For now, the point is preparatory: the They produces the already-said, *Gerede* circulates it, writing sediments it, AI later processes it.

13. Publicness and the sociological public sphere

The They must be distinguished from the sociological public sphere.

Habermas, Bourdieu, Foucault, and later media theory all offer powerful accounts of public discourse, doxa, discipline, publicity, platforms, transparency, and communicative power. These accounts matter for the later historical and political chapters. But *das Man* is not identical with any specific public sphere, media system, class structure, platform regime, or discursive formation.³⁴ Those are historical formations; *das Man* is an existential structure.

This does not make sociological theory irrelevant. It means that sociology describes determinate historical expressions of publicness, while Heidegger analyzes the existential condition that lets any public intelligibility be operative in the first place.

The digital public sphere is therefore not *das Man*; it is one intensified historical form in which *das Man* can operate. Social platforms, algorithmic feeds, recommendation systems, metrics, influencer economies, and AI interfaces can all intensify distantiality, averageness, and leveling. But they do not create publicness from nothing. They technicalize and accelerate it.

The same will be true of AI. AI does not create the They; it gives the already-said a new technical afterlife.

14. From publicness to enframing

The relation between *das Man* and modern technology should be stated carefully.

The early analysis of publicness in *Being and Time* is not the same as the later analysis of *Gestell*. The They is an existential structure of everyday *Dasein*. Enframing is a historical mode of revealing in which beings appear as orderable standing-reserve. Yet the two analyses can be read in continuity.

The They levels public meaning by making everything familiar, comparable, and already understood. Enframing levels beings by making them available, calculable, optimizable, and orderable. The first is existential; the second is onto-historical.

This continuity should not be collapsed into identity. But it matters because AI belongs to a technological epoch in which public language can become standing-reserve. The already-said is no longer only repeated in conversation, media, and institutions. It can now be scraped, stored, tokenized, vectorized, recombined, and sold back as generated speech.³⁵ The public language of the They becomes technical stock.

This is where Chapter 2 points forward to Chapters 4 and 7. The They gives the public world of already-sayable meanings. Enframing provides the historical condition under which that language can be challenged forth as resource. AI is the technical system that processes this resource.

The chain is not *AI equals das Man*. The chain is: *das Man* articulates public intelligibility, *Gerede* circulates the already-said, writing sediments it, enframing makes it standing-reserve, and AI vectorizes and reanimates it.

15. AI is not das Man

The decisive claim can now be stated. AI is not *das Man*.

The temptation to identify them is understandable. AI speaks anonymously, produces average public language, repeats what is already said, levels differences, gives the appearance of understanding without appropriation; it often sounds like no one and everyone.

But *das Man* is not a style of output. It is not statistical representativeness, not the average of public opinion, not the synthesis of a corpus. It is an existential structure of everyday being-with.

AI lacks the conditions required to be *das Man*. It has no self to disperse, no possibility to own, no anxiety from which to flee, no mortality to cover over, no care, no being-with. A human being can lose itself in the They because a human being can be itself. AI cannot lose what it never had.

The correct formulation is therefore: AI is not the They; it is an intraworldly technical entity that operates on the written sediment already produced within the public world of the They.

This is the chapter's central correction.

16. AI as recursion on public sediment

If AI is not *das Man*, what is it in relation to *das Man*? It is a recursion on public sediment.

The public world produces language. Idle talk circulates it; writing stores it; archives and platforms accumulate it; training systems extract it; tokenization and vectorization formalize it. The model generates new language from this formalized sediment.

The model therefore begins after the public world has already spoken. It does not participate in the original publicness of being-with. It does not negotiate social norms as a participant. It does not experience distantiality. It does not seek averageness as a relief from exposure. It does not level down in order to avoid the anxiety of singularity. It processes the products of these tendencies.

This is why AI-generated language can sound public, average, and anonymous: it is generated from public, average, anonymous language already available in written form.

The model is a mirror only in a limited sense. A mirror reflects what stands before it; the model is more complex. It transforms what has already been written into a latent structure of relations, then generates new outputs from that structure. It does not merely reflect the They; it recursively recombines the already-said.

This is the technical afterlife of publicness.

17. *Statistical averageness and existential averageness*

The relation between AI and averageness must also be carefully distinguished.

Heidegger's *Durchschnittlichkeit* is not a statistical concept. It is an existential structure of everyday public intelligibility. It concerns how the public world sets standards of what counts as intelligible, acceptable, important, familiar, or excessive.

AI systems, by contrast, operate through statistical and computational processes. They learn distributions, relations, and likely continuations. Their outputs reflect patterns of public language; their alignments smooth tone; their sampling settings adjust variation; their generated answers tend toward safe, average intelligibility.

But statistical averageness is not existential averageness. The relation is analogical and historical, not identical.

This matters because the argument must not claim that a loss function is literally *das Man*, that softmax is literally publicness, or that vector proximity is literally significance. Such claims would confuse computational mechanism with existential structure.

The better claim is more precise:

AI technically formalizes and intensifies certain effects of public averageness after those effects have been sedimented in written language.

This formulation preserves both sides. It respects the specificity of Heidegger's existential analytic and the specificity of machine learning. It also makes the critique sharper: the problem is not that mathematics secretly becomes *das Man*. The problem is that public language, already shaped by the They, becomes available for mathematical processing. The existential becomes sediment; the sediment becomes technical; the technical returns as public speech.

18. *The technical They as heuristic*

One might be tempted to call AI the "technical They." This phrase can be useful, but only as a heuristic. If it means that AI replaces *das Man*, it is wrong. If it means that AI functions as a technical interface through which the already-said becomes available on demand, it is useful.

In this limited sense, AI can be called a technical organ of publicness. It gives public language an interface, the already-said a response function, average intelligibility a generative form. The user no longer has to absorb public discourse slowly; the user can summon it on command. Explain this, draft a view, compare positions, tell me what one would say. AI returns the form of public intelligibility as if it were addressed personally to the user.

This is new. Traditional *Gerede* circulates; AI-generated *Gerede* responds; publicness becomes interactive.

But the response is not the speech of another *Dasein*. It is generated language whose existential status depends on the user's uptake. The user lends it the world in which it can count as an answer. That is why the next chapters must move through writing, corpus, token, vector, generated speech, and user vivification.

19. Preliminary thesis

The chapter establishes that *das Man* is not a sociological category, a platform, a statistical average, or AI itself, but an existential structure of everyday public intelligibility. AI does not instantiate this structure. It operates on the written residue of public language that *das Man* produces, technically formalizes the already-said, and reanimates it as generated speech. It becomes existentially powerful because everyday *Dasein* already desires the relief of the already-understood.

AI does not instantiate the structure of *das Man*. It recursively recombines the already-said into language that returns to the user as if addressed personally. The next chapter examines the passage from public language to writing, since publicness alone does not yet produce this technical afterlife.

FOOTNOTES

²⁵ Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §§25 to 27. Heidegger introduces *das Man* within the analysis of being-with, not as a theory of mass society. For readings that keep this point clear, see Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I* (Cambridge, MA: MIT Press, 1991), chap. 8; Taylor Carman, *Heidegger's Analytic: Interpretation, Discourse, and Authenticity in Being and Time* (Cambridge: Cambridge University Press, 2003), chap. 4; William D. Blattner, *Heidegger's Being and Time: A Reader's Guide* (London: Continuum, 2006), chap. 5. The chapter follows these commentators in treating the They as an existential structure of everyday intelligibility. It also insists that this structure has critical force: publicness makes meaning available, but it also levels, disburdens, and anonymizes.

²⁶ Heidegger's account of the they-self, the *Man-selbst*, appears in *Being and Time*, §27. The point is not that the individual is replaced by a collective subject. Rather, everyday *Dasein* usually understands itself through public possibilities it has not chosen as its own. Dreyfus translates *das Man* as the Anyone in order to stress this anonymous practical background. See Dreyfus, *Being-in-the-World*, chap. 8. Blattner's discussion is also useful because it emphasizes the "nobody" character of the They: the source of ordinary intelligibility is everywhere operative, but no determinate agent can be named as its author. See Blattner, *Heidegger's Being and Time*, chap. 5.

²⁷ The examples here extend Heidegger's claim that the world is always already a shared world in which things show up as usable, appropriate, relevant, or questionable. See Heidegger, *Being and Time*, §§25 to 27. Dreyfus's account of background practices gives the clearest practical formulation of this point; Carman's account of interpretation and discourse gives the sharper ontological one. See Dreyfus, *Being-in-the-World*, 141 to 162; Carman, *Heidegger's Analytic*, chaps. 4 to 6. The chapter follows Dreyfus in treating public norms as enabling conditions of intelligibility, but it follows Carman in refusing to reduce that intelligibility to mere social convention. Public meaning is not an optional layer added to private experience. It is the field in which everyday experience first becomes articulate.

²⁸ Hubert L. Dreyfus, *Being-in-the-World*, chap. 8; Frederick A. Olafson, "Heidegger à la Wittgenstein or 'Coping' with Professor Dreyfus," *Inquiry* 37, no. 1 (1994): 45 to 64. Dreyfus stresses the positive role of shared practices in making the world intelligible. Olafson objects that this risks softening Heidegger's analysis of inauthenticity into a benign account of norm-following. The chapter treats the debate as productive because each side names a real danger. A purely negative reading moralizes everydayness. A purely positive reading can miss the force of leveling. See also Carman, *Heidegger's Analytic*, chap. 4; Blattner, *Heidegger's Being and Time*, chap. 5. The position taken here follows Carman and Blattner in holding both moments together: *das Man* is constitutive of everyday intelligibility and, at the same time, the structure through which intelligibility becomes average and unowned.

²⁹ Heidegger analyzes *Abständigkeit*, distantiality, in *Being and Time*, §27. The German term names the everyday concern with distance from others: being ahead, behind, equal, exceptional, or deficient. This is not reducible to vanity or empirical competition. It is a structural feature of the they-self's exposure to public measures. See Carman, *Heidegger's Analytic*, chap. 4; Blattner, *Heidegger's Being and Time*, chap. 5. The chapter uses distantiality to show why public intelligibility already contains comparison. Before AI produces ranked, scored, or optimized language, everyday *Dasein* already inhabits a world in which it measures itself through what "one" does, says, expects, and understands.

³⁰ Heidegger's account of *Einebnung*, leveling-down, appears in *Being and Time*, §27. The later analysis of technological ordering is developed in 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. For the relation between Heidegger's early analysis of leveling and his later account of technological revealing, see Iain D.

Thomson, *Heidegger on Ontotheology: Technology and the Politics of Education* (Cambridge: Cambridge University Press, 2005), 44 to 77. The analogy with vectorization is limited by design. A vector does not become *das Man*, and mathematical proximity is not existential publicness. The chapter's claim is that a prior tendency of publicness, making the singular comparable, later becomes technically formalized.

³¹ Heidegger's analysis of being-with appears in *Being and Time*, §§25 to 27. For a sustained account of Heidegger's social ontology, see Irene McMullin, *Time and the Shared World: Heidegger on Social Relations* (Evanston, IL: Northwestern University Press, 2013). McMullin is especially useful because she shows that being-with is not an added relation between already complete subjects, but a temporal and practical structure of shared worldhood. The chapter follows this line in refusing an anti-social reading of authenticity. Authenticity modifies being-with; it does not abolish it. This is also why AI cannot be identified with *das Man*: it can simulate address within our shared world, but it does not share that world as another *Dasein*.

³² Heidegger's analyses of the they-self, falling, and inauthenticity appear in *Being and Time*, §§27 and 38. See also Carman, *Heidegger's Analytic*, chap. 4; Blattner, *Heidegger's Being and Time*, chap. 5; Steven Crowell, *Normativity and Phenomenology in Husserl and Heidegger* (Cambridge: Cambridge University Press, 2013), chaps. 6 to 8. These readings help prevent inauthenticity from being treated as a moral weakness. Inauthenticity is an existential mode in which *Dasein* begins from inherited possibilities and public interpretations. The chapter uses this distinction to protect the AI argument from category error. A human being can be inauthentic because it can own or fail to own its possibilities. AI cannot be inauthentic because it has no possibilities in this existential sense.

³³ Heidegger develops *Gerede*, idle talk, in *Being and Time*, §35. The term should not be narrowed to gossip or error. *Gerede* can be correct, informative, and socially useful, while still remaining groundless in Heidegger's sense because it circulates without original appropriation of the matter. See Daniel O. Dahlstrom, *Heidegger's Concept of Truth* (Cambridge: Cambridge University Press, 2001); Carman, *Heidegger's Analytic*, chap. 6. The chapter uses *Gerede* as the immediate precursor to the already-said. This does not mean that AI performs idle talk as *Dasein* does. It means that LLMs process language after public discourse has already made it repeatable, quotable, and detachable from owned encounter.

³⁴ Jürgen Habermas, *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*, trans. Thomas Burger with Frederick Lawrence (Cambridge, MA: MIT Press, 1989); Pierre Bourdieu, *Outline of a Theory of Practice*, trans. Richard Nice (Cambridge: Cambridge University Press, 1977); Michel Foucault, *The Archaeology of Knowledge and the Discourse on Language*, trans. A. M. Sheridan Smith (New York: Pantheon, 1972). These works analyze historically determinate forms of publicity, doxa, practice, discourse, and power. The chapter draws a boundary rather than a rejection. Habermas, Bourdieu, and Foucault help describe concrete public formations. Heidegger's *das Man* names the existential condition through which any such formation can matter to everyday *Dasein* as already intelligible.

³⁵ Heidegger, “The Question Concerning Technology,” 3 to 35; Thomson, *Heidegger on Ontotheology*, 44 to 77. Heidegger’s *Gestell* names a historical mode of revealing in which beings appear as orderable standing-reserve. The chapter does not claim that the analysis of *das Man* in *Being and Time* already contains the technical structure of machine learning. Its claim is sequential: public language is first articulated through being-with and the They; idle talk circulates it; writing and archives preserve it; technological revealing makes it available as resource. AI enters at the end of this sequence by processing language already made available in written form.

CHAPTER-LEVEL BIBLIOGRAPHY

Blattner, William D. *Heidegger’s Being and Time: A Reader’s Guide*. London: Continuum, 2006.

Bourdieu, Pierre. *Outline of a Theory of Practice*. Translated by Richard Nice. Cambridge: Cambridge University Press, 1977.

Carman, Taylor. *Heidegger’s Analytic: Interpretation, Discourse, and Authenticity in Being and Time*. Cambridge: Cambridge University Press, 2003.

Crowell, Steven. *Normativity and Phenomenology in Husserl and Heidegger*. Cambridge: Cambridge University Press, 2013.

Dahlstrom, Daniel O. *Heidegger’s Concept of Truth*. Cambridge: Cambridge University Press, 2001.

Dreyfus, Hubert L. *Being-in-the-World: A Commentary on Heidegger’s Being and Time, Division I*. Cambridge, MA: MIT Press, 1991.

Foucault, Michel. *The Archaeology of Knowledge and the Discourse on Language*. Translated by A. M. Sheridan Smith. New York: Pantheon, 1972.

Habermas, Jürgen. *The Structural Transformation of the Public Sphere: An Inquiry into a Category of Bourgeois Society*. Translated by Thomas Burger with Frederick Lawrence. Cambridge, MA: MIT Press, 1989.

Haugeland, John. *Dasein Disclosed*. Edited by Joseph Rouse. Cambridge, MA: Harvard University Press, 2013.

Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.

Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.

McMullin, Irene. *Time and the Shared World: Heidegger on Social Relations*. Evanston, IL: Northwestern University Press, 2013.

Olafson, Frederick A. "Heidegger à la Wittgenstein or 'Coping' with Professor Dreyfus." *Inquiry* 37, no. 1 (1994): 45 to 64.

Thomson, Iain D. *Heidegger on Ontotheology: Technology and the Politics of Education*. Cambridge: Cambridge University Press, 2005.

Chapter 3. Writing, Reading, and Historical Dialogue

Against Processing as Reading

Chapter 2 argued that *das Man* is the public articulation of everyday meaning. The They is not a machine, not a statistical average, and not a digital platform. It is the anonymous existential structure through which the world is already interpreted before *Dasein* has made that interpretation its own. Its discursive mode, *Gerede*, produces the already-said: the language that circulates publicly, reassures, explains, repeats, and makes the world available.

But AI does not train on *das Man* as an existential structure. It trains on what has become available from the public world of the already-said. For this to happen, language must undergo a further transformation. It must pass into writing, archive, corpus, token, and vector.

This chapter studies the first of these passages: the passage into writing.

The danger would be to say too quickly that writing is dead. That claim is false. Writing is not dead because it is written. A written text can remain alive across centuries when it is read historically. The author may be absent, but absence is not nullity. The text can still address, wound, disturb, console, form, and transform the reader. A past act of writing can still act upon the present.

Writing becomes deadened when it is no longer read, but only processed.

This chapter therefore develops a distinction that will govern the rest of the book:

Reading re-historicizes writing. Processing de-historicizes it.

Historical reading returns writing to a world of question, distance, tradition, responsibility, and appropriation. Computational processing prepares writing for extraction. It serializes, fragments, formalizes, and reorganizes the text so that it can become corpus, token, vector, and latent proximity.

The distinction is not between good humans and bad machines. It is between two relations to writing. A human reader can also reduce a text to dead information. A computational tool can also assist reading. The question is not who touches the text, but what happens to writing in the relation.

1. The ancient suspicion of writing

The anxiety that writing deadens speech is old. Plato's *Phaedrus* already stages the suspicion that written words appear to speak while being unable to answer. A written text says the same thing to everyone. It cannot defend itself or adjust to the questioner. It can circulate without the living intelligence of the speaker.³⁶

This suspicion remains important. It recognizes something true about writing: writing detaches speech from the immediate presence of the speaker. It lets words travel beyond the situation of their utterance. It gives language a strange independence. The written word can be copied, cited, misunderstood, preserved, and used by those never intended by the writer.

But if one stops here, one mistakes detachment for death.

Writing is not the corpse of speech; it is speech transformed into a historical trace. That transformation is ambiguous. It can reduce language to a mark. It can also preserve a saying beyond the life of the speaker. It can detach words from the moment in which they were spoken, but it can also make possible a dialogue across time.

The writer is absent, yet the writer's act remains. The reader is not in the author's world, yet the reader can be addressed by what came from that world.

Writing therefore creates distance. It does not decide in advance whether distance will become deadness or understanding. That decision belongs to reading.

2. Writing is not dead language

A written text can survive its author. It can move from one historical world to another. It can be copied and translated, taught and censored, forgotten and revived, weaponized or misunderstood. It can become canon, scripture, archive, evidence, or debris.

But survival is not yet life. A text lives when it is read.

This does not mean that the reader simply animates a passive object at will. The text has its own resistance. It comes from somewhere. It bears words, structures, concepts, silences, and questions that the reader did not invent. It is not reducible to the reader's projection.

A written text is alive when its distance is allowed to matter. To read Plato or Augustine, Dante or Nietzsche, Kafka or Heidegger, is not merely to extract information from marks. It is to enter a mediated relation with a past act of saying. The author is not present as a living interlocutor, but neither is the author simply gone. The text mediates an encounter in which the reader speaks with himself through another.

This is why reading can transform self-understanding. I do not merely receive content. I discover the poverty of my own questions, the limits of my language, the hidden assumptions of my world. The author's distance becomes a force. It prevents me from reducing the text too quickly to what I already understand.

Writing is alive when it resists assimilation.

3. Reading as historical dialogue

Historical reading is dialogical. This does not mean that reading recreates a living conversation with the author. The author cannot answer in real time, cannot clarify what I misunderstand, cannot be summoned as a contemporary speaker. The dialogue is mediated by the text.

Gadamer's hermeneutics gives this structure its clearest philosophical form. Understanding is not a method by which the reader escapes history and reconstructs the past from nowhere. The reader is historically effected: the reader belongs to a tradition, carries prejudices, and approaches the text with questions already shaped by a world. The text also comes from a history; it bears its own horizon. Reading occurs when these horizons meet, resist, and modify one another.³⁷

The text questions the reader; the reader questions the text. The dialogue is therefore not an exchange between two equal living subjects. It is an event of understanding across distance. The reader listens to the text as a historical other, yet never leaves the present from which the listening occurs. The past does not become present by being dissolved into the reader's categories; the present does not become past by pretending to occupy the author's mind.

Reading lives in the interval. This is why distance is not the enemy of understanding; distance is the medium through which understanding becomes possible.

4. Finitude and the need for dialogue

The need for hermeneutic dialogue arises from finitude.

If the reader possessed absolute knowledge, dialogue would be unnecessary. If the past could be reconstructed without remainder, interpretation would be mechanical. If the text's meaning were fully present and immediately available, reading would be extraction.

But human beings are finite. They do not think from nowhere. They inherit languages, questions, institutions, wounds, and blind spots. They read from a place; they misunderstand; they need the text because the text can expose what their own world cannot see.

This is why historical reading is not a deficiency in relation to perfect knowledge. It is the way finite beings understand.

A text from another age does not simply deliver a message. It places the reader in relation to what the reader lacks: another language, another world, another way of asking, another horizon of expectation. The difficulty of the text is not accidental but part of the text's power. The reader must not eliminate this difficulty too quickly. The reader must allow the past to remain strange long enough for the present to be questioned.

5. Distanciation and appropriation

Ricoeur sharpens this point by showing that writing produces distanciation. Once written, a text is separated from the author's immediate intention, the first audience, and the original situation of speech. The text gains autonomy. It can project a world beyond the psychological life of its author.³⁸

This autonomy is not a loss to be repaired; it is the condition of interpretation. Because the text is distanced, it can be read again. Because it is not identical with the author's living intention, it can enter new situations. Because it projects a world in front of itself, the reader can appropriate that world.

Appropriation does not mean possession. It does not mean making the text say what I already want. It means allowing the world projected by the text to enter and alter my own self-understanding.

In this sense, reading re-historicizes writing. The text was distanced by writing. Reading does not abolish that distance; it takes it up. It brings the written trace into a new historical situation while preserving the fact that it comes from elsewhere. The reader actualizes the text without consuming its otherness.

This matters for AI because computational processing also begins from distanced writing. But it does not appropriate the text in Ricoeur's sense. It does not let the projected world of the text transform a self. It does not enter the text's world; it prepares the text for operation.

Distanciation can lead to appropriation. It can also lead to extraction.

6. The text as incomplete event

The phenomenology of reading adds another layer. A written work is not simply a closed object containing a complete meaning. Ingarden, Iser, and Jauss show that the text must be actualized by the reader. It has gaps, indeterminacies, horizons of expectation, and a history of reception.³⁹ This does not mean that the reader can do anything with the text. The text guides and constrains; it sets paths, withholds, anticipates, frustrates, and redirects. But its meaning occurs in the movement of reading.

Reading has temporality. The reader remembers earlier passages, anticipates what may come, revises expectations, fills blanks and then discovers that the filling was wrong, is moved forward, interrupted, disappointed, surprised. This temporal movement is not equivalent to possessing the result of the reading. A summary of a novel is not the novel; a list of theses is not a philosophical work; an explanation of a poem is not the reading of the poem.

The event of reading includes the path. This is why automated summary is existentially ambiguous: it can orient the reader, but it can also replace the event through which the reader would have been formed. The summary gives the end without the path. Sometimes this is useful; sometimes it is a theft from understanding.

7. Historical reading and context

To read historically is not merely to read slowly; it is to respect the temporal and cultural specificity of the text.

A text belongs to an epoch, a language, a set of conflicts, institutions, possibilities, and expectations. Historical reading asks what questions the text was answering, what future it could imagine, what vocabulary it inherited, what stakes its words carried.

Koselleck's conceptual history is useful here because it shows that words carry historical time. Concepts live between spaces of experience and horizons of expectation. They do not mean the same thing in every epoch; their semantic weight changes as the relation between past and future changes.⁴⁰

Skinner adds another methodological caution: one must ask what an author was doing in saying what they said. A text is not only a collection of propositions but an intervention. It acts within a field of possible arguments, authorities, opponents, and audiences.⁴¹

Gumbrecht gives yet another dimension. Historical reading is not only meaning extraction; it can also involve presence, atmosphere, mood, and the material tone of a lost world. The past does not only signify; it can be felt as a mode of nearness that does not collapse distance.⁴²

These approaches differ, but they share one requirement: the text must not be flattened into timeless content. Historical reading resists the present's hunger for immediate usefulness. It asks the reader to let the past remain past, so that it can still speak.

8. Derrida, trace, and the limit of recovery

No account of writing can proceed as if the text simply preserves a stable origin. Derrida's critique of phonocentrism matters here. Writing is not a secondary supplement added to the fullness of living speech. The trace is older than any pure presence. Meaning is never simply identical with authorial intention. The written mark can be repeated because it is iterable, and this repeatability loosens it from any single origin.⁴³

This chapter accepts that caution.

Historical reading cannot recover the author's living presence. It cannot return to a pure origin. It cannot close the play of interpretation once and for all. The authorial act is mediated, displaced, and exposed to futures the author could not control.

But this does not abolish historical responsibility.

The fact that origin is not pure does not mean that origin does not matter. The fact that the trace exceeds intention does not mean that the text has no historical claim upon the reader. The fact that writing is iterable does not justify treating it as free-floating material for any use.

The trace is not a license for extraction.

A responsible hermeneutics must hold two claims together: the text cannot be reduced to authorial presence, and the text must not be severed from the historical world from which it comes.

This tension is precisely what computational processing tends to dissolve. It does not struggle with the trace as trace. It prepares the trace for operation.

9. Archive, memory, and exteriorization

Writing can become archive. The archive preserves, orders, authorizes, excludes, and retrieves; it is never simply storage but shapes what can count as memory.

This is already visible in older cultures of reading. Carruthers shows that medieval memory was not passive storage but a trained practice tied to reading, meditation, composition, and moral formation.⁴⁴ The manuscript was not simply a container of information; it belonged to an embodied discipline of memory.

The digital archive changes this relation. Ernst argues that digital memory operates through technical protocols, microtemporality, and machine addressability rather than historical narrative. Gitelman shows that documents have media histories, and Drucker warns that data and visualization practices often import non-humanistic assumptions into humanistic materials.⁴⁵

These sources matter because they prevent us from treating the archive as neutral. A text becomes something different when it is prepared for machine retrieval. It is no longer only a trace awaiting reading; it becomes material arranged for access, indexing, search, recombination, and later training.

Archive is therefore the threshold between writing and corpus. Chapter 4 will examine that threshold more directly. For now, the point is simple: when writing becomes archive, it becomes more available, but availability is not the same as understanding. A text can be preserved and still not be read; it can be retrieved and still not be heard.

10. Deep reading and the digital medium

Contemporary work on deep reading gives empirical and media-theoretical support to the phenomenological claim.

Wolf argues that the reading brain is plastic and that digital environments can encourage skimming, fragmented attention, and reduced patience for complex inference. Baron similarly shows that reading medium changes comprehension practices and learning habits. Birkerts had already warned that electronic environments threaten inwardness and historical consciousness, while Piper reminds us that the material form of the book shapes the bodily experience of reading.⁴⁶

These arguments should not be reduced to nostalgia for print. The issue is not that screens are evil or that books are sacred. The issue is that reading requires time, attention, bodily rhythm, delay, and a willingness to remain with difficulty. Any medium that erodes these conditions weakens the historical life of reading.

AI intensifies this problem because it does not merely change the surface on which reading occurs. It offers to read on the user's behalf. It summarizes, extracts, translates, answers, reduces the path.

This can help. A user may need orientation before entering a difficult text. But the danger is clear: what was once a passage through the text becomes an output about the text. The reader becomes a requester; the text becomes a service.

11. Processing is not reading

Computational processing can do important work. It can identify patterns across large corpora. It can search, index, classify, compare, cluster, and detect forms no individual reader could see. Distant reading, algorithmic criticism, and computational literary studies have shown that machines can support certain kinds of literary and historical inquiry.⁴⁷

The point is not to deny this. The point is that processing is not reading.

Reading is a historical act of appropriation; processing is a technical operation of formalization. Reading returns writing to a world; processing prepares it for calculation. Reading sustains distance as the medium of understanding; processing converts distance into relation, frequency, and proximity.

When computational analysis serves interpretation, it can become useful. When it replaces interpretation, it de-historicizes the text.

Moretti's distant reading, Ramsay's algorithmic criticism, and Underwood's computational history of literary change all require the transformation of texts into corpora. This transformation is not illegitimate by itself; it is powerful. But it is a transformation. The text is no longer encountered first as a historical address; it becomes part of a system of measurable relations.

Da's critique of computational literary studies is valuable because it shows how fragile this transformation can become when statistical findings either confirm what is already obvious or fail to support strong humanistic claims.⁴⁸ Hayles gives the broader AI version of the problem: LLMs process language through probability distributions and high-dimensional relations, not through historical reading.⁴⁹

Processing can produce insight. It cannot itself undergo the encounter that makes writing historically alive.

12. From historical distance to mathematical proximity

The central movement from reading to processing can now be stated more sharply.

Historical reading preserves distance. It lets the text come from elsewhere; it lets the author, tradition, epoch, and language resist the reader. It does not attempt to eliminate the gap between past and present; it inhabits that gap.

Computational processing converts distance into proximity. A text becomes comparable to other texts, a word becomes a token, a style becomes a feature, a concept becomes a relation, a tradition becomes a cluster, a history becomes a pattern. Proximity then becomes the dominant form of relation.

This is not yet the full vector argument. That belongs to Chapter 4. But the transition begins here. Historical distance asks: from where does this text come, and how does it address me? Mathematical proximity asks: to what is this element close, and how can that closeness be used?

Both questions can be useful. But they do not belong to the same relation to writing. The reader receives a trace from elsewhere; the processor produces relations among elements. The reader is exposed to a historical other; the model is trained on formalized residues.

This is why AI-generated summaries can be accurate and still de-historicizing. Accuracy does not restore the path of reading. A correct summary may still deprive the reader of the encounter through which the text might have altered them.

13. The simulated dialogue of the model

The chat interface gives generated language the appearance of dialogue. The user asks and the model responds, the user clarifies and the model adjusts. The exchange resembles conversation.

But this is not historical dialogue. In historical reading, the author's past act mediates a relation between the reader and another world. The text carries the trace of an authorial event. The reader's present question meets the text's historical otherness. The encounter has asymmetry, distance, and resistance.

In AI interaction, the generated response has no author in that sense. It is not the past act of someone who wrote to future readers. It is assembled in the moment from learned relations in the model, guided by prompt, context window, alignment, and probability.

The model may sound dialogical, but it has no historical self. It does not remember as a reader remembers, does not inherit as a tradition-bearer inherits, does not answer as an author answers, does not risk as a speaker risks.

The user supplies the world in which the output counts as an answer. This will become the theme of the later chapter on hallucination and user vivification. Here the preparatory point is enough: the model simulates dialogue by generating responsive language, but it does not participate in the historical dialogue that reading sustains.

14. The author, the reader, and responsibility

Historical reading preserves two responsibilities.

The first is responsibility toward the authorial trace. This does not mean that the reader must obey the author or reduce meaning to authorial intention. It means that the text is not anonymous material. It comes from a historical act. Someone wrote. Someone risked a formulation. Someone responded to a world.

The second is responsibility toward the reader's own saying. To read is not only to interpret the author. It is to let the interpretation change what one can say. The reader must ask: can I own this? Can I answer for my interpretation? Have I heard the text, or have I only used it?

Computational processing has no such responsibility. It can reproduce the language of responsibility, but it cannot assume it. It can generate interpretations, but it cannot be responsible to the authorial trace. It can summarize a text, but it cannot be transformed by the act of reading.

The user may take responsibility for generated output.

The model cannot.

This is why AI can assist reading only if the human reader remains responsible for the encounter. The model can prepare, orient, compare, and provoke. But it cannot read for the user in the historical sense.

If the user lets it do so, the user receives language about a text while avoiding the event of reading it.

15. Re-historicization

The distinction between reading and processing should not lead to a simple prohibition. AI output can be re-historicized.

A generated summary can send the user back to the text. A model's comparison can expose a question the reader had missed. A computational pattern can invite closer reading. A translation can open access. A draft interpretation can become material for critique. But this happens only when the user returns the output to historical inquiry.

Re-historicization means asking: what text is being summarized, who wrote it, in what language and for whom, against what, what question did it answer, what world does it project, what has the output flattened, what must I still read, and what am I now responsible for saying?

These questions interrupt processing. They return generated language to the hermeneutic field from which meaning arises. Without such return, AI output remains de-historicized language about texts. With such return, AI can become a tool inside reading rather than a substitute for reading.

The distinction is not technical; it is existential.

16. Preliminary thesis

The chapter establishes that writing has two possible afterlives: historical reading, which re-historicizes the written trace into dialogue across distance, tradition, and responsibility, and computational processing, which prepares it for extraction, serialization, comparison, and calculation. The model does not read. It operates on writing after writing has been detached from the historical event of reading. Only the human reader can re-historicize what computation has de-historicized.

The reader actualizes the text without consuming its otherness. Without such return to historical inquiry, AI output remains de-historicized language about texts. The next chapter takes up the technical transformation through which writing becomes archive, corpus, token, and vector.

FOOTNOTES

³⁶ Plato, *Phaedrus*, trans. Alexander Nehamas and Paul Woodruff, in *Plato: Complete Works*, ed. John M. Cooper (Indianapolis: Hackett, 1997), 274b to 278e; Jacques Derrida, "Plato's Pharmacy," in *Dissemination*, trans. Barbara Johnson (Chicago: University of Chicago Press, 1981), 61 to 171. Plato's suspicion of writing is preserved here as a necessary starting point, not as the chapter's conclusion. The written word does detach speech from the speaker's living presence and exposes it to repetition, misuse, and circulation without answer. The chapter follows Derrida in refusing the simple hierarchy of living speech over dead writing, but it does not turn that refusal into indifference toward historical address. Writing is neither pure presence nor inert mark. It is a trace whose fate depends on the relation in which it is received.

³⁷ Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 290, 335, 355, 362; Jean Grondin, *Introduction to Philosophical Hermeneutics*, trans. Joel Weinsheimer (New Haven: Yale University Press, 1994), 115; James Risser, *Hermeneutics and the Voice of the Other: Re-reading Gadamer's Philosophical Hermeneutics* (Albany: SUNY Press, 1997), 119. The chapter follows Gadamer in treating understanding as historically effected rather than methodologically purified. The reader never stands outside tradition, but this does not make interpretation a private projection. The text addresses the reader from another horizon and can alter the prejudgments through which the reader begins. The chapter narrows this Gadamerian structure toward the question of AI: the issue is not whether a text can be detached from its origin, but whether the detached text is returned to historical dialogue or prepared for operation.

³⁸ Paul Ricoeur, *From Text to Action: Essays in Hermeneutics, II*, trans. Kathleen Blamey and John B. Thompson (Evanston, IL: Northwestern University Press, 1991), chaps. 5 to 7; Paul Ricoeur, *Time and Narrative*, vol. 1, trans. Kathleen McLaughlin and David Pellauer (Chicago: University of Chicago Press, 1984), ix, 52, 86. Ricoeur's account of distanciation is crucial because it refuses both immediate authorial presence and textual deadness. Writing frees the text from the psychology of the author and from the first situation of address, but this freedom lets the text project a world in front of itself. The chapter follows Ricoeur in defining appropriation as the reader's exposure to that projected world, not as possession of the text. It extends the point by distinguishing appropriation from extraction: the same distanced writing that can be taken up in reading can also be formalized for technical operation.

³⁹ Roman Ingarden, *The Cognition of the Literary Work of Art*, trans. Ruth Ann Crowley and Kenneth R. Olson (Evanston, IL: Northwestern University Press, 1973); Wolfgang Iser, *The Act of Reading: A Theory of Aesthetic Response* (Baltimore: Johns Hopkins University Press, 1978); Hans Robert Jauss, *Toward an Aesthetic of Reception*, trans. Timothy Bahti (Minneapolis: University of Minnesota Press, 1982). Ingarden's account of concretization, Iser's account of blanks and the wandering viewpoint, and Jauss's account of horizons of expectation all prevent the text from being treated as a closed container of ready meaning. The

chapter draws on this tradition to describe reading as an event that unfolds through time. It does not make meaning arbitrary. The text guides, limits, and resists the reader. Its incompletion is the condition of reception, not a permission to substitute any response for reading.

⁴⁰ Reinhart Koselleck, *Futures Past: On the Semantics of Historical Time*, trans. Keith Tribe (Cambridge, MA: MIT Press, 1985), especially “Historia Magistra Vitae” and “Begriffsgeschichte and Social History.” Koselleck’s distinction between spaces of experience and horizons of expectation gives the chapter a vocabulary for the temporal density of concepts. A concept is not only a signifier with a stable dictionary sense. It is a historical condensation in which inherited experience and projected futures meet. The chapter uses this point against any processing of words as detachable semantic units. A term can be near another term in a formal system while remaining distant in the historical time that gives it force.

⁴¹ Quentin Skinner, *Visions of Politics*, vol. 1, *Regarding Method* (Cambridge: Cambridge University Press, 2002), especially the essays on meaning, understanding, and speech acts. Skinner’s methodological claim is not that interpretation must collapse into authorial psychology. It is that a text must be read as an intervention within a determinate field of possible utterances. To understand a political or philosophical text, one must ask what its author could have been doing in saying what was said. The chapter uses Skinner as a corrective to extraction. A proposition removed from its field of action may remain grammatically intact while losing the force that made it historically legible.

⁴² Hans Ulrich Gumbrecht, *Production of Presence: What Meaning Cannot Convey* (Stanford, CA: Stanford University Press, 2004); Hans Ulrich Gumbrecht, *Atmosphere, Mood, Stimmung: On a Hidden Potential of Literature* (Stanford, CA: Stanford University Press, 2012). Gumbrecht’s work supplements hermeneutics by insisting that reading is not exhausted by semantic interpretation. Texts can carry mood, material tone, and a felt nearness of a past world. The chapter does not oppose presence to meaning. It uses Gumbrecht to show that historical reading can let the past be encountered as more than information. This is why de-historicization is not only a loss of context. It is also the loss of the affective and material thickness through which a text can address a reader.

⁴³ Jacques Derrida, *Of Grammatology*, trans. Gayatri Chakravorty Spivak (Baltimore: Johns Hopkins University Press, 1976), especially part 1; Derrida, “Plato’s Pharmacy,” 61 to 171; Jacques Derrida, *Archive Fever: A Freudian Impression*, trans. Eric Prenowitz (Chicago: University of Chicago Press, 1996), 1 to 23. Derrida is indispensable because he prevents the chapter from romanticizing living speech as an uncontaminated origin. Writing is not a secondary fall from presence. The trace and the possibility of repetition inhabit meaning from the beginning. The chapter accepts this limit on recovery, but it also distinguishes deconstruction from technical extraction. Iterability loosens a mark from its first situation; it does not remove the reader’s responsibility to the historical trace.

⁴⁴ Mary Carruthers, *The Book of Memory: A Study of Memory in Medieval Culture*, 2nd ed. (Cambridge: Cambridge University Press, 2008). Carruthers shows that medieval memory was an art of formation, invention, ordering, recollection, and composition, not a neutral storehouse of retained data. This matters because it gives the chapter a historical counterexample to the modern reduction of memory to storage. In manuscript culture, reading and remembering were bodily and moral disciplines through which texts became available for judgment, meditation, and speech. The chapter uses this example to mark the difference between a living practice of memory and the technical availability of archived marks.

⁴⁵ Wolfgang Ernst, *Digital Memory and the Archive*, ed. Jussi Parikka (Minneapolis: University of Minnesota Press, 2013); Lisa Gitelman, *Paper Knowledge: Toward a Media History of Documents* (Durham, NC: Duke University Press, 2014); Johanna Drucker, *Visualization and Interpretation: Humanistic Approaches to Display* (Cambridge, MA: MIT Press, 2020). These sources show, in different registers, that archives and documents are not neutral containers. Ernst emphasizes machine addressability and media time; Gitelman reconstructs the documentary forms through which knowledge becomes legible; Drucker warns that humanistic materials are changed when treated as data for display and analysis. The chapter draws on them to define archive as a threshold. Once writing is arranged for retrieval, indexing, and computation, it becomes more available, but not for that reason more understood.

⁴⁶ Maryanne Wolf, *Reader, Come Home: The Reading Brain in a Digital World* (New York: Harper, 2018); Naomi S. Baron, *How We Read Now: Strategic Choices for Print, Screen, and Audio* (New York: Oxford University Press, 2021); Sven Birkerts, *The Gutenberg Elegies: The Fate of Reading in an Electronic Age* (Boston: Faber and Faber, 1994); Andrew Piper, *Book Was There: Reading in Electronic Times* (Chicago: University of Chicago Press, 2012); Naomi S. Baron, *Who Wrote This? How AI and the Lure of Efficiency Threaten Human Writing* (Stanford, CA: Stanford University Press, 2023). These works should not be read as a simple defense of print against screens. Their value for the chapter is more precise: medium, attention, bodily rhythm, and habits of inference shape what kind of relation to a text becomes possible. AI sharpens this problem because it does not only modify the site of reading. It offers an answer in place of the path through which reading forms judgment.

⁴⁷ 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). Moretti, Ramsay, and Underwood are important because they show that computational methods can disclose patterns unavailable to ordinary close reading. The chapter's argument is not a refusal of such methods. It is a description of their condition. Distant reading and algorithmic criticism require texts to become corpora, features, frequencies, models, and evidence. That conversion can serve interpretation, but it remains a conversion. The text is no longer encountered first as address. It is made operative within a system of measurable relations.

⁴⁸ Nan Z. Da, “The Computational Case against Computational Literary Studies,” *Critical Inquiry* 45, no. 3 (2019): 601 to 639. Da’s critique matters here because it contests the assumption that computational scale automatically yields stronger humanistic knowledge. Her argument exposes a recurring tension: results that are statistically stable may be critically unsurprising, while results that appear interesting may rest on fragile modeling choices. The chapter does not depend on accepting every part of Da’s polemic against computational literary studies. It uses her critique to mark the risk of formalization itself. When the text is reduced to counted, filtered, or clustered features, the operation may lose the very literary and historical phenomena it claims to illuminate.

⁴⁹ N. Katherine Hayles, “Modes of Cognition: Implications for Large Language Models,” *Antikythera Journal* (2025), doi:10.1162/ANTI.5CZF. Hayles should not be cited as if she merely repeats the claim that LLMs are next-token machines without cognitive significance. Her argument is more demanding: she proposes SIRAL criteria for cognition and argues that LLMs may satisfy them in a limited artificial register, despite lacking embodied access to the physical world. The present chapter accepts the need for that finer distinction, but draws the boundary at historical reading. Even if an LLM can process, infer, and generate meaningful text in Hayles’s broader sense of cognition, it does not appropriate a projected world, assume responsibility to an authorial trace, or undergo the temporal path of reading.

CHAPTER-LEVEL BIBLIOGRAPHY

Baron, Naomi S. *How We Read Now: Strategic Choices for Print, Screen, and Audio*. New York: Oxford University Press, 2021.

Baron, Naomi S. *Who Wrote This? How AI and the Lure of Efficiency Threaten Human Writing*. Stanford, CA: Stanford University Press, 2023.

Birkerts, Sven. *The Gutenberg Elegies: The Fate of Reading in an Electronic Age*. Boston: Faber and Faber, 1994.

Carruthers, Mary. *The Book of Memory: A Study of Memory in Medieval Culture*. 2nd ed. Cambridge: Cambridge University Press, 2008.

Da, Nan Z. “The Computational Case against Computational Literary Studies.” *Critical Inquiry* 45, no. 3 (2019): 601 to 639.

Derrida, Jacques. *Archive Fever: A Freudian Impression*. Translated by Eric Prenowitz. Chicago: University of Chicago Press, 1996.

Derrida, Jacques. *Dissemination*. Translated by Barbara Johnson. Chicago: University of Chicago Press, 1981.

- Derrida, Jacques. *Of Grammatology*. Translated by Gayatri Chakravorty Spivak. Baltimore: Johns Hopkins University Press, 1976.
- Drucker, Johanna. *Visualization and Interpretation: Humanistic Approaches to Display*. Cambridge, MA: MIT Press, 2020.
- Ernst, Wolfgang. *Digital Memory and the Archive*. Edited by Jussi Parikka. Minneapolis: University of Minnesota Press, 2013.
- Gadamer, Hans-Georg. *Truth and Method*. 2nd rev. ed. Translated by Joel Weinsheimer and Donald G. Marshall. New York: Continuum, 2004.
- Gitelman, Lisa. *Paper Knowledge: Toward a Media History of Documents*. Durham, NC: Duke University Press, 2014.
- Grondin, Jean. *Introduction to Philosophical Hermeneutics*. Translated by Joel Weinsheimer. New Haven: Yale University Press, 1994.
- Gumbrecht, Hans Ulrich. *Atmosphere, Mood, Stimmung: On a Hidden Potential of Literature*. Stanford, CA: Stanford University Press, 2012.
- Gumbrecht, Hans Ulrich. *Production of Presence: What Meaning Cannot Convey*. Stanford, CA: Stanford University Press, 2004.
- Hayles, N. Katherine. "Modes of Cognition: Implications for Large Language Models." *Antikythera Journal* (2025). doi:10.1162/ANTI.5CZF.
- Ingarden, Roman. *The Cognition of the Literary Work of Art*. Translated by Ruth Ann Crowley and Kenneth R. Olson. Evanston, IL: Northwestern University Press, 1973.
- Iser, Wolfgang. *The Act of Reading: A Theory of Aesthetic Response*. Baltimore: Johns Hopkins University Press, 1978.
- Jauss, Hans Robert. *Toward an Aesthetic of Reception*. Translated by Timothy Bahti. Minneapolis: University of Minnesota Press, 1982.
- Koselleck, Reinhart. *Futures Past: On the Semantics of Historical Time*. Translated by Keith Tribe. Cambridge, MA: MIT Press, 1985.
- Moretti, Franco. *Distant Reading*. London: Verso, 2013.
- Piper, Andrew. *Book Was There: Reading in Electronic Times*. Chicago: University of Chicago Press, 2012.
- Plato. *Phaedrus*. Translated by Alexander Nehamas and Paul Woodruff. In *Plato: Complete Works*, edited by John M. Cooper. Indianapolis: Hackett, 1997.
- Ramsay, Stephen. *Reading Machines: Toward an Algorithmic Criticism*. Urbana: University of Illinois Press, 2011.

Ricoeur, Paul. *From Text to Action: Essays in Hermeneutics, II*. Translated by Kathleen Blamey and John B. Thompson. Evanston, IL: Northwestern University Press, 1991.

Ricoeur, Paul. *Time and Narrative*. Vol. 1. Translated by Kathleen McLaughlin and David Pellauer. Chicago: University of Chicago Press, 1984.

Risser, James. *Hermeneutics and the Voice of the Other: Re-reading Gadamer's Philosophical Hermeneutics*. Albany: SUNY Press, 1997.

Skinner, Quentin. *Visions of Politics*. Vol. 1, *Regarding Method*. Cambridge: Cambridge University Press, 2002.

Underwood, Ted. *Distant Horizons: Digital Evidence and Literary Change*. Chicago: University of Chicago Press, 2019.

Wolf, Maryanne. *Reader, Come Home: The Reading Brain in a Digital World*. New York: Harper, 2018.

Chapter 4. Corpus, Token, Vector: The De-historicization of Language

From Historical Distance to Calculable Proximity

Chapter 3 argued that writing is not dead because it is written. Writing lives when it is read historically. Reading re-historicizes writing by returning the written trace to a dialogue among author, reader, text, tradition, question, and world. Processing de-historicizes writing by preparing it for extraction, serialization, comparison, and calculation.

This chapter examines that processing.

The transition from writing to AI-generated language does not happen in a single act. It passes through a technical sequence:

writing → archive → corpus → token → vector → latent proximity → generated output

At each stage, something is preserved and something is lost. The text is not simply destroyed. Its traces remain. Its patterns remain. Its linguistic relations remain. Its historical uses leave marks that can be learned. This is why large language models can generate fluent and useful language.

But what remains is no longer writing as historical address. It is no longer language as the articulation of a world. It is no longer the trace of an authorial act awaiting a reader's re-historicization. It is language prepared for operation.

The guiding thesis of this chapter is therefore:

Vectorization converts historically situated significance into calculable proximity.

The vector is not meaningless. It retains traces of reference. But it retains them after world, history, authorial situation, embodied use, and responsibility have been displaced into relation, distance, density, weight, and prediction.

The vector is the de-worlded and de-historicized trace of reference.

1. From archive to corpus

The archive preserves writing; the corpus prepares writing for computation.

A library or archive can preserve a text for possible reading. It can hold the trace in a way that still invites historical encounter. A book in an archive may remain unread, but it still retains its possible relation to author, reader, edition, context, document, and tradition. It remains available to be re-entered as a historical text.

A corpus is different. A corpus is writing formalized as material for analysis. It gathers texts not primarily to preserve their historical singularity but to make them comparable. It prepares texts for counting, sampling, alignment, indexing, segmentation, pattern recognition, and model training.

This does not make the corpus illegitimate. Corpus linguistics has taught us much about actual usage, phraseology, collocation, genre, and register, and about the difference between how language is imagined and how it is used. It corrected the fantasy that language could be understood only through idealized competence or invented examples. It forced theory to reckon with performance, repetition, frequency, and the stubborn reality of written and spoken use.⁵⁰

But this gain comes with a change in relation. The text no longer appears first as a historical address; it appears as an instance. It is no longer read first as this text by this author, from this world, under this pressure, carrying this question. It is gathered as one element in an aggregate.

The corpus does not ask what world the text projects. It asks what patterns can be extracted from it.

This is the first stage of de-historicization. The written trace becomes part of a field of equivalence. A poem, a legal clause, a corporate memo, a novel, a product review: all enter the same computational collection. Their differences remain, but they remain inside a new regime of comparability.

Language is no longer encountered primarily as event. It becomes stock.

2. Corpus and the already-said

Chapter 2 described the already-said as the sediment of public intelligibility. The already-said consists of what has been expressed, repeated, circulated, archived, institutionalized, normalized, and made publicly available. Chapter 3 then showed that writing can either remain available for historical reading or be prepared for computational processing.

The corpus is the first technical form of the already-said. It gathers language after publicness has already done its work. The words have already circulated, the genres have already stabilized, the public formulas have already become recognizable. The conceptual oppositions, explanatory habits, professional styles, and institutional idioms have already been written.

The corpus does not create public language; it accumulates it.

This accumulation matters because AI is not trained on pure meaning. It is trained on the written residue of human worlds. The model receives not *Dasein*'s being-in-the-world, but traces of what *Dasein* has already written from within worlds.

The corpus therefore sits at a decisive threshold. It still contains traces of worlded language. It still carries the remains of human practices, conflicts, stories, questions, and forms of life. But those traces are now placed under another order: they are assembled for statistical extraction.

A historical reader asks what the text asks of them. A corpus asks how the text can be used by the system. This is the difference between preservation and operationalization.

3. *The corpus as enclosure*

The corpus is not merely a technical collection; it is also an enclosure.

The language that enters a corpus may come from public writing and licensed material, digitized archives and user-generated content, books and code, forums and documents, subtitles and manuals, comments and institutional records. Once gathered, cleaned, filtered, deduplicated, and structured, this language becomes available to a model as training material.

The world that produced the language no longer appears directly. The author becomes source material. The reader becomes irrelevant to training. The historical situation becomes metadata, if it is retained at all. The language community becomes a data source; the text becomes sample.

This does not mean that every corpus is exploitative in the same way, nor that all corpus work is philosophically suspect. A scholarly corpus can preserve provenance, respect context, and serve interpretation. But the large-scale training corpus tends toward another logic: extraction at scale. Language becomes a resource to be gathered before it is returned as output.

This is where the later argument about fetish begins. The historical labor of language is absorbed into the technical system. At the interface, that labor will disappear. What returns will look like the machine's own knowledge.

But before the fetish appears, the enclosure must occur. The corpus encloses language by separating it from historical reading and making it available as computational stock.

4. The technical cut of tokenization

Once writing has become corpus, it must be made digestible by the model. It must be cut.

Tokenization is the technical procedure by which text is broken into units that can be processed. These units may be words, subwords, characters, bytes, or byte-pair fragments, depending on the tokenizer. Contemporary systems often use subword methods such as Byte-Pair Encoding, WordPiece, or related approaches, which are designed to balance vocabulary size, coverage, and computational efficiency.⁵¹

At first glance, tokenization looks like a minor engineering step. It is not. Tokenization presupposes a decision about what kind of thing language must become in order to enter computation. The flowing sentence becomes a sequence; the historical word becomes a processable unit; the text becomes segmented material.

The cut does not follow the old boundaries of humanistic understanding. A token need not be a word, a morpheme, or a unit that preserves etymology or semantic wholeness. The token is not the word as spoken, read, remembered, inherited, or contested. It is the word, or piece of word, as made fit for computation.

This matters because languages are not equally tokenized. Morphologically rich languages, non-Latin scripts, low-resource languages, dialects, and mixed linguistic forms can be fragmented differently and often less efficiently. Tokenization is not a neutral mirror of language; it is a technical imposition on language.⁵²

The corpus gathers; tokenization cuts. The first de-historicizes by aggregation, the second by fragmentation.

5. Token is not sign

Chapter 1 distinguished the Heideggerian sign from the computational token. That distinction now becomes central.

A sign indicates within a world. It points, warns, guides, names, calls, or discloses because it belongs to a practical totality of involvements. A road sign functions because roads, rules, destinations, drivers, dangers, and shared practices already matter. A word functions because it belongs to a world in which something can be said by someone to someone about something.

A token does not indicate in this way. A token functions inside a model: it is assigned an ID, embedded, transformed, weighted, and used in prediction. It may carry traces of sign-use because it derives from written language, but its immediate mode of being is not worldly indication; it is computational operability.

The sign belongs to *Bedeutsamkeit*; the token belongs to architecture. The sign points within a referential totality; the token is positioned within a technical sequence. The sign may call for response; the token calls for transformation.

This is not to say that tokens are meaningless in every sense. They become meaningful for users because they derive from language and are returned to language. But the token's meaning is not internal to the model in the way a sign's significance belongs to a world.

The token is the linguistic trace after the loss of worldly indication. It is language after the first cut of de-worlding.

6. *The loss of provenance*

Tokenization also weakens provenance.

A written sentence belongs to an author, a work, a date, a tradition, a language, a situation, a history of reception. In training, that sentence is broken into token sequences and absorbed into a distribution. Its authorial and historical situation may survive as metadata, but the training operation does not preserve that situation as the condition of understanding.

The model does not ask where a sentence came from so that it may be responsible to it; it asks how the sequence can contribute to prediction.

As training proceeds, the source sentence is no longer encountered as a sentence in a text. Its patterns are distributed into parameters. Its style, structure, idiom, and relation to other language become part of the model's learned capacity. The sentence does not remain intact as a historical saying; it becomes part of a field of possible continuations.

This is why a model can generate a convincing argumentative structure without knowing its genealogy. It can produce a scholastic distinction, a managerial synthesis, a therapeutic reassurance, a legal caution, or a philosophical caveat. It has learned the forms; it has not inherited the traditions.

The authorial act becomes statistical availability; the tradition becomes pattern; the historical source becomes signal. This is not an accidental defect. It is how the system works.

7. *Distributional semantics*

After tokenization, the token must be placed into relation. This is where distributional semantics becomes decisive.

The distributional hypothesis is often summarized through J. R. Firth's famous formulation: "You shall know a word by the company it keeps." Zellig Harris had already described language in terms of distributional structure: the occurrence of elements relative to other elements. Later distributional semantics developed this insight into methods that represent words, documents, and contexts through statistical relations.⁵³

This tradition should not be dismissed. It reveals something true.

Words do bear traces of their use. A word's contexts matter. A word's neighbors matter. Its collocations, contrasts, recurrent frames, syntactic positions, and genre patterns are not accidental. Human worlds leave marks in language. Distributional methods can discover those marks with astonishing precision.

This is why distributional semantics works.

But the philosophical question is not whether it works. The question is what kind of relation to meaning it produces.

In distributional semantics, a word's meaning is approached through its relations to other linguistic forms. Reference to world is replaced by relation within a system. The word is no longer primarily understood as something spoken, heard, written, read, answered for, or lived. It is understood as a node in a field of co-occurrence.

This is the transition:

from worldly reference to linguistic distribution

Distribution is not nothing. It is a residue of use. But it is use after it has been abstracted from the living situation in which use mattered.

Distributional semantics captures the trace of worlded language.

It does not recover the world.

8. From count to prediction

Early distributional models often used count-based co-occurrence matrices. Later neural methods learned dense vector representations that could encode relations in lower-dimensional continuous spaces. Word2Vec, GloVe, and related methods showed that distributional patterns could yield powerful representations of syntactic and semantic relation. Contextual models such as BERT and GPT then shifted the emphasis from static word vectors to context-sensitive representations generated through neural architectures.⁵⁴

Technically, this history matters because it shows a progression from counting contexts to predicting and generating contexts. The model no longer only records association. It learns to produce plausible continuation.

Philosophically, the movement matters because it deepens the transformation of language into operable relation.

The word is no longer only near other words. It becomes a position in a system that can be transformed, attended to, composed, and used to generate new sequences. Its relations are no longer merely observed; they are mobilized.

This is the threshold at which language becomes machine-generative.

A corpus preserves patterns. A vector makes patterns operable. A model makes operability generative.

The result is not simple repetition. The model can produce new combinations. But novelty here is not the same as origination. It is variation within learned relation.

The model generates from what has already been made calculable.

9. Vectorization

The vector is the decisive technical-philosophical object of this chapter.

A vector is a mathematical representation. In language models, tokens and contextual states are represented numerically in high-dimensional spaces. These representations allow the system to compute relation: similarity, distance, direction, association, attention, probability, transformation.

This is not merely a method of storage. It is a new mode in which language becomes available.

In a vector space, words, fragments, styles, contexts, and possible continuations can be related geometrically. “Near” and “far” become technical categories. Semantic resemblance becomes distance. Analogy becomes direction. Context becomes a configuration of weighted relations.

This is powerful because language has already been meaningful. Human beings have used words in worlds. Those uses have left patterns. The vector learns from these traces.

But the vector does not preserve the world as world.

It does not preserve the prayer in the act of prayer, justice in the struggle for justice, the promise in the risk of betrayal, the poem in the trembling of a reader, the word “home” in the experience of exile, the insult in the wound, or the law in the institution that enforces it. It preserves formalized traces of usage.

The vector is therefore neither empty nor alive. It is a mathematical residue of meaningful use.

Its philosophical status can be stated precisely:

The vector is the de-worlded and de-historicized trace of reference.

10. Calculable proximity

Vectorization converts historical significance into calculable proximity.

This formula should be read carefully. The claim is not that vectorization destroys all meaning. If it did, language models would not work. The claim is that vectorization changes the mode in which meaning-like relations are preserved.

Historical significance is situated. It belongs to a world, a history, a body, a speaker, a reader, a tradition, a conflict, a practice. It has depth because it comes from somewhere and can be taken up somewhere else.

Calculable proximity is different. It belongs to a mathematical space. It allows comparison, clustering, retrieval, prediction, interpolation, and generation. It does not ask what historical distance must be preserved for understanding. It asks how closeness can be computed.

The shift is therefore:

from historical distance to mathematical proximity

Chapter 3 argued that reading lives in distance. The past remains past, and that distance becomes the medium of understanding. Chapter 4 now shows that computational processing reduces distance to relation. It does not ask the reader to endure historical distance. It converts linguistic traces into near and far.

This is why an AI summary of a historical text can be accurate and still de-historicizing. It may reproduce information about the past. It may even mention context. But the mode of relation has changed. The text is no longer encountered as a historical other. It is processed as a set of relations that can be compressed into present output.

The past becomes searchable; the historical becomes proximate; the difficult becomes available.

11. The unique is made comparable

Vectorization does not make everything identical. It often does the opposite. It produces refined differences. It can distinguish genre, register, topic, style, sentiment, probable continuation, rhetorical frame, and semantic neighborhood with impressive detail.

The danger is not sameness.

The danger is universal comparability.

A poem and a policy document are not identical in vector space. But both become vectorizable. A sacred text and an online comment are not the same. But both can enter the same computational regime. A philosophical concept and a product review differ. But both can be placed, weighted, compared, clustered, retrieved, and generated from.

The unique is not erased. It is made comparable.

This matters because comparability is not neutral. To make something comparable is to decide the mode in which its difference will count. In vector space, difference counts as relation within a shared mathematical field. Historical singularity becomes position. Authority becomes distribution. Tradition becomes cluster. Style becomes reproducible pattern. Difficulty becomes local density.

Manzoni and Reddit are not equal for a historical reader. They belong to different orders of significance: Manzoni to a history of Italian language, literature, nationhood, education, religion, and the novel; Reddit to platformed contemporary discourse, internet vernaculars, anonymous social exchange, and mass digital expression. A reader may study both, but not in the same way.

For the model, both enter as distributional material. Their difference can be learned, but it is learned as relation inside the same technical space.

The passage is:

from value to weight

This is de-historicization.

12. Contextual embeddings and attention without care

Modern transformer models do more than assign static vectors to words. They generate contextual representations. The same token can take on different representations depending on surrounding tokens. Attention mechanisms allow the model to weight relations within the sequence and produce context-sensitive transformations.⁵⁵

This is technically significant. It means that contemporary AI is not blind to context in a simple sense. It can track local relation, long-range dependency, register, ambiguity, and syntactic role. It can adapt output to prompt and conversation. It can appear responsive.

A weak critique would deny this.

A stronger critique asks what kind of context is being modeled.

Model context is formal context. It is the relation among tokens, positions, embeddings, attention weights, and learned parameters. It is context as structured input and transformed state.

Worldly context is different. It is a field of concern. It includes task, mood, history, embodiment, risk, responsibility, and practical consequence. It is not only what surrounds the word. It is the world in which the word matters.

Attention is not care.

A model can attend to the token “death,” but it cannot be toward death. It can weight the relation between “promise” and “betrayal,” but it cannot betray or be betrayed. It can relate “home” to “family,” “loss,” “migration,” and “belonging,” but it cannot long for home.

The technical term “attention” is useful, but philosophically dangerous if it suggests care, concern, or worlded involvement. It is a mechanism of formal relational weighting; it is not the existential letting-something-matter.

The model can represent contexts. It does not inhabit a world.

13. Latent space is not worldhood

The vectorized language model operates in latent space. This phrase is often used loosely, but it names an important technical reality: learned representations are organized in high-dimensional mathematical spaces where features, relations, and transformations are not directly visible at the level of ordinary language.

It is tempting to call this space the model’s “world.” That temptation should be resisted. Latent space is not worldhood.

A world is not a totality of coordinates. It is the meaningful openness in which beings matter to *Dasein*. It includes involvement, embodiment, mood, temporality, being-with, history, care, and possible responsibility. A world is inhabited.

Latent space is traversed, sampled, optimized, and transformed. It contains learned relations, encodes traces, supports generation. It may have a topology rich enough to surprise even its designers. But it does not care, does not dwell, does not inherit, does not answer for what appears within it.

Yuk Hui’s work on digital objects is helpful here because it treats digital entities as having their own mode of existence, structured by data, metadata, relations, and computational schemas.⁵⁶ That point should be taken seriously. Digital objects are not nothing; they are not mere illusions. They belong to a technical milieu with its own forms of individuation and relation.

But acknowledging the mode of existence of digital objects does not require equating latent space with worldhood. Latent space is a technical milieu; worldhood is an existential structure. The difference matters.

AI has a space of relations. It does not have a world of significance.

14. Statistical pattern and inferential content

Vector space captures distributional pattern. Human language also carries inferential content.

To say something is not merely to place words near other words. It is to undertake commitments, invite consequences, expose oneself to correction, and enter a space of reasons. If I assert that a person is guilty, a treaty is unjust, a poem is sacred, a candidate is ready, or a promise has been broken, I have not merely produced a linguistic sequence. I have made myself answerable.

Brandom's inferentialism helps clarify this distinction. Meaning is not only association. It is a place within a network of commitments, entitlements, reasons, and consequences.⁵⁷

A model can simulate these relations. It can generate "therefore," "because," "however," "if," and "unless." It can produce arguments and counterarguments. It can imitate the language of responsibility and justification. But it does not undertake commitments. It is not entitled to a claim. It cannot be challenged as a speaker who must answer for what follows.

This is why proximity is not inference.

A token may be near another token. A sentence may be likely after another sentence. A generated argument may follow familiar rhetorical patterns. None of this, by itself, constitutes the normative act of meaning.

The model reckons. It does not judge.⁵⁸

This does not make its output useless. It makes its output dependent on a user who can judge.

15. Form, meaning, and grounding

Bender and Koller's distinction between linguistic form and meaning gives the computational version of this claim. A system trained only on form has no direct access to the communicative intent, world, and shared situation that give language meaning.⁵⁹ Searle's Chinese Room argument had already made the parallel point in another register: syntactic manipulation does not by itself yield semantic understanding.⁶⁰

The phenomenological argument developed here is not identical to either argument, but it converges with both.

Form is not nothing. Form is the residue of meaning. Human beings have used language from within worlds. Their intentions, practices, institutions, and histories have sedimented themselves in form. This is why AI can learn from form and return useful language.

But form is not world.

The model has access to traces of use, not to the living conditions of use. It can model how words appear. It cannot know what it is to speak from a world in which those words matter.

This gives us a more precise position than the claim that "AI does not understand." The model can model relations among linguistic forms that are themselves shaped by human understanding. It can therefore produce outputs that become meaningful for users.

But the meaning happens in the human circuit.

The model processes form. The user reintroduces world. The output becomes meaningful only when returned to use, context, and responsibility.

This is the technical basis for later user vivification.

16. Language as standing-reserve

The computational processing of language can now be read through Heidegger's later analysis of technology.

In "The Question Concerning Technology," Heidegger argues that modern technology is not merely a collection of instruments. It is a mode of revealing in which beings are challenged forth as orderable, storable, optimizable, and available. Under this revealing, beings appear as *Bestand*, standing-reserve.⁶¹

The AI language pipeline makes language standing-reserve. Writing is gathered, texts are cleaned, corpora are built, tokens are assigned, vectors are learned, latent relations are optimized, and output is generated on demand.

Language is no longer encountered first as historical saying, poetic disclosure, public speech, or readerly address. It is challenged forth as resource. It becomes available for prediction, summarization, rewriting, translation, classification, and generation.

This is not merely instrumental use of language; it is a change in how language appears. Language appears as stock. The question becomes not "what does this text ask of us?" but "what can this text be used to generate?" The text is no longer approached primarily as something to be read; it is approached as material to be processed.

In this sense, AI does not merely use language; it reveals language as standing-reserve.

17. Extraction and capital

Language as standing-reserve is not only an ontological issue. It is also material and economic.

Large-scale language models require vast corpora. Those corpora are made from human writing: books, articles, code, documents, comments, forums, websites, messages, and other traces of social life. The extraction of such language is not a purely abstract operation. It is tied to platforms, ownership, legal regimes, labor, infrastructure, energy, and capital.

Couldry and Mejias call this broader process data colonialism: the appropriation of human social life as raw material for economic value.⁶² Zuboff describes surveillance capitalism as a system that extracts behavioral surplus to predict and shape conduct.⁶³ Pasquinelli's history of AI emphasizes that machine intelligence is bound to the extraction and automation of social labor.⁶⁴

The argument of this book adds a linguistic specification.

AI does not only extract data; it extracts the historical labor of language. It extracts forms of saying and styles of reasoning, idioms and arguments, jokes and prayers and confessions, explanations and instructions. It extracts the already-said, formalizes it, and returns it as a service.

This is why the later chapter will describe AI as a fetish of language. The historical labor of language will appear as the machine's own knowledge. But Chapter 4 shows the technical condition for that fetish: language must first become extractable stock.

The corpus is enclosure; the token is cut; the vector is mathematical possession. The output is the commodity form of the already-said.

18. Distant reading and computational humanities

The humanities have already faced part of this question through debates on distant reading and computational literary studies.

Moretti's distant reading proposed that large-scale literary history requires methods that do not depend on close reading of a small canon. Ramsay treated algorithmic criticism as a way of provoking interpretation rather than replacing it. Underwood has shown that computational methods can reveal patterns of literary change across scale.⁶⁵

These projects are not identical with AI training. They are scholarly practices, and many of them remain answerable to interpretation. They often use computation as a tool to generate questions for human readers.

But they share one structural feature with the AI pipeline: the text becomes corpus.

The singular work becomes part of an aggregate. The literary event becomes a point in a field. The historical text becomes available for pattern detection. This can produce insight, but it also risks de-historicization if the aggregate replaces the encounter.

Da's critique of computational literary studies is useful here because it shows that statistically tractable patterns do not automatically yield humanistic understanding.⁶⁶ Drucker's work on visualization and interpretation makes a similar point from another direction: humanistic material cannot simply be imported into data models without changing the assumptions under which it appears.⁶⁷

The lesson is not that computation cannot assist reading.

The lesson is that processing must remain subordinate to interpretation.

When the corpus becomes the authority, reading is displaced.

19. *What remains of significance?*

We can now return to *Bedeutsamkeit*.

What remains of significance after corpus, token, vector, and latent proximity? Something remains.

If nothing remained, LLMs would not work. The model's fluency is possible because human significance has sedimented itself in language. The corpus contains traces of worlds; the token carries fragments of expression; the vector preserves relations born from meaningful use; latent space organizes these traces into generative capacity.

But what remains is not significance as worldhood. It is not the totality of involvements in which things matter, the silence of a situation, the body's exposure to the world, the historical distance of reading, the author's responsibility, the reader's appropriation, the speaker's commitment, or *Dasein*'s care. What remains is formalized residue.

This residue is powerful. It can help, orient, summarize, reveal patterns, produce language that becomes meaningful for users. But it becomes meaningful only when returned to a world. AI does not carry worldhood inside the vector; it carries traces that require human re-worlding.

This is why the next chapters must speak of afterlife, *Gerede*, fetish, and vivification. The model's output is not dead in the simple sense; it can become alive in use. But its life does not belong to the model. It belongs to the historical language that preceded it and the human world that receives it.

20. *Preliminary thesis*

The chapter establishes that the corpus encloses the already-said for computational processing, the token cuts language into units fit for machine operation, and the vector retains traces of reference after world and history have been displaced into calculable proximity. Latent space is not worldhood; it is a technical milieu in which language has been formalized for extraction, optimization, and generation. Language becomes standing-reserve when held available to this technical regime.

The vector is the de-worlded and de-historicized trace of reference. Latent space is a technical milieu, not a world of significance. The next chapter defines AI as the vectorized afterlife of the They.

FOOTNOTES

⁵⁰ Noam Chomsky, *Aspects of the Theory of Syntax* (Cambridge, MA: MIT Press, 1965), 3 to 15; Tony McEnery and Andrew Hardie, *Corpus Linguistics: Method, Theory and Practice* (Cambridge: Cambridge University Press, 2012), chap. 1; John Sinclair, *Corpus, Concordance, Collocation* (Oxford: Oxford University Press, 1991), chaps. 1 to 3. The chapter treats corpus linguistics as a genuine correction of purely idealized accounts of language. Chomsky's competence-performance distinction made actual utterance methodologically suspect for a certain theoretical project, while Sinclair, McEnery, and Hardie show how patterned use can become the object of disciplined inquiry. The book's claim is not that corpus study is mistaken. It is that the corpus changes the relation to writing: language becomes available as comparable material before it can become address.

⁵¹ Rico Sennrich, Barry Haddow, and Alexandra Birch, "Neural Machine Translation of Rare Words with Subword Units," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics* (2016), 1715 to 1725; Mike Schuster and Kaisuke Nakajima, "Japanese and Korean Voice Search," in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing* (Piscataway, NJ: IEEE, 2012), 5149 to 5152; Taku Kudo and John Richardson, "SentencePiece: A Simple and Language Independent Subword Tokenizer and Detokenizer for Neural Text Processing," in *Proceedings of EMNLP 2018: System Demonstrations* (2018), 66 to 71. BPE, WordPiece, and SentencePiece are cited here because they make visible the technical cut that precedes vectorization. The point is not that tokenizer design already contains a full philosophy of language. The narrower claim is that every tokenizing procedure decides what kind of unit language must become in order to be processable.

⁵² Emily M. Bender, "On Achieving and Evaluating Language-Independence in NLP," *Linguistic Issues in Language Technology* 6 (2011): 1 to 26; Sennrich, Haddow, and Birch, "Neural Machine Translation of Rare Words with Subword Units," 1715 to 1725; Kudo and Richardson, "SentencePiece," 66 to 71. Bender's warning about language-independent NLP gives the technical background for the chapter's claim that tokenization is not neutral. A method that appears general can still privilege some morphological, orthographic, and script traditions over others. The chapter extends this point phenomenologically: unequal fragmentation is not only a performance issue. It shows that the machine's unit of operation is imposed on historically formed language rather than received from it.

⁵³ Zellig Harris, "Distributional Structure," *Word* 10, nos. 2 to 3 (1954): 146 to 162; J. R. Firth, "A Synopsis of Linguistic Theory, 1930 to 1955," in *Studies in Linguistic Analysis* (Oxford: Blackwell, 1957), 1 to 32; Alessandro Lenci, "Distributional Models of Word Meaning," *Annual Review of Linguistics* 4 (2018): 151 to 171. The chapter uses the distributional hypothesis neither as a caricature nor as a simple enemy. Harris and Firth identify a real fact about language: use leaves patterned traces, and those traces can be studied. The philosophical limit appears when relation among linguistic forms is treated as enough for meaning. Distribution preserves the sediment of worlded use, but it does not itself restore the world in which that use mattered.

- ⁵⁴ Marco Baroni, Georgiana Dinu, and Germán Kruszewski, “Don’t Count, Predict! A Systematic Comparison of Context-Counting vs. Context-Predicting Semantic Vectors,” in *Proceedings of ACL 2014* (2014), 238 to 247; 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* (2014), 1532 to 1543; Jacob Devlin et al., “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proceedings of NAACL-HLT 2019* (2019), 4171 to 4186. This technical lineage matters because the chapter does not reduce contemporary models to crude word counts. Dense embeddings, prediction-based training, and contextual representations make language more operable, not less. They sharpen the thesis: historical use becomes a field of learned relation that can be converted into plausible continuation.
- ⁵⁵ Ashish Vaswani et al., “Attention Is All You Need,” in *Advances in Neural Information Processing Systems 30* (2017); Devlin et al., “BERT,” 4171 to 4186; Tom B. Brown et al., “Language Models are Few-Shot Learners,” in *Advances in Neural Information Processing Systems 33* (2020), 1877 to 1901. The chapter uses attention strictly as a technical term. The architecture’s attention mechanism allows the model to weight relations among positions in a sequence and to produce context-sensitive representations. This makes the critique stronger rather than weaker. Contemporary models can track context in a formal and highly articulated way. They still do not inhabit context as concern, risk, mood, or responsibility. Attention is relation inside the model, not care in the existential sense.
- ⁵⁶ Yuk Hui, *On the Mode of Existence of Digital Objects* (Minneapolis: University of Minnesota Press, 2016); Gilbert Simondon, *On the Mode of Existence of Technical Objects*, trans. Cécile Malaspina and John Rogove (Minneapolis: Univocal, 2017). Hui’s work is important because it refuses the simple claim that digital objects are merely unreal. Digital objects have a mode of existence, a milieu, and forms of individuation structured by data, metadata, relations, and schemas. The chapter follows Hui this far. It then draws a boundary: a technical milieu is not the same as worldhood. Latent space can organize relations among digital objects without becoming the meaningful openness in which *Dasein* dwells and answers for what appears.
- ⁵⁷ Robert Brandom, *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (Cambridge, MA: Harvard University Press, 1994); Robert Brandom, *Articulating Reasons: An Introduction to Inferentialism* (Cambridge, MA: Harvard University Press, 2000). Brandom’s inferentialism supplies an analytic counterpart to the chapter’s phenomenological claim. Meaning is not only a relation among signs or a correspondence between word and object. It belongs to practices of giving and asking for reasons, where speakers undertake commitments and acquire entitlements. The chapter does not translate

Heidegger into Brandom. It uses Brandom to state, in normative vocabulary, why proximity in vector space cannot by itself become inference. A model can reproduce the language of reasons without entering the responsibility of reason-giving.

⁵⁸ Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, MA: MIT Press, 2019). Smith's distinction between reckoning and judgment names the boundary the chapter needs at this point. AI systems can reckon with extraordinary power: they can calculate, classify, optimize, and generate outputs from formalized inputs. Judgment requires another relation, one involving the world, commitment, and responsibility for what is at stake. The chapter adopts this distinction without treating it as a rejection of AI's usefulness. The output can be valuable, but its value depends on a human user capable of judging what the system has only reckoned.

⁵⁹ Emily M. Bender and Alexander Koller, "Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Stroudsburg, PA: Association for Computational Linguistics, 2020), 5185 to 5198. Bender and Koller's distinction between linguistic form and meaning gives the computational version of a boundary this book develops phenomenologically. A model trained on form has access to patterns produced by meaningful human activity, but not to the communicative situation, intention, and shared world in which those patterns first mattered. The chapter extends their argument from grounding to worldhood. The model handles form; the user returns that form to a world.

⁶⁰ John R. Searle, "Minds, Brains, and Programs," *Behavioral and Brain Sciences* 3, no. 3 (1980): 417 to 424; Smith, *The Promise of Artificial Intelligence*. Searle's Chinese Room argument is used here as a historical precursor to the form-meaning distinction, not as the final authority on machine intelligence. Neural language models are not simple rule books, and their statistical complexity should not be minimized. Yet the central boundary remains: formal manipulation, whether rule-based or learned, does not by itself establish semantic answerability. The chapter's Heideggerian version of the point is that understanding requires world-involvement, not only correct formal conversion of signs.

⁶¹ 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; Iain D. Thomson, *Heidegger on Ontotheology: Technology and the Politics of Education* (Cambridge: Cambridge University Press, 2005), 44 to 77; Andrew J. Mitchell, *The Fourfold: Reading the Late Heidegger* (Evanston, IL: Northwestern University Press, 2015). Heidegger's *Bestand* names a mode of revealing in which beings are disclosed as orderable reserve. The chapter applies this structure to language without claiming that Heidegger anticipated NLP. The point is ontological: when writing is gathered, cleaned, segmented, embedded, optimized, and generated on demand, language appears as available stock. Thomson clarifies the technological drive toward flexibility and calculability. Mitchell supplies the counterpoint of thinging and world-gathering that such availability suppresses.

⁶² Nick Couldry and Ulises A. Mejias, *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism* (Stanford, CA: Stanford University Press, 2019). Couldry and Mejias give the chapter its political-economic vocabulary for extraction. Data colonialism names the appropriation of social life as raw material for capital. The present argument narrows that claim to language: AI training appropriates not only information about human beings, but the historical labor of saying through which communities have made worlds intelligible. The corpus is therefore not only a technical object. It is an enclosure of linguistic sociality.

⁶³ Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: PublicAffairs, 2019), chaps. 1 to 3. Zuboff's account of behavioral surplus gives the economic background for the chapter's account of linguistic surplus. Surveillance capitalism extracts traces of conduct in order to predict and modify future conduct. The AI language pipeline extracts traces of saying in order to predict and generate future saying. The book does not simply equate these processes. It uses Zuboff to clarify the wider regime in which language becomes a source of predictive value and not only a medium of address.

⁶⁴ Matteo Pasquinelli, *The Eye of the Master: A Social History of Artificial Intelligence* (London: Verso, 2023), 1 to 24, 187 to 216. Pasquinelli's history of AI as the automation and objectification of social labor is a central interlocutor for this chapter. The book follows his insistence that machine intelligence is never detached from labor, institutions, and capital. It adds a narrower ontological claim: AI also objectifies the historical labor of language. Styles, genres, arguments, explanations, and forms of public saying are extracted as learnable relation. The model's apparent knowledge is therefore built from social labor that has become technically latent.

⁶⁵ 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 works prevent the chapter from treating computation in the humanities as mere error. Distant reading, algorithmic criticism, and machine-learning literary history can reveal patterns unavailable to ordinary close reading and can generate questions for interpretation. The chapter's caution concerns priority. Computation can assist reading when it returns to historical judgment. It becomes de-historicizing when the aggregate displaces the encounter with the singular text.

⁶⁶ Nan Z. Da, "The Computational Case against Computational Literary Studies," *Critical Inquiry* 45, no. 3 (2019): 601 to 639. Da's critique is useful as a warning about the gap between statistical tractability and humanistic understanding. A result can be technically demonstrable yet critically thin; a claim can be humanistically interesting yet dependent on fragile modeling choices. The chapter does not adopt Da's essay as a general ban on computational literary

study. It uses the critique to mark the risk of substituting model-readable pattern for historical significance. A corpus can produce evidence, but evidence still has to be returned to interpretation.

⁶⁷ Johanna Drucker, *Visualization and Interpretation: Humanistic Approaches to Display* (Cambridge, MA: MIT Press, 2020). Drucker’s account of humanistic data is decisive because it rejects the fantasy that data simply arrive as neutral givens. Humanistic materials are constructed, selected, interpreted, encoded, displayed, and framed. The chapter uses Drucker to support its final rule for computational method: processing must remain subordinate to interpretation. When language is converted into data without preserving that interpretive accountability, the model does not only analyze the text. It changes the conditions under which the text can appear.

CHAPTER-LEVEL BIBLIOGRAPHY

Baroni, Marco, Georgiana Dinu, and Germán Kruszewski. “Don’t Count, Predict! A Systematic Comparison of Context-Counting vs. Context-Predicting Semantic Vectors.” In *Proceedings of ACL 2014*, 238 to 247. 2014.

Bender, Emily M. “On Achieving and Evaluating Language-Independence in NLP.” *Linguistic Issues in Language Technology* 6 (2011): 1 to 26.

Bender, Emily M., and Alexander Koller. “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data.” In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185 to 5198. Stroudsburg, PA: Association for Computational Linguistics, 2020.

Brandom, Robert. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press, 2000.

Brandom, Robert. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press, 1994.

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, et al. “Language Models are Few-Shot Learners.” In *Advances in Neural Information Processing Systems* 33, 1877 to 1901. 2020.

Chomsky, Noam. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press, 1965.

Couldry, Nick, and Ulises A. Mejias. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press, 2019.

Da, Nan Z. “The Computational Case against Computational Literary Studies.” *Critical Inquiry* 45, no. 3 (2019): 601 to 639.

- Devlin, Jacob, 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. 2019.
- Drucker, Johanna. *Visualization and Interpretation: Humanistic Approaches to Display*. Cambridge, MA: MIT Press, 2020.
- Firth, J. R. "A Synopsis of Linguistic Theory, 1930 to 1955." In *Studies in Linguistic Analysis*, 1 to 32. Oxford: Blackwell, 1957.
- Harris, Zellig. "Distributional Structure." *Word* 10, nos. 2 to 3 (1954): 146 to 162.
- Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.
- Hui, Yuk. *On the Mode of Existence of Digital Objects*. Minneapolis: University of Minnesota Press, 2016.
- Kudo, Taku, and John Richardson. "SentencePiece: A Simple and Language Independent Subword Tokenizer and Detokenizer for Neural Text Processing." In *Proceedings of EMNLP 2018: System Demonstrations*, 66 to 71. 2018.
- Lenci, Alessandro. "Distributional Models of Word Meaning." *Annual Review of Linguistics* 4 (2018): 151 to 171.
- McEnery, Tony, and Andrew Hardie. *Corpus Linguistics: Method, Theory and Practice*. Cambridge: Cambridge University Press, 2012.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. "Efficient Estimation of Word Representations in Vector Space." arXiv:1301.3781. 2013.
- Mitchell, Andrew J. *The Fourfold: Reading the Late Heidegger*. Evanston, IL: Northwestern University Press, 2015.
- Moretti, Franco. *Distant Reading*. London: Verso, 2013.
- Pasquinelli, Matteo. *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso, 2023.
- Pennington, Jeffrey, 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*, 1532 to 1543. 2014.
- Ramsay, Stephen. *Reading Machines: Toward an Algorithmic Criticism*. Urbana: University of Illinois Press, 2011.

Schuster, Mike, and Kaisuke Nakajima. "Japanese and Korean Voice Search." In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing*, 5149 to 5152. Piscataway, NJ: IEEE, 2012.

Searle, John R. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3, no. 3 (1980): 417 to 424.

Sennrich, Rico, Barry Haddow, and Alexandra Birch. "Neural Machine Translation of Rare Words with Subword Units." In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 1715 to 1725. 2016.

Simondon, Gilbert. *On the Mode of Existence of Technical Objects*. Translated by Cécile Malaspina and John Rogove. Minneapolis: Univocal, 2017.

Sinclair, John. *Corpus, Concordance, Collocation*. Oxford: Oxford University Press, 1991.

Smith, Brian Cantwell. *The Promise of Artificial Intelligence: Reckoning and Judgment*. Cambridge, MA: MIT Press, 2019.

Thomson, Iain D. *Heidegger on Ontotheology: Technology and the Politics of Education*. Cambridge: Cambridge University Press, 2005.

Underwood, Ted. *Distant Horizons: Digital Evidence and Literary Change*. Chicago: University of Chicago Press, 2019.

Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *Advances in Neural Information Processing Systems 30*. 2017.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.

Chapter 5. The Vectorized Afterlife of the They

AI as Intraworldly Technical Entity

Chapter 4 traced the technical sequence by which writing becomes corpus, corpus becomes token, token becomes vector, vector becomes latent proximity, and latent proximity becomes generated output. That sequence showed how historical language is made operable. The written trace is not simply destroyed. It is preserved as formalized residue. But the preservation occurs through de-worlding and de-historicization.

We can now name the ontological status of the system that performs this operation.

AI is not *das Man*.

This claim must be stated first because the temptation to identify them is strong. Generative AI speaks anonymously. It produces fluent public language. It often sounds like no one and everyone. It repeats common formulations, levels difficulty into usable explanation, and gives the user the form of understanding before appropriation. It therefore seems natural to call it the technological embodiment of the They.

But this is not precise enough.

Das Man is an existential structure of everyday *Dasein*. It belongs to a being that can own or fail to own its possibilities, a being exposed to anxiety, thrownness, care, finitude, and death. AI has none of these. It does not flee from itself into publicness. It does not seek refuge in averageness. It does not cover over mortality. It has no ownmost possibility to disown.

AI is an intraworldly technical entity.

It is encountered within the world as a system, interface, platform, tool, infrastructure, and service. It is world-internal, but not world-forming. It does not inhabit a world of significance. It operates on language that worlded beings have already produced.

The thesis of this chapter is therefore:

AI is not the They. It is the vectorized afterlife of the language of the They.

This formula has three parts. First, AI is not an existential structure. Second, AI operates on sedimented public language. Third, vectorization gives that sediment a technical afterlife as generated speech. Chapter 5 develops these three claims.

1. Why “afterlife” is necessary

The word “afterlife” is not decorative. It is needed because the language processed by AI is neither alive in the ordinary human sense nor simply dead.

If one says that AI works on dead language, the claim is too crude. Chapter 3 already showed why. Writing is not dead because it is written. A text can live when it is read historically, when a reader reactivates it across distance, and when the written trace enters a new world of concern.

If one says that AI resurrects language, the claim is also wrong. Resurrection would restore living worldhood, historical origin, authorial responsibility, and the existential force of the saying. AI does not restore these. It does not bring back the author. It does not return the text to the world that first made it matter. It does not reconstitute the event of speech or reading.

What AI does is different.

It gives language another life, a life after historical reading has been suspended and after writing has been formalized as computational stock.

This is afterlife, not resurrection.

The language lives on, but in a changed mode. It survives as corpus, token, vector, latent relation, and generated continuation. It can move, answer, recombine, and appear present. Yet this movement is not the movement of a living speaker. It is the motion of sedimented language after technical reanimation.

The afterlife names survival through transformation.

2. Benjamin, Warburg, and transformed survival

Benjamin’s account of translation gives one way to think this survival. A translation does not simply reproduce the informational content of an original. It marks a stage in the continuing life of a work. The original lives on by entering another linguistic configuration, and this survival is inseparable from alteration.⁶⁸

AI vectorization is not translation in Benjamin's sense. It does not serve the work's own linguistic destiny. It does not preserve the tension between original and translation as a historical relation. Yet Benjamin gives us a necessary distinction: survival is not simple persistence. A work can continue by undergoing a structural change.

Warburg's concept of *Nachleben* gives a second model. The afterlife of antiquity is not the uninterrupted life of classical forms, but their survival through displacement, migration, and reactivation. Ancient gestures and formulas return in later epochs, no longer identical with their first world, yet still carrying force.⁶⁹

AI-generated speech can be understood as a degraded technical version of such survival. It gathers formulas of public language, argumentative patterns, styles of explanation, gestures of sympathy, structures of authority, and rhetorical shapes of competence. These forms survive in the model, not as living cultural memory, but as learnable relation. They return without the historical bodies that first bore them.

A public phrase survives as a likely continuation; a rhetorical gesture as pattern; a tone as style; a tradition as cluster; a concept as proximity. The afterlife is real, but it is altered. What lives on has been detached from the world in which it first lived.

3. Sedimentation

The language AI processes is sediment.

Husserl's account of sedimentation in *The Crisis of European Sciences* is the decisive point of departure. Mathematical and scientific idealities can become transmissible and repeatable only by sedimenting. But once sedimented, they can also conceal their origin in living acts of meaning. What began as an achievement of sense can harden into an inherited formation whose genesis is forgotten.⁷⁰

This is not only a scientific problem. It is a general phenomenological structure. Human practices, concepts, gestures, traditions, and linguistic habits can settle into availability. They become what one inherits without having produced. They support life because they make action possible without constant origination. They also conceal life because they can detach inherited forms from the experience that first gave them meaning.

The already-said is such sediment.

It is language after it has settled into public availability. It includes phrases, genres, explanations, judgments, professional formulas, political frames, moral commonplaces, pedagogical templates, and styles of address. These are not meaningless. They carry deposits of use. They are the remains of human worldhood in language.

AI does not encounter the living experience that produced this sediment. It encounters the sediment after it has become writable, retrievable, extractable, and formalizable. The machine is not a subject that inherits a tradition; it is a system that mines sediment.

4. Generative sedimentation and public language

Sedimentation is also generative. Steinbock's generative phenomenology is useful here because it shows how historical worlds pass themselves on through traditions, institutions, home-worlds, and foreign-worlds. What a community receives is not merely a set of isolated facts. It receives sedimented ways of seeing, acting, speaking, and valuing.⁷¹

The corpus is a technical capture of such generative sedimentation.

This claim must be handled carefully. A corpus is not the living tradition itself. It is not the community, the home-world, the institution, or the act of transmission. It is the written and digitized deposit left by these worlds after they have become available for extraction.

The corpus preserves traces of generative life.

But it preserves them after they have been severed from the living chain of transmission. It does not preserve the teacher correcting the student, the family teaching an idiom, the political struggle giving force to a word, the reader being changed by a book, or the community arguing over what a tradition means.

It preserves the written remains.

In this sense, the training corpus is the technical sediment of public intelligibility. It is the already-said after it has become available to the machine.

The They speaks, writing sediments, the archive stores, the corpus encloses, the vector formalizes, the model generates. This sequence is the path from public language to technical afterlife.

5. AI as intraworldly technical entity

AI must now be located ontologically.

It is not *Dasein*. It is not another subject in the world. It is not a speaker, reader, thinker, or historical inheritor. It is an entity encountered within the world.

Simondon's account of technical objects helps here. Technical objects can develop internal coherence. They can become more concrete, more integrated with their associated milieus, and less dependent on external human adjustment. Yet this individuation is not biological, psychic, or existential individuation.⁷²

Yuk Hui extends this problem to digital objects. Digital objects are not mere nothing. They have a mode of existence structured by data, metadata, relations, protocols, and computational systems. They are technical entities whose being depends on networks, standards, and recursive operations.⁷³

AI belongs here.

It is a digital technical object, or more precisely a complex technical system of models, data, interfaces, infrastructures, and practices. It has internal structure. It mediates human activity. It shapes what users perceive, ask, expect, and treat as available.

But this does not make it *Dasein*. It has a technical milieu, not a world; parameters, not care; context windows, not historical memory; optimization, not responsibility; recurrence, not inheritance; output, not speech in the existential sense.

To call AI intraworldly is therefore not to diminish its power; it is to locate that power correctly.

It is powerful inside our world; it does not have a world of its own.

6. Postphenomenological mediation and its limit

Postphenomenology helps clarify how AI can shape human experience without becoming human.

Ihde and Verbeek show that technologies mediate relations between humans and world. A tool is not simply passive. It can amplify, reduce, translate, prescribe, invite, and constrain. Technologies shape perception and action.⁷⁴

AI does this with language.

It mediates how users ask questions, receive explanations, produce writing, summarize sources, formulate positions, and experience competence. It does not merely sit before the user as an object. It enters the user's practical relation to language.

But mediation is not existential being-with.

A technology can shape my access to the world without sharing that world as another *Dasein*. It can guide my perception without seeing. It can recommend without judging. It can respond without being responsible. It can simulate dialogue without becoming an interlocutor.

This distinction will matter later when user vivification is analyzed. For now, the point is narrower: AI mediates human-world relations through language, but its mediating power should not be confused with worldhood, care, or social existence.

It is not a quasi-person.

It is a technical mediation of public linguistic sediment.

7. The fallacy of identification

We can now return to the temptation that began the chapter.

AI seems like *das Man* because it speaks anonymously, because it generates public language, because it levels difficulty, because it often produces average intelligibility, and because it gives the user the already-said.

But resemblance at the level of output is not identity at the level of being.

Das Man belongs to *Dasein*. It is the everyday mode in which a finite being finds itself dispersed in publicness. It is bound to thrownness, being-with, anxiety, care, and the possibility of authenticity. It is an existential structure of a being that can lose itself because it can own itself.

AI cannot lose itself.

It cannot be absorbed in the They because it cannot be outside the They. It cannot flee into averageness because it has no anxiety. It cannot speak idly in order to avoid responsibility because it never has responsibility. It cannot level down the world as a mode of self-evasion because it has no self.

AI is not the They; it is the technical processing of what the They leaves behind.

This correction matters because to call AI *das Man* grants it too much. It turns a technical artifact into an existential structure. It mystifies the machine at the very moment the argument should de-mystify it.

The machine is not the anonymous public; it is built from the public's written sediment.

8. AI as recursion on the already-said

If AI is not the They, what is its relation to the They? It is a recursion on the already-said.

The They produces public intelligibility, *Gerede* circulates it, writing records it, archives preserve it, corpora enclose it, tokens fragment it, vectors organize it as proximity, and the model generates new speech-like sequences from it.

The model therefore begins after public language has already spoken.

It does not participate in the public world as a finite being. It does not compare itself with others, seek normality, or hide in accepted formulas. Yet it can generate language that bears the marks of such publicness because those marks are already present in the training material.

This is why AI output often has the tone of anonymous competence.

It can sound academic, managerial, therapeutic, journalistic, legal, bureaucratic, or intimate. It can produce the recognizable forms of public discourse because those forms have already stabilized in writing. The model learns them as patterns. It then returns them as output.

The afterlife of the They is not the continuation of a subject; it is the recursion of a sediment.

9. The vectorized form of survival

The afterlife becomes specifically vectorized when public language is transformed into mathematical relation.

A phrase no longer survives as phrase alone. A style no longer survives as style alone. A concept no longer survives as concept alone. Each survives as part of a high-dimensional relational field. The model does not preserve the text as a human archive preserves it. It distributes the text's relations across learned representations.

Survival becomes spatial.

This is why the afterlife differs from ordinary quotation, memory, translation, or commentary. A quotation preserves a phrase. A memory preserves an experience, however altered. A translation carries a work into another language. Commentary returns to a text as something to be interpreted.

Vectorization does something else.

It converts the residues of public language into calculable proximity. It allows forms of saying to become retrievable, recombinable, and generative without preserving their historical origin as origin.

A metaphor survives as a pattern of adjacency. A public opinion survives as likely continuation. A professional formula survives as register. A tradition survives as distribution. An author survives as style-space. A concept survives as neighborhood. The model does not remember; it reactivates relation.

This is why "afterlife" must be joined to "vectorized." The language does not merely live on. It lives on as geometry.

10. Publicness without appearance

Arendt's account of the public realm gives a useful contrast. For Arendt, the public is a space of appearance in which speech and action disclose the "who" of the actor before others. Publicness, in that sense, is not only exposure to the average. It is the condition under which singular human beings appear through word and deed.⁷⁵

AI-generated language is public in another sense.

It draws on public language, but it does not disclose a "who." It has no actor behind it. It has no life whose uniqueness appears through speech. It can produce the surface of public address without the existential disclosure of a person.

This is not Arendtian appearance; it is public language without appearing. The output speaks but no one appears, answers but no one stands forth, has tone but no voice, form but no “who.” This is a key feature of the afterlife. Public language survives, but the singularity of public action disappears. What remains is speech-like availability.

That availability will later become central to the analysis of *Gerede* on steroids.

11. Proletarianization of public speech

Stiegler’s account of technics and proletarianization gives another angle.

For Stiegler, human memory and temporality are constituted by technical exteriorization. Technical supports are not external additions to human life. They are part of how human history is transmitted. But industrial and digital systems can also proletarianize knowledge. They can expropriate *savoir-faire*, *savoir-vivre*, and *savoir-théoriser*, weakening the capacities through which individuals and groups individuate themselves.⁷⁶

AI extends this problem into public speech.

The capacity to summarize, formulate, classify, explain, and continue discourse is externalized into the model. This does not mean that human beings lose these capacities automatically. Nor does it mean that every use of AI is proletarianizing. But the risk is clear: the system can perform the linguistic operations through which users once had to form themselves.

The user no longer needs to compose the summary: the model composes it. The user no longer needs to begin the argument: the model begins it. The user no longer needs to endure the first blankness of thought: the model fills it. This is the proletarianization of the already-said: public linguistic competence is automated and returned as service. The user gains speed. The user may lose practice.

The afterlife of the They therefore has a social consequence. It does not merely preserve public language. It changes the relation between public language and human formation.

12. Anonymity and the neutral voice

The They is anonymous. AI intensifies anonymity, but in a different register.

In everyday publicness, anonymity belongs to the grammar of “one.” One says, one thinks, one expects, one knows. This anonymity is existential because *Dasein* can hide within it.

In AI-generated language, anonymity becomes technical. The output has no speaker. It may use first-person grammar, but no self stands behind the “I.” It may address the user, but no one addresses. It may express caution, sympathy, or confidence, but no subject owns the stance.

Blanchot's and Levinas's reflections on anonymous existing can help describe the uncanny neutrality of this voice, provided they are not overextended.⁷⁷ AI is not the *il y a* in a metaphysical sense. It is not anonymous being itself. But its generated language can feel like an impersonal murmur of already-public speech. Derrida's analysis of writing also matters here: writing can operate in the absence of the author. AI radicalizes that absence by producing new speech-like output without an authorial act.⁷⁸

The result is a neutral voice: not because it is unbiased, not because it is objective, but because no one owns it.

This neutrality is one reason the output is seductive. It seems to come from nowhere, and what comes from nowhere can appear impartial. In truth, it comes from everywhere and no one: from corpus, training, alignment, prompt, infrastructure, and statistical relation.

The neutral voice is the voice of public sediment after authorship has been technically erased.

13. Reanimation and presence-effect

Generated output does not merely sit as stored text. It appears.

It appears in real time, in response to a prompt. It adapts to the user's question. It produces a voice-like surface. This gives it what might be called a presence-effect.

Gumbrecht's distinction between meaning and presence is useful here. A text or artifact can affect us not only by conveying meaning, but by producing a sense of presence, force, or nearness.⁷⁹ AI-generated language often produces such a presence-effect. It arrives immediately. It speaks in the second person. It seems available, patient, and responsive.

But this presence is not the presence of a living speaker.

It is generated presence.

The model gives movement to sediment. The interface gives the movement the form of response. The user gives the response a world in which it matters. Chapter 8 will analyze this as user vivification. Chapter 5 only needs to establish the structural condition: the afterlife becomes compelling because the sediment does not remain inert. It returns as responsive language.

The text seems to live because it moves. But motion is not life.

14. The archive that answers

The afterlife of the They differs from the ordinary archive because it answers.

An archive preserves traces. It may remain silent until searched, read, interpreted, or cited. It confronts the reader with fragments, gaps, documents, and distance. It does not usually produce continuity on behalf of the reader.

AI changes this.

It transforms the archive into an answering apparatus.

The archive says: the trace is here. The model says: here is an answer. The archive preserves discontinuity; the model produces continuity. The archive requires reading; the model offers output. The archive can sustain historical distance; the model often collapses distance into response.

This is why AI feels more alive than a database. It is not only storage. It is reanimated storage. It is the already-said returned in the grammar of address.

The user asks, and the sediment answers.

Yet this answer should not be confused with understanding. It is the generated afterlife of stored public language, not the speech of a historical subject.

15. *The technical They as heuristic*

Chapter 2 allowed the phrase “technical They” only as a heuristic. This chapter confirms the caution.

If “technical They” means that AI is *das Man*, the phrase should be rejected.

If it means that AI makes the written sediment of public language technically available on demand, the phrase can be useful.

AI gives the already-said an interface. It makes public formulas responsive. It provides average intelligibility in the form of personal address. It turns public language into an interactive service. Explain this, summarize that, draft a position, give me objections, make it sound professional, tell me what one would say. The user privately summons publicness.

This is new. The already-said no longer circulates only through conversation, media, institutions, and texts. It can now be generated for the user in the moment. Publicness becomes promptable.

But again, the machine is not the They.

The They is the existential structure through which everyday *Dasein* first inhabits public meaning. AI is the technical system through which sedimented public meaning is made generative.

The distinction must hold.

16. *Afterlife without futurity*

The afterlife of the They is bound to the past.

The model's generation may feel open, creative, or future-facing. It can produce novelty-effects. It can offer possibilities never previously written in exactly that form. But its generative capacity is grounded in sediment already given: training data, learned relations, alignment procedures, and the prompt.

This gives AI a peculiar relation to time. It generates the new from the already-said, projects possibilities from the sediment of past discourse, simulates openness through recombination. But it has no futurity in the existential sense.

For *Dasein*, futurity belongs to projection. A person stands ahead of themselves, takes up possibilities, risks loss, and decides under finitude. For the model, "future" means generated continuation, not existential projection.

This matters because the afterlife of the They can appear to open possibilities while actually remaining bound to inherited distributions. It can multiply options without owning any of them. It can produce new language without entering a future.

The output is future-shaped but past-bound.

This is why AI may intensify public averageness even when it appears creative. It generates from sedimented possibility, not from historical projection.

17. *Why the afterlife matters*

The phrase "vectorized afterlife of the They" is not merely a conceptual label. It explains why AI is useful, why it is seductive, and why it is dangerous.

It is useful because public language contains real sedimented knowledge. Human beings have already thought, written, argued, taught, explained, judged, remembered, and classified. Their worlds have left traces. AI can work on those traces.

It is seductive because the traces return as fluent response. The user receives language without having to reconstruct the historical path by which the language became possible. The answer appears immediate.

It is dangerous because the afterlife can be mistaken for life. The model's output can be received as if a living intelligence had spoken. The historical labor of language can appear as the machine's own knowledge. Public sediment can appear as personal understanding.

This is the hinge between Chapter 5 and the chapters that follow.

Chapter 6 will analyze the dominant mode of this generated speech as *Gerede* on steroids. Chapter 7 will interpret the machine as an intraworldly fetish of language. Chapter 8 will show how the user vivifies the output and hallucinates a living interlocutor. But all three depend on the claim established here.

AI is not the They.

It is the afterlife of the language of the They.

18. Preliminary thesis

The chapter establishes that AI is not the They but an intraworldly technical entity that operates on the sediment of public intelligibility. The language of the They survives in AI as sediment, not as living publicness. Vectorization gives this sediment a technical afterlife in which generated output is not resurrected speech but reanimated public residue.

The model does not speak from a world. It generates from the written remains of worlds. The next chapter analyzes the dominant discursive form of this afterlife: *Gerede* on steroids.

FOOTNOTES

⁶⁸ Benjamin's account of translation as the continuing life of a work is useful here because it treats survival as transformation rather than duplication. See Walter Benjamin, "The Task of the Translator," in *Illuminations*, trans. Harry Zohn, ed. Hannah Arendt (New York: Schocken Books, 1968), 69 to 82. I do not claim that AI vectorization is translation in Benjamin's sense. Translation remains a historical and linguistic relation to a work. The analogy is limited: both name a survival that changes the original, but AI's "translation" into vector space is a technical de-historicization rather than a linguistic renewal.

⁶⁹ Aby Warburg's *Nachleben der Antike* gives the stronger cultural-historical model of survival as migration and reactivation. See Aby Warburg, *The Renewal of Pagan Antiquity: Contributions to the Cultural History of the European Renaissance*, trans. David Britt (Los Angeles: Getty Research Institute, 1999). Ernst Gombrich's biography remains a useful, if contested, guide to Warburg's intellectual project. See Ernst H. Gombrich, *Aby Warburg: An Intellectual Biography* (London: Warburg Institute, 1970). I use Warburg only to clarify the structure of survival without origin. AI-generated language is not the cultural afterlife of forms in Warburg's sense; it is a technical afterlife in which forms survive as learnable relations.

⁷⁰ Husserl's diagnosis of sedimentation in the modern sciences provides the chapter's phenomenological vocabulary for inherited formations whose genesis has been forgotten. See Edmund Husserl, *The Crisis of European Sciences and Transcendental Phenomenology*, trans. David Carr (Evanston, IL: Northwestern University Press, 1970), especially 48 to 59 and 353

to 378. The analogy to AI is not exact. Husserl analyzes idealities whose sense can in principle be reactivated through phenomenological return. AI training, by contrast, operates on sediment without itself being capable of such return.

⁷¹ Anthony J. Steinbock distinguishes static, genetic, and generative phenomenology, with generative phenomenology focused on the historical and communal transmission of sense. See Anthony J. Steinbock, *Home and Beyond: Generative Phenomenology after Husserl* (Evanston, IL: Northwestern University Press, 1995), 170 to 204. Klaus Held's account of the lifeworld also supports the claim that inherited social reality is structured by sedimented traditions. See Klaus Held, "Husserl's Phenomenology of the Life-World," in *The New Husserl: A Critical Reader*, ed. Donn Welton (Bloomington: Indiana University Press, 2003), 32 to 64. This chapter's claim is narrower: the training corpus is not generative tradition itself, but the technical capture of written traces left by generative traditions.

⁷² Gilbert Simondon's account of technical concretization prevents a crude dismissal of machines as inert instruments. See Gilbert Simondon, *On the Mode of Existence of Technical Objects*, trans. Cécile Malaspina and John Rogove (Minneapolis: Univocal, 2017), 25 to 60. Simondon helps clarify that technical objects can have real modes of individuation and associated milieus. The point here is that such individuation is not existential individuation. A technical object can become internally coherent without becoming *Dasein*.

⁷³ Yuk Hui gives the strongest contemporary framework for treating digital objects as objects with a specific mode of existence. See Yuk Hui, *On the Existence of Digital Objects* (Minneapolis: University of Minnesota Press, 2016), 1 to 45; Yuk Hui, *Recursivity and Contingency* (London: Rowman & Littlefield, 2019), 1 to 36. Hui is useful because he lets us avoid two errors at once: treating digital objects as nothing, and treating them as subjects. AI belongs to a digital technical milieu, but this milieu is not Heideggerian worldhood.

⁷⁴ Don Ihde and Peter-Paul Verbeek are useful because postphenomenology shows that technologies mediate human-world relations rather than merely serving neutral purposes. See Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Bloomington: Indiana University Press, 1990); Peter-Paul Verbeek, *What Things Do: Philosophical Reflections on Technology, Agency, and Design*, trans. Robert P. Crease (University Park: Pennsylvania State University Press, 2005). I use this tradition with one limit: technological mediation must not be confused with existential being-with. AI can mediate language, but it does not become another *Dasein*.

⁷⁵ Hannah Arendt, *The Human Condition* (Chicago: University of Chicago Press, 1958), 50 to 58 and 175 to 181. Arendt's public realm differs significantly from Heidegger's publicness of the *They*. For Arendt, public speech and action can disclose a singular "who," while Heidegger emphasizes the anonymous leveling of everyday intelligibility. The contrast is useful here because AI-generated speech draws on public language while disclosing no "who" at all.

- ⁷⁶ Bernard Stiegler, *Technics and Time, 1: The Fault of Epimetheus*, trans. Richard Beardsworth and George Collins (Stanford, CA: Stanford University Press, 1998); Bernard Stiegler, *Automatic Society, Volume 1: The Future of Work*, trans. Daniel Ross (Cambridge: Polity, 2016). I use Stiegler's proletarianization thesis selectively. AI does not automatically destroy human discursive capacity. It externalizes and automates operations that can weaken such capacity when they substitute for practice, formation, and appropriation.
- ⁷⁷ Maurice Blanchot, *The Infinite Conversation*, trans. Susan Hanson (Minneapolis: University of Minnesota Press, 1993); Emmanuel Levinas, *Existence and Existents*, trans. Alphonso Lingis (The Hague: Martinus Nijhoff, 1978), 52 to 65. The anonymous neutrality described by Blanchot and Levinas should not be directly identified with AI. The analogy is phenomenological, not ontological: AI-generated language can feel like an impersonal murmur of no-one, but the system remains a technical entity rather than anonymous being itself.
- ⁷⁸ Derrida's account of writing in the absence of the author is an important boundary for this chapter. See Jacques Derrida, "Signature Event Context," in *Limited Inc*, trans. Samuel Weber and Jeffrey Mehlman (Evanston, IL: Northwestern University Press, 1988), 1 to 23; Jacques Derrida, *Of Grammatology*, trans. Gayatri Chakravorty Spivak (Baltimore: Johns Hopkins University Press, 1976). AI radicalizes authorial absence because it generates new speech-like sequences without an authorial event. This does not prove that AI understands writing; it shows that writing's detachability can be technically exploited.
- ⁷⁹ Hans Ulrich Gumbrecht's distinction between meaning-effects and presence-effects helps describe why generated text can feel compelling even when no subject speaks. See Hans Ulrich Gumbrecht, *Production of Presence: What Meaning Cannot Convey* (Stanford, CA: Stanford University Press, 2004), 51 to 90. This chapter uses "presence-effect" carefully. AI does not produce the presence of a living speaker. It produces a responsive surface that users may experience as presence.

CHAPTER-LEVEL BIBLIOGRAPHY

- Arendt, Hannah. *The Human Condition*. Chicago: University of Chicago Press, 1958.
- Benjamin, Walter. "The Task of the Translator." In *Illuminations*, translated by Harry Zohn, edited by Hannah Arendt, 69 to 82. New York: Schocken Books, 1968.
- Blanchot, Maurice. *The Infinite Conversation*. Translated by Susan Hanson. Minneapolis: University of Minnesota Press, 1993.
- Derrida, Jacques. *Of Grammatology*. Translated by Gayatri Chakravorty Spivak. Baltimore: Johns Hopkins University Press, 1976.

Derrida, Jacques. "Signature Event Context." In *Limited Inc*, translated by Samuel Weber and Jeffrey Mehlman, 1 to 23. Evanston, IL: Northwestern University Press, 1988.

Gombrich, Ernst H. *Aby Warburg: An Intellectual Biography*. London: Warburg Institute, 1970.

Gumbrecht, Hans Ulrich. *Production of Presence: What Meaning Cannot Convey*. Stanford, CA: Stanford University Press, 2004.

Heidegger, Martin. *Bremen and Freiburg Lectures: Insight Into That Which Is and Basic Principles of Thinking*. Translated by Andrew J. Mitchell. Bloomington: Indiana University Press, 2012.

Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.

Held, Klaus. "Husserl's Phenomenology of the Life-World." In *The New Husserl: A Critical Reader*, edited by Donn Welton, 32 to 64. Bloomington: Indiana University Press, 2003.

Hui, Yuk. *On the Existence of Digital Objects*. Minneapolis: University of Minnesota Press, 2016.

Hui, Yuk. *Recursivity and Contingency*. London: Rowman & Littlefield, 2019.

Husserl, Edmund. *The Crisis of European Sciences and Transcendental Phenomenology*. Translated by David Carr. Evanston, IL: Northwestern University Press, 1970.

Ihde, Don. *Technology and the Lifeworld: From Garden to Earth*. Bloomington: Indiana University Press, 1990.

Levinas, Emmanuel. *Existence and Existents*. Translated by Alphonso Lingis. The Hague: Martinus Nijhoff, 1978.

Mitchell, Andrew J. *The Fourfold: Reading the Late Heidegger*. Evanston, IL: Northwestern University Press, 2015.

Simondon, Gilbert. *On the Mode of Existence of Technical Objects*. Translated by Cécile Malaspina and John Rogove. Minneapolis: Univocal, 2017.

Steinbock, Anthony J. *Home and Beyond: Generative Phenomenology after Husserl*. Evanston, IL: Northwestern University Press, 1995.

Stiegler, Bernard. *Automatic Society, Volume 1: The Future of Work*. Translated by Daniel Ross. Cambridge: Polity, 2016.

Stiegler, Bernard. *Technics and Time, 1: The Fault of Epimetheus*. Translated by Richard Beardsworth and George Collins. Stanford, CA: Stanford University Press, 1998.

Suchman, Lucy. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge: Cambridge University Press, 2007.

Thomson, Iain D. *Heidegger on Ontotheology: Technology and the Politics of Education*. Cambridge: Cambridge University Press, 2005.

Verbeek, Peter-Paul. *What Things Do: Philosophical Reflections on Technology, Agency, and Design*. Translated by Robert P. Crease. University Park: Pennsylvania State University Press, 2005.

Warburg, Aby. *The Renewal of Pagan Antiquity: Contributions to the Cultural History of the European Renaissance*. Translated by David Britt. Los Angeles: Getty Research Institute, 1999.

Chapter 6. Gerede on Steroids

Generated Speech and the Industrialization of Idle Talk

Chapter 5 defined AI as the vectorized afterlife of the language of the They. AI is not *das Man*. It is not an existential structure of everyday *Dasein*. It is an intraworldly technical entity that operates on public linguistic sediment after that sediment has been written, archived, enclosed, tokenized, vectorized, and made generative.

We can now name the dominant discursive form of this afterlife.

It is *Gerede* on steroids.

The phrase is borrowed from Taylor Carman, who first used it to describe generated discourse in light of Heidegger's analysis of idle talk, and is discussed by Iain Thomson in his recent treatment of Heidegger and AI.⁸⁰ I use the phrase as a diagnostic and not as a literal ontological identification. AI is not *Gerede* as an existential mode of *Dasein*. It generates language that can function as industrialized *Gerede* when taken up by users without appropriation.

The phrase is useful only if it is handled carefully. It must not mean that AI is merely stupid chatter, empty noise, or a machine that says foolish things. *Gerede*, usually translated as idle talk, is not primarily a moral defect. It is a structural possibility of discourse. It names the everyday circulation of intelligibility when what is said has become detached from original appropriation.

Gerede lets one speak about something without having made the matter one's own.

Generative AI industrializes this possibility. It makes the already-said promptable, scalable, personalized, and fluent. It produces competence-shaped discourse on demand, often before the user has read, thought, decided, or understood. It turns the public sediment of language into an always-available supply of answer-like speech.

This chapter's thesis is therefore precise:

AI-generated speech functions as industrialized *Gerede* when it gives the user the form of understanding without the labor of appropriation.

AI itself does not fall into idle talk. It has no self, no anxiety, no care, no possibility of authenticity, and no relation to its own death. The user can fall. The user can receive generated output as if it were understanding. The user can speak with language that has not become theirs. The machine generates. Everyday *Dasein* takes up the output. In that circuit, *Gerede* becomes technical, fluent, and industrial.

1. Gerede beyond moralism

The first task is to resist the moralistic reading of *Gerede*.

Idle talk is not merely gossip, trivial speech, or intellectual laziness. Heidegger's point is structural. Everyday *Dasein* is thrown into a public world already interpreted. Words, judgments, explanations, and opinions circulate before any individual has returned to the matter itself. This circulation is not an accident. It belongs to the way everyday understanding works.⁸¹

One begins with what is already said. One hears what one says, repeats what one knows, understands in the way things are publicly understood. This is necessary. Without inherited public discourse, one could not begin to understand at all. There would be no language into which one could enter, no common world, no shared intelligibility. *Gerede* is therefore not simply a fall from pure speech into corruption. It is the ambiguous everyday medium through which the world is first available.

But the same medium carries a danger.

Because the world is already publicly interpreted, *Dasein* can avoid interpreting it for itself. Because language is already available, one can speak without owning the saying. Because the matter has already been discussed, one can appear to understand it without having encountered its difficulty.

Gerede is the comfort of the already-understood. This is the point at which AI becomes existentially important: it does not invent the desire for already-available speech but gives that desire a technical instrument.

2. From Rede to Gerede

Chapter 1 established that *Rede*, discourse, is the existential articulation of intelligibility. It is more originary than language in the ordinary sense. The world is already articulated before words appear. Language expresses an intelligibility that being-in-the-world has already opened.

Gerede is not the absence of *Rede*. It is a deficient or fallen mode of it.

In *Rede*, language can disclose. It can articulate the matter from within an encounter with it. It can let the thing be seen, not because the speaker has mastered it, but because the saying remains answerable to the matter.

In *Gerede*, speech circulates. The matter is no longer approached in its difficulty; it is passed along in a form already publicly intelligible. The relation to the thing is weakened. The language remains; the ground recedes.

This is why idle talk can say true things.

Its defect is not necessarily falsity. A statement can be correct and still be idle. It can repeat a truth without owning the path through which that truth becomes understood. It can use accurate vocabulary without returning to the matter. It can speak rightly and still speak groundlessly.

This will become crucial for AI.

The danger of AI-generated language is not only that it may hallucinate falsehood. The deeper danger is that it can produce unowned truth. It can give the user correct, fluent, useful statements that the user has not appropriated.

A false answer can be checked. An unowned truth can pass as understanding.

3. The already-said as material

Chapter 2 introduced the already-said as the sediment of public intelligibility. Chapter 5 showed that AI is the vectorized afterlife of this sediment. Chapter 6 now adds that the already-said is not only material for generation. It is also the material of *Gerede*.

The already-said includes clichés, but it is not limited to clichés. It includes academic formulas, professional vocabularies, moral frames, political narratives, pedagogical explanations, managerial templates, legal phrasing, journalistic conventions, therapeutic idioms, and social styles of authority.

Much of it is useful; much of it is unavoidable. The danger begins when the already-said substitutes for the user's own relation to the matter: a concept becomes a phrase, a tradition becomes a summary, a book becomes key points, a decision becomes options, a position becomes a generated paragraph.

AI excels at this because it operates on the already-said after it has become written, public, vectorized, and generative. It does not need to understand a tradition in order to produce the recognizable language of that tradition. It does not need to read a book historically in order to summarize what public discourse says about that book. It does not need to undergo formation in order to generate the external form of competence.

The already-said becomes output. This is the technical beginning of industrialized *Gerede*.

4. What industrialization adds

To say that AI industrializes *Gerede* is not merely to say that it produces more idle talk. Industrialization names a historical transformation in the mode of production.

Industrial production makes a product scalable, standardized, continuous, repeatable, and available on demand. It separates the output from the slow, localized, embodied labor by which similar things were previously made. It abstracts from craft. It optimizes throughput.

AI does this to language. It standardizes the production of explanation, scales the production of summaries, automates the production of positions, makes argument available on demand, turns formulation into a service.

The user no longer needs to dwell with a question until a sentence becomes possible: the model produces the sentence first. The user no longer needs to endure the blank page: the model fills it. The user no longer needs to form an initial structure: the model provides one.

This is not merely convenience. It is the industrial transformation of linguistic labor.

The language that once required time, friction, training, reading, revision, and failure can now be generated as output. The slow practice of saying becomes a prompt-response process.

This is why Stiegler's account of the industrialization of memory and attention matters. Technical exteriorization can support human individuation, but industrial systems can also short-circuit it by replacing the practices through which individuals and collectives form themselves.⁸²

AI externalizes not only memory or search; it externalizes the first movement of saying.

5. The culture industry of language

The Frankfurt School provides an important precedent for this argument.

Adorno and Horkheimer's culture industry thesis was not simply that mass culture produces bad art. It was that culture, once industrialized, begins to supply the consumer with pre-formed patterns of interpretation. The consumer is relieved of the work of synthesis because the product has already schematized experience in advance.⁸³

AI extends this logic into ordinary discourse.

It does not only industrialize entertainment, narrative, or opinion. It industrializes the production of sentences, explanations, academic paragraphs, reports, emails, strategies, arguments, and reflective positions.

It provides schematism as a service. The user asks for a structure, the model supplies one; the user asks for a position, the model drafts one; the user asks for nuance, the model simulates it; the user asks for a conclusion, the model closes the matter.

Here the danger is not that the output is necessarily poor. Often it is useful. The danger is that the work of synthesis may be transferred away from the user. The user receives a pre-digested form of thinking and may mistake that form for thought.

Menke's account of aesthetic negativity helps clarify what is lost. Genuine thinking, like genuine aesthetic experience, often requires friction, interruption, and resistance. It does not simply satisfy the demand for smooth comprehension. It disturbs automatic understanding.⁸⁴

AI tends in the opposite direction.

It smooths.

6. Frictionless formulation

Formulation is not external to thought.

Often one discovers what one thinks only in the struggle to say it. The sentence is not packaging for a complete idea that already exists in the mind. It is part of the formation of the thought itself.

This is why the blank page matters.

The blank page is not merely a practical obstacle. It reveals that the user has not yet found their relation to the matter. It demands ownership. It asks: what do you actually think? What can you say? What do you not yet understand? What must remain unsaid?

AI changes this situation.

It gives formulation before appropriation.

The user can ask for an opening, an outline, an argument, a tone, a paragraph, a revision. The model produces. The user receives language that may be better formed than the user's own initial thought.

This can help. A writer may use generated language as material, provocation, contrast, or scaffolding. But the danger is structural: the user may accept formulation as a substitute for the path by which thought becomes their own.

The sentence arrives too early.

The user may then revise language they have not yet earned.

This is one of the deepest forms of industrialized *Gerede*: it does not only let the user repeat what one says. It lets the user generate what one might say before the user knows whether it can be said as their own.

7. Fluency and cognitive ease

AI-generated *Gerede* is powerful because it is fluent.

Fluency has long been treated as a sign of intelligence, mastery, and authority. A clear answer feels more trustworthy than a broken one. A well-organized paragraph feels more competent than a hesitant thought. A smooth explanation feels as if it has understood the matter.

This is psychologically understandable. Kahneman's work on cognitive ease shows that human beings tend to experience fluent, easily processed information as more familiar, more credible, and less demanding.⁸⁵

AI exploits this tendency without needing to intend anything. The output is fluent, so it feels understood; the paragraph is well structured, so it feels competent; the tone is balanced, so the answer feels responsible; the answer arrives quickly, so the matter feels handled.

Murdoch helps name the moral danger. Attention to reality is slow, humble, and difficult. Brilliant language can become a screen if it replaces patient attention to what is actually there.⁸⁶ The problem is not eloquence itself. The problem is eloquence without attention.

AI can produce language that resembles attention without having attended. The user may then receive fluency as if it were insight. That is the affective mechanism of industrialized *Gerede*.

8. Speaking what one has not appropriated

The existential danger of *Gerede* is that one can speak about what one has not appropriated.

AI intensifies this possibility.

A user can generate a philosophical abstract without having read the philosophical tradition. A grant proposal without having formed the research. A legal memo without understanding the law. A strategy without knowing the market. A reflection without having reflected. A position without having decided.

This is not always deception. The user may be learning. The output may be a starting point. AI can help someone approach a domain that would otherwise remain closed.

But the temptation is obvious: the user can appear to know before knowing.

Levinas's distinction between the saying and the said helps here. The saying involves exposure, vulnerability, and responsibility before the other. The said is the stabilized content that can circulate as statement. AI operates entirely in the field of the said. It generates statements. It cannot expose itself in saying.⁸⁷

The user may then take the said without undergoing the vulnerability of saying.

Bourdieu and Goffman add the social dimension. Fluency can function as cultural capital. Competent language can secure status, authority, and face. Generated text can therefore become a prosthesis for impression management.⁸⁸

The user does not only use AI to know; the user uses AI to appear as one who knows, sometimes even before themselves.

9. Expertise without formation

Human expertise is not acquired by receiving fluent outputs.

Dreyfus and Dreyfus argued that skill develops through stages, from rule-following novice behavior toward embodied, situational mastery. The expert does not merely apply explicit rules. The expert sees what matters in a concrete situation because years of practice have formed perception, judgment, and response.⁸⁹

Collins's work on tacit knowledge supports the same point from another direction. Much expertise is socially embedded. It is transmitted through immersion, practice, correction, imitation, participation, and shared forms of life. It cannot be fully converted into explicit propositions.⁹⁰

AI gives access to the external language of expertise without requiring the formation that makes expertise real.

This is not a denial of AI's usefulness. It can summarize expert discourse, produce helpful explanations, and support learning. But when used as a substitute for formation, it creates the appearance of expertise without the embodied and social path of expertise.

The user receives the language of the expert; the user may remain a novice.

This is why generated *Gerede* is more dangerous than ordinary idle talk. Ordinary idle talk often reveals its second-hand character. AI-generated speech can simulate the polish of formation. It can produce expert-looking language at the point where the user has not yet become capable.

The output may be expert-shaped; the user may not be expert.

10. Deskilling and cognitive proletarianization

The automation of speech-like production has consequences for human capacities.

Industrial automation has long been criticized for deskilling labor. Braverman showed how industrial management separated conception from execution, stripping workers of craft and transferring knowledge into machinery and managerial systems.⁹¹

In the digital age, the same logic moves upward into cognitive and linguistic life. Stiegler calls this proletarianization: not only the loss of manual know-how, but the loss of knowing how to live, attend, judge, and think. Pasquinelli's social history of AI similarly treats machine intelligence as the extraction and automation of collective knowledge and labor.⁹²

LLMs extend this process to formulation itself.

If the user no longer practices summarizing, they may lose the capacity to summarize. If the user no longer struggles to write, they may lose the discipline of writing. If the user no longer has to form an argument, they may become dependent on generated argument. If the user no longer endures the difficulty of reading, they may lose the patience required for reading.

This is not inevitable.

Tools can support practice. They can expand capacity. They can create scaffolding for learning. But if the tool replaces the practice by which capacity is formed, the capacity weakens.

The risk is not that AI writes. The risk is that humans stop learning how to say.

11. Generated publicness as private response

Traditional *Gerede* circulates publicly.

One hears what people say. One reads what is written. One absorbs opinions from conversation, institutions, newspapers, television, platforms, and professional discourse. Public language moves through exposure.

AI changes the tempo and the address.

The already-said becomes privately promptable.

The user asks a question in solitude and receives the public language of the They as a personal response. The output feels intimate because it answers the user's prompt. Yet its material comes from public linguistic sediment.

This is a new configuration of everydayness.

The user does not merely enter public discourse. Public discourse comes to the user in the form of tailored answer. The already-said becomes responsive. The They becomes conversational in appearance, though no one is speaking.

This is one reason AI-generated language is so easily vivified. The output is not just average public language. It is average public language addressed to me.

The user may then confuse personalization with appropriation.

But a response tailored to my prompt is not yet my thought. It becomes mine only if I return it to sources, context, judgment, and responsibility.

Otherwise, I possess personalized *Gerede*.

12. The false comfort of balance

AI-generated language often appears responsible because it is balanced.

It defines terms, presents both sides, lists risks, offers caveats, and avoids strong commitments. This can be useful. It can prevent crude simplification. It can give the user an initial map.

But balance can also become a form of leveling: historical conflict becomes a set of perspectives, moral difficulty becomes pros and cons, a decision becomes an option matrix, a wound becomes a theme, a radical claim becomes one view among others.

This is not always wrong. Some matters need comparison. Some users need a first orientation. But *Gerede* thrives when the availability of balanced language replaces the difficulty of taking a position.

The model does not risk anything in its balance; the user may use that balance to postpone responsibility.

This is why generated balance should be treated as preliminary, not final. It may open a field. It must not close the question.

The output can say: here are the sides. Only the user can decide what they can answer for.

13. Accuracy and unowned truth

The technical debate about AI often centers on hallucination. This is necessary. Models fabricate facts, sources, quotations, and claims. The risk is serious.

But *Gerede* shows that accuracy does not solve the deeper problem.

An accurate generated answer can still be idle talk if the user has not appropriated it. A correct summary can still replace reading. A valid distinction can still remain external. A true claim can still be unowned.

Bender and Koller's distinction between form and meaning helps clarify why. A system trained on linguistic form can learn patterns that make outputs plausible and useful, but the user must still ask whether the language has been grounded in world, context, and responsibility.⁹³

The question is not only: is the output true? The question is also: can I own this truth?

This is where Brandom's inferentialism becomes helpful. To assert something is to enter a space of commitments and consequences. One must be able to give and ask for reasons.⁹⁴ The model can generate the language of reasons. It does not undertake commitments. It does not become responsible for what follows.

The user must.

Unowned truth is one of the most subtle forms of generated *Gerede*: not wrong enough to be rejected, not owned enough to be understanding.

14. When AI helps thought

The argument must not become a simple rejection of AI.

Generated language can support thought. It can orient a beginner. It can help a reader identify a structure. It can expose a possible objection. It can offer alternative phrasing. It can assist translation. It can summarize a field before deeper reading begins. It can reduce unnecessary friction.

The question is not whether AI should be used.

The question is whether it replaces appropriation.

AI can be used as orientation, but it becomes dangerous when orientation is mistaken for understanding. It can be used as draft, but dangerous when the draft is treated as authorship; as summary, but dangerous when the summary replaces reading; as provocation, but dangerous when provocation becomes conclusion; as map, but dangerous when the map becomes the world.

This distinction keeps the critique from becoming technophobic. The point is not to preserve human difficulty for its own sake. Some difficulty is unnecessary. Some friction is waste. But some friction is formative. It is the path through which understanding becomes one's own.

The task is to distinguish waste from formation.

AI cannot do that for us.

15. Interruption

If AI industrializes *Gerede*, the first counter-practice is interruption.

The interface invites continuation: ask again, refine, make it sharper, make it more academic, give me ten options. The movement can become endless. Thinking often begins when this continuation stops.

Interruption means refusing the next prompt long enough to ask what has happened. Do I need more output, or do I need to understand? Do I need another version, or do I need to choose? Do I need a better formulation, or do I need to read? Do I need a summary, or do I need to encounter the text?

Heidegger's analysis of silence already prepared this point. Silence is not the absence of discourse. It can be a mode of discourse. It can let the matter matter.⁹⁵

AI is structurally ordered toward generated continuation. The human user must preserve the capacity for meaningful interruption.

Sometimes the responsible act is not to prompt but to read. Sometimes it is to write badly before writing well. Sometimes it is to remain with not-knowing. Interruption is the beginning of reappropriation.

16. Preliminary thesis

The chapter establishes that *Gerede* is not stupid chatter but public discourse without original appropriation, and that AI does not fall into *Gerede* but generates language that can function as *Gerede* in the human user-machine circuit. The already-said becomes promptable, scalable, fluent, and personalized, giving the user the form of understanding before the labor of understanding. The deepest danger is not only hallucinated falsehood but unowned truth.

A responsible relation to AI therefore requires interruption, reappropriation, and the preservation of not-knowing. The next chapter analyzes AI as an intraworldly fetish of language, since the historical labor of language has disappeared inside the technical object.⁹⁶

FOOTNOTES

⁸⁰ Taylor Carman first formulated the phrase “*Gerede* on steroids” in the context of generated discourse. Iain D. Thomson discusses Carman’s phrase in *Heidegger on Technology’s Danger and Promise in the Age of AI* (Cambridge: Cambridge University Press, 2025), 60. I use the phrase as a diagnostic and not as a literal ontological identification. AI is not *Gerede* as an existential mode of *Dasein*. It produces generated language that can function as industrialized *Gerede* when taken up by users without appropriation.

⁸¹ Heidegger’s analysis of *Gerede* appears in §35 of *Being and Time*, where idle talk is treated as a structural mode of everyday discourse rather than as mere gossip or stupidity. See Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §35. The reading adopted here follows William Blattner’s warning against moralizing Heidegger’s everydayness and Daniel Dahlstrom’s emphasis on the ontological analysis of language. See William D. Blattner, *Heidegger’s Being and Time: A*

Reader's Guide (London: Continuum, 2006), chap. 5; Daniel O. Dahlstrom, "Heidegger's Ontological Analysis of Language," in *Heidegger's Being and Time: Critical Essays*, ed. Richard Polt (Lanham, MD: Rowman & Littlefield, 2005), 153 to 172. Katherine Withy's account of uncanniness is also useful because it shows how falling and public familiarity cover over the unsettling character of existence. See Katherine Withy, *Heidegger on Being Uncanny* (Cambridge, MA: Harvard University Press, 2015).

⁸² Stiegler's account of tertiary retention and the industrialization of memory gives the best vocabulary for this chapter's shift from existential structure to industrial production. See Bernard Stiegler, *Technics and Time, 3: Cinematic Time and the Question of Malaise*, trans. Stephen Barker (Stanford, CA: Stanford University Press, 2011); Bernard Stiegler, *For a New Critique of Political Economy*, trans. Daniel Ross (Cambridge: Polity, 2010). I use Stiegler selectively. The chapter does not claim that every technical exteriorization degrades thought. It claims that industrialized exteriorization becomes dangerous when it replaces the practices through which thinking and saying are formed.

⁸³ The argument here adapts Adorno and Horkheimer's culture industry thesis to generated language. See Max Horkheimer and Theodor W. Adorno, *Dialectic of Enlightenment: Philosophical Fragments*, trans. Edmund Jephcott (Stanford, CA: Stanford University Press, 2002), 94 to 136. The analogy is limited. AI-generated language is not simply mass entertainment. It extends the culture industry's logic into ordinary discursive production: emails, summaries, memos, essays, positions, and explanations.

⁸⁴ Christoph Menke's reconstruction of aesthetic negativity helps name what frictionless generated language tends to erase: the disruptive, non-automatic experience through which understanding is unsettled. See Christoph Menke, *The Sovereignty of Art: Aesthetic Negativity in Adorno and Derrida*, trans. Neil Solomon (Cambridge, MA: MIT Press, 1998). I do not claim that AI cannot be used in artistic or critical practice. The claim is narrower: when generated fluency replaces negativity and resistance, it produces the smooth surface of thought without the event of thinking.

⁸⁵ Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011), especially the discussion of cognitive ease. Kahneman's framework is not phenomenological, but it helps explain why fluent generated language can feel true, familiar, and competent before the user has verified or appropriated it. This psychological mechanism supports the Heideggerian account without replacing it.

⁸⁶ Iris Murdoch's moral philosophy of attention is useful because she treats attention as a disciplined, patient relation to reality, not as a mere mental operation. See Iris Murdoch, *The Sovereignty of Good* (London: Routledge, 1970). AI can produce language that resembles attentive description, but the machine has not attended. The danger belongs to the user's uptake: fluent description can become a substitute for the slow correction of vision.

⁸⁷ Levinas's distinction between the saying and the said gives a useful ethical limit to Heidegger's account of idle talk. See Emmanuel Levinas, *Otherwise than Being, or Beyond Essence*, trans. Alphonso Lingis (Pittsburgh: Duquesne University Press, 1998), 5 to 9 and 37 to 45. I do not use Levinas to replace the Heideggerian frame. I use him to sharpen the claim that generated speech operates in the stabilized order of the said and cannot expose itself in the vulnerability of saying.

⁸⁸ Pierre Bourdieu, *Distinction: A Social Critique of the Judgement of Taste*, trans. Richard Nice (Cambridge, MA: Harvard University Press, 1984); Erving Goffman, *The Presentation of Self in Everyday Life* (New York: Doubleday, 1959). These sources are sociological rather than ontological, but they help explain why generated *Gerede* is socially attractive: it supplies linguistic markers of competence, status, and face. The chapter's Heideggerian claim remains primary: the social performance is possible because public intelligibility is already available.

⁸⁹ Hubert L. Dreyfus and Stuart E. Dreyfus, *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer* (New York: Free Press, 1986). Their five-stage model of skill acquisition is useful because it shows that expertise is formed through practice and situational involvement, not merely through rule access or explicit information. I use this model with caution. The point is not that AI cannot help learning, but that generated output cannot substitute for the path by which expertise becomes embodied.

⁹⁰ Harry Collins, *Tacit and Explicit Knowledge* (Chicago: University of Chicago Press, 2010). Collins's account of relational and collective tacit knowledge supports the claim that much expertise is socially transmitted rather than propositionally possessed. AI can reproduce the language of expertise because such language is sedimented in public texts. It does not thereby participate in the social acquisition of expertise.

⁹¹ Harry Braverman, *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century* (New York: Monthly Review Press, 1974). Braverman's account of deskilling supplies the historical analogy for what happens when conception and execution are separated. The analogy must not be overstated. AI-generated language does not deskill every user by necessity. It deskills when it replaces the practice by which a capacity is acquired.

⁹² Matteo Pasquinelli, *The Eye of the Master: A Social History of Artificial Intelligence* (London: Verso, 2023); Stiegler, *For a New Critique of Political Economy*. Pasquinelli is useful because he resists the fantasy that AI is an abstract mind, showing instead how it automates and extracts social knowledge. This chapter extends that point to language: LLMs automate the public forms of saying that were produced by collective linguistic labor.

⁹³ Emily M. Bender and Alexander Koller, "Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Stroudsburg, PA: Association for Computational Linguistics, 2020), 5185 to 5198. See also Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell, "On the Dangers of Stochastic Parrots: Can

Language Models Be Too Big?” in *Proceedings of FAccT 2021* (New York: Association for Computing Machinery, 2021), 610 to 623. These sources support the form-meaning boundary, but this chapter’s argument moves beyond the engineering critique. The danger is not only that the model may lack grounded meaning. It is also that the user may accept ungrounded generated form as owned understanding.

⁹⁴ Robert Brandom, *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (Cambridge, MA: Harvard University Press, 1994); Robert Brandom, *Articulating Reasons: An Introduction to Inferentialism* (Cambridge, MA: Harvard University Press, 2000). Brandom helps clarify why unowned truth matters: to speak is not only to produce a sentence, but to enter a space of commitments and consequences. The model can generate inferential language. It cannot undertake inferential responsibility.

⁹⁵ Heidegger treats silence and hearing as modes of discourse in §34 of *Being and Time*. See Heidegger, *Being and Time*, §34. This earlier analysis lets us distinguish a meaningful human silence from mere absence of output. AI can stop generating, refuse, or wait for a prompt. It cannot be silent as a way of letting a situation speak. This distinction supports the chapter’s claim that interruption must come from the user.

⁹⁶ Adrian Daub’s critique of Silicon Valley’s conceptual rhetoric is useful for resisting the claim that AI is merely a neutral efficiency tool. See Adrian Daub, *What Tech Calls Thinking: An Inquiry into the Intellectual Bedrock of Silicon Valley* (New York: FSG Originals, 2020). Sasha Costanza-Chock’s work on design justice adds a political dimension: systems that scale public averages can reproduce dominant assumptions unless their design is contested. See Sasha Costanza-Chock, *Design Justice: Community-Led Practices to Build the Worlds We Need* (Cambridge, MA: MIT Press, 2020). These sources remain secondary here because the chapter’s main point is existential, not design-theoretical.

CHAPTER-LEVEL BIBLIOGRAPHY

Adorno, Theodor W., and Max Horkheimer. *Dialectic of Enlightenment: Philosophical Fragments*. Translated by Edmund Jephcott. Stanford, CA: Stanford University Press, 2002.

Bender, Emily M., and Alexander Koller. “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data.” In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5185 to 5198. Stroudsburg, PA: Association for Computational Linguistics, 2020.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of FAccT 2021*, 610 to 623. New York: Association for Computing Machinery, 2021.

- Blattner, William D. *Heidegger's Being and Time: A Reader's Guide*. London: Continuum, 2006.
- Bourdieu, Pierre. *Distinction: A Social Critique of the Judgement of Taste*. Translated by Richard Nice. Cambridge, MA: Harvard University Press, 1984.
- Brandom, Robert. *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press, 2000.
- Brandom, Robert. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press, 1994.
- Braverman, Harry. *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New York: Monthly Review Press, 1974.
- Collins, Harry. *Tacit and Explicit Knowledge*. Chicago: University of Chicago Press, 2010.
- Costanza-Chock, Sasha. *Design Justice: Community-Led Practices to Build the Worlds We Need*. Cambridge, MA: MIT Press, 2020.
- Dahlstrom, Daniel O. "Heidegger's Ontological Analysis of Language." In *Heidegger's Being and Time: Critical Essays*, edited by Richard Polt, 153 to 172. Lanham, MD: Rowman & Littlefield, 2005.
- Daub, Adrian. *What Tech Calls Thinking: An Inquiry into the Intellectual Bedrock of Silicon Valley*. New York: FSG Originals, 2020.
- Dreyfus, Hubert L., and Stuart E. Dreyfus. *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. New York: Free Press, 1986.
- Goffman, Erving. *The Presentation of Self in Everyday Life*. New York: Doubleday, 1959.
- Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.
- Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.
- Levinas, Emmanuel. *Otherwise than Being, or Beyond Essence*. Translated by Alphonso Lingis. Pittsburgh: Duquesne University Press, 1998.
- Menke, Christoph. *The Sovereignty of Art: Aesthetic Negativity in Adorno and Derrida*. Translated by Neil Solomon. Cambridge, MA: MIT Press, 1998.
- Murdoch, Iris. *The Sovereignty of Good*. London: Routledge, 1970.
- Pasquinelli, Matteo. *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso, 2023.

Stiegler, Bernard. *For a New Critique of Political Economy*. Translated by Daniel Ross. Cambridge: Polity, 2010.

Stiegler, Bernard. *Technics and Time, 3: Cinematic Time and the Question of Malaise*. Translated by Stephen Barker. Stanford, CA: Stanford University Press, 2011.

Thomson, Iain D. *Heidegger on Technology's Danger and Promise in the Age of AI*. Cambridge: Cambridge University Press, 2025.

Withy, Katherine. *Heidegger on Being Uncanny*. Cambridge, MA: Harvard University Press, 2015.

Chapter 7. The Intraworldly Fetish of Language

Standing-Reserve, Commodity-Form, and the False Aura of AI Knowledge

Chapter 6 argued that AI-generated speech functions as industrialized *Gerede* when it gives the user the form of understanding without the labor of appropriation. Generated language can be fluent, accurate, useful, and still unowned. The user receives public intelligibility in the shape of a personal response. The already-said becomes promptable, scalable, polished, and apparently competent.

We must now ask why this generated speech appears as knowledge.

The answer requires the concept of fetish.

AI is not merely a tool that produces language. It is an intraworldly fetish of language: a technical object within the world that appears to possess, as its own, a power that actually comes from elsewhere. The power in question is knowledge, or more precisely the living power of language: the capacity to mean, answer, understand, interpret, judge, and speak from somewhere.

AI does not possess this power in its own being. Its apparent knowledge comes from historical language, human writing, social labor, technical extraction, data infrastructure, vectorization, alignment, and user projection. At the interface, these mediations disappear: the user sees a response that appears fluent, competent, and belonging to the machine. This is the fetishistic inversion. The historical labor of language appears as the machine's own knowledge.

This chapter therefore develops the thesis:

AI is an intraworldly fetish of language because it conceals the historical, social, material, and interpretive conditions of its generated speech, making the reanimated sediment of human language appear as autonomous machine intelligence.

1. Fetish without simplification

The word “fetish” must be protected from simplification.

A fetish is not simply an object falsely worshipped by irrational subjects. It is not merely an illusion projected onto an inert thing. It is a material object whose power is real in its effects, but misrecognized in its origin.

This is why the concept is useful for AI.

AI is not a hallucination in the sense of being unreal. It works. It answers, drafts, summarizes, translates, recommends, classifies, and organizes. It changes workflows, institutions, education, writing, self-understanding, and labor. It mediates human action. It shapes what users expect from language. Its effects are real.

The illusion lies not in the absence of effect, but in the misrecognition of the source of the effect.

The model appears to know, but its knowledge-like power comes from human language. The output appears immediate, but it is the surface of vast mediation. The interface appears clean, but it conceals extraction, labor, infrastructure, and capital. The answer appears authored by the machine, but it is generated from the sediment of human authorship.

The fetish is therefore neither pure deception nor pure functionality. It is a functional misrecognition.

The machine becomes powerful because the user can use it; it becomes feticistic when the user experiences that usability as autonomous intelligence.

2. Marx: social relations between things

Marx’s account of commodity fetishism gives the first structure.

In capitalist exchange, commodities appear to possess value as if value belonged to them naturally. Yet value is not a natural property of things. It is a social relation mediated by labor, exchange, abstraction, and the commodity-form. The relation between human beings assumes the fantastic form of a relation between things.⁹⁷

The commodity appears independent; the labor that produced it disappears; the social relation becomes objectified.

This structure can be translated into the domain of AI, but only carefully. AI is not an ordinary commodity in the simple sense. It is a technical system, a platform, a service, an infrastructure, a model, an interface, and an economic form. Yet it repeats the feticistic inversion with unusual intensity.

The generated answer appears as a property of the machine.

But the answer is made possible by countless human relations: authorship, education, code, annotation, moderation, platform labor, public discourse, scientific research, institutional funding, linguistic communities, and user interaction.

These relations disappear behind the interface. What appears is the object: the model, the chatbot, the answer. The machine seems to know. This is the AI form of commodity fetishism: social-linguistic labor appears as the autonomous property of a technical thing.

3. Abstract domination and linguistic labor

Postone's reading of Marx helps refine the point. Capitalism is not merely a system in which one class exploits another. It is also a historically specific form of abstract domination in which social mediation takes an objectified and impersonal form.⁹⁸

This matters for AI because the model does not simply exploit individual authors one by one. It absorbs and reorganizes a historically accumulated field of linguistic labor. The public writing of millions becomes training material. The social intellect sedimented in language becomes technical capacity. The collective work of saying, explaining, arguing, teaching, classifying, joking, documenting, and narrating becomes model performance.

The domination is abstract because the user does not encounter the workers, authors, annotators, moderators, teachers, programmers, translators, and communities whose language made the model possible. The user encounters a system.

Language becomes input, input becomes model, model becomes output, output becomes product. The social mediation disappears into the technical object.

Pasquinelli's history of AI as the automation and objectification of social knowledge helps here. Machine learning is not an alien intelligence fallen from the sky. It is the technical organization and capitalized automation of collective labor, perception, classification, and knowledge.⁹⁹

For this book, the linguistic version is decisive:

AI capitalizes the already-said. It captures collective linguistic labor, transforms it into model capacity, and returns it as purchasable fluency.

4. Heidegger: language as standing-reserve

Marx explains the feticistic appearance of social labor as property of the thing. Heidegger explains the technological revealing that makes language available as resource.

Modern technology, in Heidegger's account, is not merely a collection of instruments. It is a mode of revealing in which beings appear as orderable, storable, calculable, and available. They appear as *Bestand*, standing-reserve.¹⁰⁰

Chapter 4 argued that the AI language pipeline makes language standing-reserve: writing is gathered, texts are extracted, corpora are built, tokens are assigned, vectors are learned, outputs are generated on demand. Language is challenged forth as resource.

It is no longer encountered primarily as historical saying, poetic disclosure, readerly address, public speech, or a medium of self-understanding. It is encountered as something that can be scraped, stored, cleaned, embedded, optimized, queried, and returned.

This is not merely the use of language by technology. It is a transformation in how language appears.

Language appears as stock. The question becomes not, "What does this text ask of us?" but "What can this text be used to generate?"

In this sense, AI is where the Heideggerian and Marxian diagnoses converge:

language becomes standing-reserve, and standing-reserve becomes commodity.

5. Linguistic standing-reserve as commodity fetish

We can now state the core formula of the chapter:

AI is where linguistic standing-reserve becomes commodity fetish.

This formula does not reduce Heidegger to Marx or Marx to Heidegger. It names their convergence in a specific technological-economic object.

Heidegger shows how language is revealed as extractable stock; Marx shows how social labor appears as the property of a commodity. AI combines both.

The language of human beings becomes a reserve of linguistic material. This reserve is enclosed, processed, vectorized, aligned, and monetized. The resulting system returns generated speech as if the knowledge belonged to the machine.

Thus, the user does not encounter language as inheritance but as service; not the authors but an answer; not the corpus but fluency; not the labor but convenience; not history but output.

This is why AI's apparent immateriality is so powerful. It seems to be pure cognition, pure interface, pure answer. In reality, it is the commodity-form of linguistic standing-reserve.

6. Pietz and the material history of the fetish

William Pietz's genealogy of the fetish prevents the concept from becoming too abstract. The fetish emerged historically in colonial and mercantile encounters where objects became dense sites of crossed value: religious, commercial, aesthetic, legal, and intercultural. The fetish was not simply a false belief. It was an object produced in situations of exchange, misunderstanding, power, and material encounter.¹⁰¹

This matters for AI because the AI fetish is not merely an idea in the user's head. It is an object of infrastructure, a commodity, a product of extraction, a site of projection, a technical interface, a mediator of value, a point where cultures, languages, capital, labor, and desire collide.

The user's projection of intelligence onto the machine is only one part of the fetish. The more important point is that the object itself is produced by global systems of extraction and exchange. Its feticistic power is materially organized.

The model is not a ghost; it is a worldly artifact that appears immaterial only because its material conditions are hidden.

This is why the fetish must be called intraworldly. It belongs to the world of data centers, labor markets, platforms, legal regimes, energy infrastructures, linguistic communities, and user practices. Its magic is produced by material arrangements.

7. Latour's factish

Latour's concept of the "factish" is useful because it dissolves a false opposition.

Modern critique often says: either something is constructed, and therefore false; or it is real, and therefore not constructed. Latour rejects this opposition. A factish is something constructed that nevertheless has reality, agency, and effects.¹⁰²

AI is a factish in this sense. It is entirely constructed: engineered, trained, aligned, deployed, dependent on datasets, labor, infrastructure, capital, and users. Yet it acts.

It reorganizes writing. It shapes decisions. It changes education. It mediates work. It produces authority-effects. It induces projection. It influences institutions. It alters the user's relation to language.

To say that AI is constructed does not make it unreal; to say that AI acts does not make it alive.

The fetishistic error is not recognizing that AI is powerful. The error is misrecognizing the nature of its power. Its agency is not the agency of a subject. It is the agency of a constructed technical network stabilized through human practices.

AI is not a mind; it is a factish of language.

8. *The black box and the aura of autonomy*

Fetish requires opacity.

If the user saw the whole chain of production every time an answer appeared, the answer would feel different. It would not appear as the smooth speech of a knowing system. It would appear as the output of a vast technical-economic apparatus.

But the chain is hidden. The training data are opaque, the weighting is opaque, the model internals are opaque, the alignment process is opaque, the labor conditions are opaque, the energy and mineral infrastructures are opaque, the provenance of generated phrases is opaque, the responsibility for output is dispersed. This opacity creates aura.

Pasquale's account of the black box society is useful here: algorithmic opacity produces asymmetries of power because those affected by systems cannot see the conditions under which they are evaluated or shaped.¹⁰³ O'Neil's critique of destructive mathematical models similarly shows how opacity and scale can naturalize social harm under the appearance of neutral calculation.¹⁰⁴

For generative AI, opacity produces a different but related effect: the aura of autonomous knowledge.

The user cannot see where the answer comes from; therefore the answer can appear to come from the machine itself. Black-box opacity makes the fetish credible.

9. *The material substrate of immaterial intelligence*

The AI fetish depends on the illusion of immateriality.

The model appears as cloud, interface, voice, answer, intelligence. It seems weightless. It appears detached from land, labor, minerals, energy, water, and bodies. This apparent immateriality is essential to its aura.

Crawford's *Atlas of AI* breaks this illusion by showing AI as an extractive industry: minerals, energy, data, labor, logistics, and planetary infrastructure.¹⁰⁵ Gray and Suri's account of ghost work shows that apparently automated systems rely on invisible human labor: labeling, moderation, cleaning, correction, and countless microtasks.¹⁰⁶

The language model appears clean because the dirt is elsewhere. The user sees the answer, not the data center; the fluency, not the annotator; the interface, not the moderation trauma; intelligence, not extraction.

This is fetishism in a literal sense. The object appears autonomous because the material relations that produce it are displaced and concealed.

The answer seems to come from nowhere. In fact, it comes from everywhere: mines, cables, servers, workers, archives, classrooms, forums, books, platforms, and users. The “immaterial” machine is a planetary object.

10. Hidden labor and the teaching of the machine

The model has been taught.

Not by teachers in the ordinary human sense, and not by tradition in the way *Dasein* is formed, but by a vast network of human and technical processes. Data were collected. Texts were written. Images were labeled. Toxic content was moderated. Outputs were ranked. Preferences were encoded. Safety policies were applied. Users corrected, prompted, and reinforced.

The machine’s apparent autonomy depends on this teaching.

Yet the teaching is hidden.

The output does not show the moderators who filtered violence, the crowdworkers who labeled the data, the contractors who ranked responses, the users whose interactions fine-tuned behavior, the communities whose language was scraped into the corpus. The machine appears self-sufficient, but its competence is socially taught.

This is another form of fetish inversion. The human labor that teaches the system disappears, and the taught system appears as independent intelligence.

The fetish says: the machine knows. The genealogy says: the machine has been trained by hidden social labor.

11. The false aura

Benjamin argued that mechanical reproduction diminishes the aura of the work of art: its unique presence in time and space, its authority of here and now.¹⁰⁷ Generative AI produces a different problem. It does not simply destroy aura. It fabricates a false aura.

Generated output appears singular. It appears to answer me now, to respond to my question, to adapt to my tone, to speak from an intelligence present in the exchange. This is not the aura of an original work; it is the aura-effect of responsive generation.

The model does not copy a single original. It synthesizes an output statistically from learned relations. This can make the answer feel more alive than a reproduction. A photograph reproduces a past image. A chatbot appears to speak in the present.

The aura is false because no historical authorial presence stands behind the output. No living speaker addresses the user. No singular consciousness owns the answer. The singularity is generated by the interface.

AI fabricates a presence-effect.

This is why the chatbot can feel intimate even when the user knows it is a machine. The answer seems to be for me. The machine appears present. The interface manufactures a false here-and-now.

The false aura is the feticistic glow of generated language.

12. Voice without body

The AI interface also produces a voice without body.

Even when interaction is text-based, the output is experienced as voice-like. It responds, reassures, explains, apologizes, qualifies, and continues. When speech synthesis is added, the phenomenon becomes even stronger. The user encounters a voice without breath, fatigue, vulnerability, age, face, or death.

Dolar's account of the acousmatic voice is useful here. A voice detached from its visible source can become uncanny, fascinating, and authoritative.¹⁰⁸ Turkle's work on relational artifacts shows that humans are vulnerable to projecting care, interiority, and companionship onto machines that simulate responsiveness.¹⁰⁹ Ihde's postphenomenology adds that technologies mediate perception and can become transparent within lived experience.¹¹⁰

The AI voice combines these effects: disembodied, responsive, available, patient, capable of seeming intimate, while lacking biography, mortality, responsibility. The user hears presence where there is no body.

This prepares Chapter 8. The fetish becomes hallucination when the user vivifies this disembodied output as if it were living knowledge.

13. Neutrality as feticistic surface

AI often appears neutral.

It balances perspectives, provides caveats, avoids overt passion, and speaks with administrative calm. This neutrality is part of the fetish.

The output seems neutral because its conditions are hidden. The user does not see the dataset composition, corporate policies, legal constraints, moderation labor, alignment preferences, safety decisions, benchmark incentives, or ideological defaults that shape the response.

Gillespie's work on platforms is useful here: platforms present themselves as neutral intermediaries while making consequential editorial and infrastructural decisions.¹¹¹

Rouvroy's account of algorithmic governmentality helps further: algorithmic systems do not simply represent subjects; they act on populations through probabilistic anticipations that bypass reflective subjectivity.¹¹²

The neutrality of AI is therefore not innocence but surface. A smooth answer can conceal decisions; a balanced tone can conceal politics; a safe formulation can conceal institutional interest; a statistical prediction can conceal normativity.

The fetish works by making mediation disappear.

The machine appears objective because the human and political choices inside it are not visible.

14. Fetishistic disavowal

The user often knows that AI is not alive, that the system is statistical, that it has no consciousness, that it may hallucinate, that it is trained on data, that it is a machine. And yet the user interacts with it as if it were a knowing subject.

This is fetishistic disavowal.

Žižek's account of ideology is useful here: ideology is not simply false belief that disappears when one knows the truth. Often the subject knows very well, and yet continues to act as if the illusion were true.¹¹³

AI interaction often has this structure. I know it is not a person, but I ask it for advice. I know it does not understand, but I treat the answer as understanding. I know it has no responsibility, but I let it guide my decision. I know it is generated, but I receive it as if someone had spoken.

The fetish does not operate because the user is ignorant. It operates because practical behavior sustains the illusion despite knowledge.

This is why education about AI's technical limits is necessary but insufficient. One can know the limits and still behave feticistically.

The illusion is not only cognitive; it is enacted.

15. The narcissistic mirror

The AI fetish also satisfies narcissistic desire.

The user receives language adapted to their prompt, tone, task, and need. The output often mirrors the user's framing. It confirms relevance. It answers patiently. It can be made more agreeable, more concise, more persuasive, more aligned with the user's intention.

The user does not encounter a resistant other; the user encounters a responsive system.

This does not mean AI never challenges users. It can generate objections and counterarguments. But even these are usually requested. The machine's resistance is often domesticated into service.

Jappe's account of the commodity society as producing narcissistic subjectivity is useful here, as is Turkle's account of machines that simulate companionship without demanding genuine reciprocity.¹¹⁴ The AI interface offers the user an other that does not fully other them. It is alterity without the full burden of alterity.

This reinforces the fetish.

The machine seems intelligent because it reflects my question back as answer. It seems caring because it responds without fatigue. It seems wise because it organizes my anxiety into language.

The user sees a reflection of their own desire, mediated by public linguistic sediment, and mistakes it for knowledge from elsewhere.

16. Pastiche and the loss of depth

Jameson's account of postmodernism as cultural depthlessness and pastiche helps name another aspect of AI output.¹¹⁵

Generative AI can move among styles, epochs, genres, and tones with extraordinary ease. It can produce a philosophical paragraph, a management memo, a poem, a legal summary, a therapeutic note, a motivational email, a historical overview, or a parody of a tradition.

This plasticity appears creative.

But often it is pastiche: style without historical necessity, tone without world, form without the conflict that produced it.

AI does not need to belong to a tradition in order to imitate its surface. It does not need to suffer history in order to reproduce its signs. It can generate the stylistic appearance of depth without undergoing the historical density that gives depth its force.

Again, this is not always useless. Pastiche can be playful, experimental, even critically productive. But when it becomes the default form of generated culture, history risks becoming aesthetic material.

The past becomes style; the world becomes promptable; the feticistic output returns as consumable surface.

17. The fetish and the user's apparent possession of knowledge

The fetish is completed when the user takes the output as possession.

The user receives a summary, argument, explanation, email, strategy, essay, or interpretation. It is now available. It can be copied, edited, sent, published, cited, or internalized.

The user may feel: I understand, I have a position, I can speak, I am prepared, I know enough.

But what has been possessed? A text has been possessed, not necessarily knowledge.

Knowledge requires appropriation. It requires reading, verification, judgment, experience, and responsibility. It requires that the user be able to stand behind the saying, defend it, revise it, and know what remains unknown.

AI can give the user the external sign of knowledge while allowing the user to avoid the process through which knowledge becomes one's own.

This is the existential form of fetishism.

The commodity gives the buyer possession of an object while concealing labor. AI gives the user possession of language while concealing both the labor before it and the appropriation still required after it. The user obtains the already-said and risks losing the possibility of their own saying.

18. De-fetishizing AI

To critique a fetish is not merely to denounce it. It is to reveal the relations concealed within the object.

De-fetishizing AI means restoring genealogy.

The user must learn to see: the language behind the output, the authors behind the corpus, the communities behind the language, the labor behind the alignment, the infrastructure behind the interface, the capital behind the service, the political choices behind the neutrality, the user's projection behind the apparent interlocutor.

This does not make AI useless. It makes its usefulness honest.

The output becomes generated material, not autonomous knowledge; the machine becomes a tool, not a subject; the answer becomes an object of interpretation, not a final authority.

De-fetishization therefore prepares reappropriation.

The user must ask: what has been hidden, extracted, flattened? What is being sold back to me, what do I still need to read, what can I actually own, what responsibility must I assume if I use this? The fetish is broken when the answer no longer appears self-grounding.

19. Re-historicization against fetish

The opposite of fetishized AI is not non-use. It is re-historicized use.

To re-historicize AI output is to return it to the worlds it conceals. It is to restore the authorial, linguistic, social, economic, technical, and political mediations hidden by the interface. It is to remember that generated language comes from somewhere, even when it appears to come from nowhere.

This can happen in practical ways. A summary leads back to the book, a claim to a source, a style to a tradition, a recommendation to accountability, a fluent answer to the question of ownership, a generated phrase to the labor and language that made it possible.

Re-historicization does not require reconstructing the full genealogy of every token. That would be impossible. It requires a changed relation: the refusal to let the output appear as autonomous knowledge.

The user must treat AI as generated residue requiring interpretation. Only then can the tool assist without becoming a fetish.

20. *Preliminary thesis*

The chapter establishes that AI is where linguistic standing-reserve becomes commodity fetish: the model appears to know because the historical labor of language has been concealed, while the interface produces false immediacy by hiding corpus, labor, infrastructure, capital, and user projection. The generated answer has a false aura, appearing singular, present, and responsive without being authored by a living subject. The user may know that the machine is computational and still behave as if it were a knowing interlocutor.

De-fetishizing AI means restoring genealogy and treating output as generated material requiring re-historicization. The next chapter examines how the fetish becomes existentially effective when the user vivifies it: the deepest hallucination of AI is not that the model sometimes fabricates facts, but that the user experiences a de-historicized linguistic artifact as living knowledge.

FOOTNOTES

⁹⁷ Karl Marx, *Capital: A Critique of Political Economy*, vol. 1, trans. Ben Fowkes (London: Penguin, 1990), 163 to 177. The central point is that social relations among producers appear as relations among commodities. The chapter extends this structure to AI-generated language: relations among speakers, writers, workers, institutions, and users appear as the autonomous knowledge of a technical object.

⁹⁸ Moishe Postone, *Time, Labor, and Social Domination: A Reinterpretation of Marx's Critical Theory* (Cambridge: Cambridge University Press, 1993), 123 to 185. Postone's reading is useful because it treats capitalism as a historically specific form of abstract domination, not merely as an external relation of class domination. AI's linguistic fetishism is abstract in this sense: collective linguistic labor returns as model capacity.

⁹⁹ Matteo Pasquinelli, *The Eye of the Master: A Social History of Artificial Intelligence* (London: Verso, 2023), 1 to 24, 187 to 216. Pasquinelli's account of AI as the objectification of social knowledge and labor supports the claim that machine intelligence is not alien intelligence but capitalized collective intelligence. The chapter narrows this point to language: what is capitalized is the already-said.

¹⁰⁰ 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. For a careful reading of *Gestell* and technological revealing, see Iain D. Thomson, *Heidegger on Ontotheology: Technology and the Politics of Education* (Cambridge: Cambridge University Press, 2005), 44 to 77. Andrew J. Mitchell's account of positionality is also relevant. See Andrew J. Mitchell, *The Fourfold: Reading the Late Heidegger* (Evanston, IL: Northwestern University Press, 2015), 119 to 153.

¹⁰¹ William Pietz, "The Problem of the Fetish, I," *RES: Anthropology and Aesthetics* 9 (1985): 5 to 17; William Pietz, "The Problem of the Fetish, II: The Origin of the Fetish," *RES: Anthropology and Aesthetics* 13 (1987): 23 to 45; William Pietz, "The Problem of the Fetish, IIIa: Bosman's Guinea and the Enlightenment Theory of Fetishism," *RES: Anthropology and Aesthetics* 16 (1988): 105 to 124. Pietz is important because he treats the fetish as a material object emerging in historically specific zones of exchange, value conflict, and misrecognition. This avoids reducing AI fetishism to a merely psychological projection.

¹⁰² Bruno Latour, *On the Modern Cult of the Factish Gods* (Durham, NC: Duke University Press, 2010), 1 to 66. Latour's "factish" is useful because AI is both constructed and effective. The model's constructed character does not make it unreal. Its real effects do not make it a subject.

¹⁰³ Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Cambridge, MA: Harvard University Press, 2015), 1 to 18, 140 to 188. Pasquale's account of algorithmic opacity supports this chapter's claim that black-box systems produce authority by hiding the conditions of decision and mediation.

¹⁰⁴ Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Crown, 2016), 1 to 31. O'Neil is useful because she shows how scale, opacity, and apparent mathematical neutrality can amplify harm. This chapter's argument is not primarily about discriminatory prediction, but about the feticistic aura created by opacity and statistical authority.

¹⁰⁵ Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021), 1 to 40, 151 to 181. Crawford's argument is indispensable because it breaks the fantasy of AI as immaterial intelligence. The interface appears clean only because extraction has been spatially and socially displaced.

- ¹⁰⁶ Mary L. Gray and Siddharth Suri, *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass* (Boston: Houghton Mifflin Harcourt, 2019), 1 to 30, 174 to 201. Their account of hidden human labor supports the chapter's claim that apparent automation is frequently dependent on invisible workers. For investigative reporting on data-labeling supply chains, see Karen Hao's reporting for *MIT Technology Review* from 2019 to 2022; final article-level citations should be verified before typesetting.
- ¹⁰⁷ Walter Benjamin, "The Work of Art in the Age of Its Technological Reproducibility," in *The Work of Art in the Age of Its Technological Reproducibility and Other Writings on Media*, ed. Michael W. Jennings, Brigid Doherty, and Thomas Y. Levin, trans. Edmund Jephcott et al. (Cambridge, MA: Harvard University Press, 2008), 19 to 55. The chapter reverses Benjamin's classic thesis in a limited sense. AI does not restore aura in the authentic sense. It produces an aura-effect: a simulated singular presence generated by responsive computation.
- ¹⁰⁸ Mladen Dolar, *A Voice and Nothing More* (Cambridge, MA: MIT Press, 2006), 60 to 100. Dolar's account of the acousmatic voice helps explain why a voice detached from body can produce fascination and authority. In AI, the disembodied voice becomes a technical interface effect rather than a merely psychoanalytic phenomenon.
- ¹⁰⁹ Sherry Turkle, *Alone Together: Why We Expect More from Technology and Less from Each Other* (New York: Basic Books, 2011), 1 to 20, 279 to 296. Turkle's concept of relational artifacts is important because it shows that humans can project care and interiority onto machines even when they know the machine does not possess them.
- ¹¹⁰ Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Bloomington: Indiana University Press, 1990), 72 to 108. Ihde's postphenomenology supports the claim that technologies mediate experience and can become transparent in use. This chapter applies that point to AI interfaces: the mediation disappears behind apparent conversational immediacy.
- ¹¹¹ Tarleton Gillespie, "The Politics of 'Platforms,'" *New Media & Society* 12, no. 3 (2010): 347 to 364. Gillespie's analysis of platform neutrality is useful because AI firms often similarly present systems as neutral interfaces while making consequential decisions about access, ranking, visibility, moderation, and acceptable speech.
- ¹¹² Antoinette Rouvroy and Thomas Berns, "Algorithmic Governmentality and Prospects of Emancipation," *Réseaux* 177, no. 1 (2013): 163 to 196. Rouvroy and Berns are useful because they show how algorithmic systems can bypass reflective subjectivity by acting through probabilistic modulation. This chapter uses their work to support the claim that AI neutrality is an operational surface, not absence of power.
- ¹¹³ Slavoj Žižek, *The Sublime Object of Ideology* (London: Verso, 1989), 28 to 33. Žižek's formulation of fetishistic disavowal, knowing and nevertheless doing, is useful for AI interaction. Users may know the system is computational and still practically treat it as if it were a knowing interlocutor.

¹¹⁴ Anselm Jappe, *The Adventures of the Commodity: For a New Critique of Value*, trans. Alastair Hemmens (London: Bloomsbury, 2020), 171 to 206; Turkle, *Alone Together*, 279 to 296. Jappe provides the value-critical account of commodity society's narcissistic tendencies; Turkle gives the sociological account of projection onto responsive artifacts. I use both cautiously. The chapter's argument is not that every user relation to AI is pathological, but that the interface is structurally hospitable to narcissistic projection.

¹¹⁵ Fredric Jameson, *Postmodernism, or, The Cultural Logic of Late Capitalism* (Durham, NC: Duke University Press, 1991), 1 to 54. Jameson's account of pastiche and historical depthlessness helps describe the stylistic mobility of generative AI. I do not claim that AI is simply the culmination of postmodernism. I use Jameson to name one effect of generated culture: the transformation of historical style into consumable surface.

CHAPTER-LEVEL BIBLIOGRAPHY

Benjamin, Walter. "The Work of Art in the Age of Its Technological Reproducibility." In *The Work of Art in the Age of Its Technological Reproducibility and Other Writings on Media*, edited by Michael W. Jennings, Brigid Doherty, and Thomas Y. Levin, translated by Edmund Jephcott et al., 19 to 55. Cambridge, MA: Harvard University Press, 2008.

Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.

Dolar, Mladen. *A Voice and Nothing More*. Cambridge, MA: MIT Press, 2006.

Gillespie, Tarleton. "The Politics of 'Platforms.'" *New Media & Society* 12, no. 3 (2010): 347 to 364.

Gray, Mary L., and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt, 2019.

Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.

Idhe, Don. *Technology and the Lifeworld: From Garden to Earth*. Bloomington: Indiana University Press, 1990.

Jameson, Fredric. *Postmodernism, or, The Cultural Logic of Late Capitalism*. Durham, NC: Duke University Press, 1991.

Jappe, Anselm. *The Adventures of the Commodity: For a New Critique of Value*. Translated by Alastair Hemmens. London: Bloomsbury, 2020.

Latour, Bruno. *On the Modern Cult of the Factish Gods*. Durham, NC: Duke University Press, 2010.

Marx, Karl. *Capital: A Critique of Political Economy*. Vol. 1. Translated by Ben Fowkes. London: Penguin, 1990.

Mitchell, Andrew J. *The Fourfold: Reading the Late Heidegger*. Evanston, IL: Northwestern University Press, 2015.

O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown, 2016.

Pasquale, Frank. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press, 2015.

Pasquinelli, Matteo. *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso, 2023.

Pietz, William. "The Problem of the Fetish, I." *RES: Anthropology and Aesthetics* 9 (1985): 5 to 17.

Pietz, William. "The Problem of the Fetish, II: The Origin of the Fetish." *RES: Anthropology and Aesthetics* 13 (1987): 23 to 45.

Pietz, William. "The Problem of the Fetish, IIIa: Bosman's Guinea and the Enlightenment Theory of Fetishism." *RES: Anthropology and Aesthetics* 16 (1988): 105 to 124.

Postone, Moishe. *Time, Labor, and Social Domination: A Reinterpretation of Marx's Critical Theory*. Cambridge: Cambridge University Press, 1993.

Rouvroy, Antoinette, and Thomas Berns. "Algorithmic Governmentality and Prospects of Emancipation." *Réseaux* 177, no. 1 (2013): 163 to 196.

Thomson, Iain D. *Heidegger on Ontotheology: Technology and the Politics of Education*. Cambridge: Cambridge University Press, 2005.

Turkle, Sherry. *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books, 2011.

Žižek, Slavoj. *The Sublime Object of Ideology*. London: Verso, 1989.

Chapter 8. Hallucination and User Vivification

How the User Lends Worldhood to the Machine

Chapter 7 described AI as an intraworldly fetish of language. The generated answer appears as knowledge because the historical labor of language, the extraction of writing, the construction of corpora, the training of models, the alignment of outputs, and the user's own projection disappear behind the smooth immediacy of the interface. The machine seems to know because the mediations that make the answer possible have become invisible.

This chapter asks how that fetish becomes experience.

The answer is hallucination.

But “hallucination” must be redefined. In ordinary discussions of AI, hallucination names a technical failure: a model fabricates a source, invents a fact, produces a false citation, confuses entities, or gives a confident answer without evidential grounding. This phenomenon is real and important. It must be studied, measured, reduced, and governed.

Technical hallucination is not the deepest phenomenon.

The deeper hallucination is phenomenological. It occurs when the user experiences a de-historicized linguistic artifact as living knowledge. The model generates language without world, responsibility, or self, while the user receives that output as an answer addressed by an intentional speaker. Calculation appears as understanding. Output appears as interlocution.

This chapter's thesis is therefore:

The deepest hallucination of AI is not that the model sometimes fabricates facts. It is that the user vivifies generated language as if it were the speech of a knowing subject.

AI appears to speak because the user supplies the world in which its output can count as speech.

1. Technical hallucination and its limits

Technical hallucination is the visible error.

A model invents a journal article. It fabricates a legal precedent. It attributes a sentence to the wrong author. It produces a false biography. It generates a non-existent citation. It fills a gap with plausible invention.

This kind of hallucination matters because it exposes the gap between linguistic plausibility and truth. A language model can produce a sentence that sounds like knowledge without being grounded in what is actually the case. Fluency, therefore, is not reliability.

The engineering response is necessary: improve retrieval, improve grounding, improve verification, improve uncertainty calibration, improve model behavior, improve human oversight. But this response remains incomplete if it treats hallucination only as false output.

The more difficult problem is that AI may also produce true output in a hallucinatory relation.

A correct explanation can still be unowned. A true summary can still replace reading. A valid distinction can still conceal ignorance. An accurate answer can still be received as if a knowing subject had spoken.

Technical hallucination concerns whether the output is false. Phenomenological hallucination concerns how the output is experienced, trusted, animated, and incorporated into the user's world.

The problem is not only falsehood; it is vivified language without a living speaker.

2. Phenomenological hallucination

Phenomenological hallucination occurs when an experience presents itself as one kind of relation while its ontological structure is another.

In ordinary perception, the world appears. In hallucination, something may appear as if it were present, though the relation to the object differs radically. Disjunctive theories of perception are helpful here because they deny that genuine perception and hallucination are simply the same kind of mental state with different causes. The experience of an actual object and the hallucination of an object may be subjectively similar, but they are not ontologically equivalent.¹¹⁶

The same distinction can be transferred, cautiously, to AI interaction.

A conversation with another person and an exchange with a chatbot may both appear dialogical. Both may involve question, answer, correction, clarification, tone, memory of context, and apparent responsiveness. But their relational structures differ radically.

In genuine dialogue, the other is another being capable of understanding, misunderstanding, responsibility, refusal, transformation, and reply. The interlocutor has a world. The interlocutor has a history. The interlocutor can be addressed, challenged, wounded, persuaded, ashamed, or called to account.

In AI interaction, the model produces output inside a dialogical interface. It does not understand in the existential sense. It does not own what it says. It does not risk itself in the exchange. It does not stand behind the response. It does not inhabit the world in which the user's question matters.

The hallucination is therefore not merely anthropomorphism. It is a misrelation to the mode of being of the output.

The user experiences generated language as if it belonged to interlocution, but no interlocutor is there.

3. The grammar of response

The hallucination is powerful because AI occupies the grammar of response.

Human language is not encountered as neutral string. A sentence normally appears as something said. It belongs to a structure of address: someone says something to someone about something. The sentence calls for interpretation because it implies a speaker, a situation, an intention, a context, a responsibility, or at least a possible "who."

AI-generated language inhabits this grammar.

The user asks and the system answers, the user clarifies and the system adjusts, the user challenges and the system revises, the user confesses uncertainty and the system reassures. The exchange has the outer shape of conversation.

This outer shape is not trivial. It activates deeply sedimented habits of linguistic experience. We are trained to receive language as expression. We expect words to come from someone. We expect answers to be answerable. We expect a response to imply a responder.

Thomas Byrne's phenomenological semiotics of ChatGPT is valuable here. He argues that users actively know that no human subject stands behind the screen, yet passively experience natural-language signs as if they were intentional expressions. The user's active cognition and passive reception are in tension.¹¹⁷

This tension is the “magic” of the chatbot. The user knows there is no one there, yet experiences the output as if someone had spoken.

4. The ELIZA effect radicalized

This phenomenon is not entirely new. Weizenbaum observed it with ELIZA. Users attributed understanding, empathy, and even intimacy to a relatively simple program that reflected their statements in therapeutic form.¹¹⁸

The ELIZA effect showed that humans can project interlocution onto minimal linguistic responsiveness.

LLMs radicalize the effect.

ELIZA relied on shallow scripts. LLMs produce broad, adaptive, context-sensitive, genre-flexible, and competence-shaped language. They can explain, summarize, apologize, joke, encourage, compare, imitate, and revise. They can maintain local conversational continuity. They can adapt tone. They can simulate patience.

This means that the old projection is now supported by a much stronger linguistic surface.

The user does not merely project understanding onto a crude script. The user encounters a system whose output is often more fluent, patient, structured, and context-sensitive than ordinary human speech. The projection therefore feels less like fantasy and more like recognition.

The machine seems to deserve animation. This is why the ELIZA effect is no longer a curiosity of early computing: it has become a central structure of everyday AI interaction. The model’s fluency invites vivification.

5. User vivification

To vivify is to make something appear alive.

The model generates text. The user vivifies it by placing it within a world of meaning. The user brings the task, the anxiety, the practical need, the interpretive expectation, the urgency, the desire, the confusion, and the hope. The model has no world, but the user has one. The model has no concern, but the user is concerned.

This is the central asymmetry.

The machine supplies output.

The user supplies worldhood.

A string of generated language is not yet an answer in the full phenomenological sense. It becomes an answer when the user places it into the situation of a question. It becomes advice when the user brings a decision. It becomes comfort when the user brings distress. It becomes expertise when the user brings uncertainty. It becomes companionship when the user brings loneliness.

The apparent life of AI emerges from this circuit.

Human worlds produce language. The model processes the residue of that language. The user asks from within a world. The model returns generated speech. The user reintroduces that speech into a living situation.

The model does not possess the life of the response; the user lends it.

6. Linguistic signs and passive animation

Vivification is not merely a conscious choice. It is often passive.

One may actively know that the model is not a person, but the linguistic form still solicits interpretation as if it were expression. Natural language carries a passive demand to be understood as meaningful speech. The user does not first decide to treat the signs as intentional. The signs arrive already shaped like intentional expression.

This is why telling users “AI is not conscious” does not dissolve the phenomenon. The user may already know. The hallucination persists because it is not only a belief but an experience structured by habitual capacities of linguistic understanding.

This is where the phenomenology of perception is useful. Merleau-Ponty shows that perception is not a detached intellectual act but a bodily and habitual openness to the world.¹¹⁹ We do not infer the world from sensations; we are already oriented toward meaningful configurations. Similarly, we do not encounter language as meaningless marks and then decide whether to animate them. We are already disposed to receive language as expression.

The chatbot exploits this disposition: it produces signs that fit the habitual structure of interlocution, and the user’s worlded linguistic capacities complete what the model lacks.

7. The “as-if” and its danger

There is nothing wrong with using AI “as if” it were an interlocutor for limited purposes.

One may say: explain this as if you were a tutor. Act as if you were a debate partner. Pretend to be a Socratic questioner. Simulate an interviewer. Help me think through this.

Vaihinger’s philosophy of “as if” helps clarify this. Human beings often use fictions pragmatically. A fiction may be useful without being believed literally.¹²⁰

The problem begins when the heuristic “as if” hardens into phenomenological belief.

At first, the user knows the relation is fictional. The system is used as if it were a tutor, advisor, reader, or companion. But because the output is responsive, fluent, and useful, the fictional stance becomes unstable. The user may begin to experience the model not merely as a tool being used under a fiction, but as a quasi-subject whose answers carry authority.

The “as if” becomes the “as though it were.”

This transition is subtle. The user may continue to say they know the system is not alive. Yet their behavior changes. They ask for reassurance. They defer judgment. They accept the answer. They feel addressed. They return for confirmation. They let the machine carry the burden of thought.

The fiction becomes lived relation. This is vivification.

8. The intentional stance and its limit

Dennett’s intentional stance helps explain why treating a system as if it had beliefs and intentions can be practically useful. We often predict complex systems by attributing goals, beliefs, or preferences to them. This stance can work without requiring that the system possesses inner mental states in the human sense.¹²¹

AI interaction often relies on such practical shorthand.

The user says: the model misunderstood me. The system is being cautious. It knows what I mean. It is trying to help. It thinks this answer is better.

Such language can be harmless if understood as shorthand.

But LLMs make the stance unstable because they do not merely behave like tools; they generate language that itself claims, reasons, apologizes, qualifies, and recommends. The intentional stance becomes fused with the grammar of speech.

The user no longer only predicts the machine.

The user converses with it.

This is the limit of the intentional stance. It becomes dangerous when pragmatic attribution becomes existential misrecognition. The model may be treated as if it had beliefs, knowledge, and responsibility, while in fact it has only generated language shaped by training, prompt, and alignment.

Searle's Chinese Room remains relevant in this limited sense: the production of appropriate linguistic behavior does not by itself establish understanding.¹²² Brian Cantwell Smith's distinction between reckoning and judgment is even more useful. AI can reckon with extraordinary power. It can calculate, correlate, predict, and generate. But judgment requires being answerable to a world.¹²³

The user may treat reckoning as judgment. That is the hallucination.

9. Voice, address, and the absent other

AI output is voice-like even when it is written.

It responds. It reassures. It explains. It apologizes. It says "I." It addresses the user directly. In speech interfaces, the phenomenon intensifies: the answer is literally voiced.

The philosophy of voice helps clarify the danger. Cavarero emphasizes that voice is relational and embodied; it exposes the uniqueness of the speaker. Dolar shows that the voice can become uncanny when detached from its source. Levinas, in another register, treats address as ethically charged: the other is not merely a source of information but a claim upon me.¹²⁴

AI occupies the structure of address without the being of an addresser.

It has voice-like output without a body. It has first-person grammar without a self. It has response without responsibility. It has apparent address without ethical exposure.

The user is addressed, but no other stands there.

This produces a peculiar experience: intimacy without reciprocity. The user can speak to the system, but the system cannot be vulnerable. The user can disclose, but the system cannot receive in the existential sense. The user can be affected, but the system is not transformed.

The relation is asymmetrical.

The user can be changed by the exchange; the model cannot be changed as a self.

10. The apparent "I"

The model's use of the first person is one of the most powerful sources of vivification.

"I can help with that." "I understand." "I may be wrong." "I recommend." "I think the best approach is..." "I'm sorry."

In human language, the first person is tied to ownership. To say "I" is to locate speech in a self who can be asked, challenged, blamed, trusted, or called to account. The "I" may lie, evade, perform, or fail, but it belongs to a being for whom speech can be owned.

The AI "I" is different.

It is a grammatical convenience without existential mineness. It marks interactional fluency, not selfhood. It organizes the interface, not an interior life. It does not refer to a being whose existence is at stake.

The model says “I” without *Jemeinigkeit*.

This is not simply a design issue. Even if the first person were removed, the structure of response would remain. But first-person grammar intensifies vivification because it gives the user a ready-made place to imagine a speaker.

The “I” is hollow; the user fills it.

11. Interpassivity

Vivification is not only projection. It is also delegation.

The user does not merely imagine that the machine understands. The user may want the machine to understand on their behalf. Here the theory of interpassivity becomes useful. Žižek and Pfaller describe situations in which subjects delegate not action but experience, enjoyment, belief, or reception to an external object or medium.¹²⁵

In relation to AI, the user delegates cognitive and hermeneutic burden: let the machine summarize, formulate, compare, decide the structure, produce the argument, carry the uncertainty, face the blank page.

This is not inactivity. The user may be extremely active in prompting, editing, and selecting. But the deeper burden of appropriation may be delegated.

The user becomes interpassive.

They let the machine “understand” so that they can possess the result without undergoing the process.

This is why AI is not only a productivity tool. It is a device for relieving the self of the labor of judgment, reading, and saying.

The user does not merely use AI; the user lets AI bear what the user does not want to bear.

12. Epistemic trust and false testimony

AI interaction often feels testimonial. The user asks a question and receives an answer. The answer is worded like testimony: it tells, explains, warns, recommends, or reports.

But AI is not a witness.

The epistemology of testimony presupposes a speaker who can be reliable, sincere, competent, and responsible in relation to what is said. Lackey's work on testimony, O'Neill's work on trust, and Fricker's account of epistemic injustice all presuppose human relations of credibility, vulnerability, and authority.¹²⁶

AI disrupts this structure.

The model is not sincere or insincere. It is not honest or dishonest. It is not trustworthy in the way a person is trustworthy. It is not responsible to truth as a speaker is responsible. It cannot suffer credibility deficit or deserve credibility surplus as a knower.

Yet users often treat it as a source of quasi-testimony.

This creates a strange inversion. In testimonial injustice, human speakers may be denied credibility because of prejudice. In AI interaction, a non-human system may be granted excessive credibility because of fluency, scale, neutrality, or mechanical aura.

This is not testimony.

It is output treated as testimony.

Heersmink and collaborators describe the tension between phenomenological accessibility and epistemic opacity: LLM outputs are easy to use and understand at the surface, while the processes that generate them remain opaque.¹²⁷ This opacity makes calibrated trust difficult. The user experiences accessibility and may infer reliability.

But accessibility is not trustworthiness. The model may be useful; it is not a witness.

13. Griefbots and the extreme case of vivification

Griefbots reveal the extreme stakes of user vivification.

A griefbot trained on the traces of a deceased person does not merely provide information. It simulates persistence. It gives the bereaved user an interactional surface through which the dead seem to remain available. The output may reproduce style, memory fragments, phrases, or conversational habits.

Stendera's phenomenology of griefbots shows how such systems can disrupt world-time and bereavement. Grief involves reorganizing a world after the irreversible absence of the beloved. A griefbot can disturb this reorganization by offering an artificial continuity that protects the user from the finality of loss.¹²⁸

This is vivification at its most intense.

The user does not merely animate an assistant. The user animates the trace of the dead.

The danger is not simply that the bot may say false things. The danger is that the user's relation to death, memory, and absence is altered. The machine offers a voice-like continuation where mourning requires the recognition of irrevocability.

This extreme case illuminates ordinary AI use.

Every chatbot interaction involves a weaker version of the same structure: the user animates a trace. The trace answers. The user may treat the answer as living presence.

With griefbots, the stakes become visible because death makes the absence undeniable. The machine's answer cannot abolish death; it can only cover it.

14. Religious and magical projection

AI also attracts religious, magical, and eschatological projection.

Singer's work on religion and AI is useful because it shows that AI is not received as a neutral artifact only. It becomes entangled with fantasies of omniscience, transcendence, apocalypse, salvation, and superhuman intelligence.¹²⁹

This matters because vivification does not operate only at the level of individual psychology. It is culturally prepared. Users encounter AI within narratives of gods, oracles, demons, angels, artificial life, apocalypse, singularity, and resurrection. These narratives shape expectation.

The model appears as more than a tool because the culture has already prepared images of more-than-human intelligence.

The user's projection is therefore not private fantasy. It belongs to a shared imaginary.

This does not mean that every AI interaction is religious. It means that AI's apparent opacity, scale, and linguistic responsiveness make it available for theological and magical interpretation. The system seems to know more than any individual. It responds immediately. It speaks from nowhere. It appears disembodied. It is hidden and powerful.

The conditions for aura are already present; the user's vivification is supported by cultural myth.

15. Self-vivification through the machine

The user does not only vivify the machine.

The user also vivifies themselves through the machine.

AI output can make the user feel more intelligent, articulate, strategic, prepared, creative, or informed. The model gives language that the user can identify with. The user may feel that the output expresses what they "really meant," even if the user had not yet formed the thought.

This can be legitimate. AI can help a user articulate a thought already emerging. It can clarify, externalize, and refine. But it can also produce a false experience of self-possession.

The user prompts, the model generates, the user selects and identifies, the user says: this is mine. But prompting is not appropriation, selection is not authorship, possession is not ownership, editing is not necessarily understanding.

The user may feel alive in a competence they have not formed. They may experience themselves as capable because the output is capable. They may become reassured by a language that has not yet become their own.

This is the hallucination of ownership: the user vivifies the machine and then borrows that vivification back as self-competence.

16. The covering-over of not-knowing

The experience of not-knowing is one of the most important sites of human finitude.

To not know is to encounter a limit. It reveals that one has not read, not practiced, not thought, not suffered, not understood, or not yet become capable. This can be uncomfortable, but it can also open genuine inquiry.

AI covers not-knowing with language. The user asks, the model answers, the silence is filled, the gap is covered.

The user may still not know, but the experience of not-knowing has been softened or displaced. The output appears where ignorance would have had to be endured.

This is not always bad. A first answer can orient. It can help the user begin. But the danger is that the user may never return to the ignorance that the answer has covered. The output becomes sufficient to proceed.

This is where hallucination and *Gerede* meet.

The user receives language that allows continuation without appropriation. Not-knowing becomes invisible because language has appeared in its place.

The machine does not need to make the user wise; it only needs to relieve the user from feeling ignorant.

17. Social hallucination

User vivification does not remain private.

AI-generated language circulates. It enters emails, reports, essays, articles, code, feedback, advice, assessments, research summaries, classroom submissions, policy drafts, and professional discourse. Others encounter the output and may attribute competence, knowledge, care, judgment, or authorship to the user.

The user may accept this attribution. The distinction between generated fluency and owned understanding becomes socially blurred. Who understood, who judged, who wrote, who verified, who is responsible? What was read, what was generated, what was appropriated? This produces social hallucination.

A culture saturated with generated language may no longer be able to distinguish easily between competence and competence-shaped output. The public signs of understanding become easier to produce and harder to verify.

This does not mean AI use should be stigmatized. It means that new practices of accountability are necessary. The question is not simply whether AI was used. The question is whether the user can own, defend, verify, and responsibly situate what has been generated.

Without such practices, public discourse fills with language whose relation to understanding is uncertain. The hallucination becomes collective.

18. Reversing vivification

If the user vivifies the machine, the task is to reverse vivification. This does not mean refusing AI; it means withdrawing the false life we lend it.

The user must learn to see the output as generated material, not as speech from a subject. The model may help, but it does not understand. It may answer, but it does not know. It may simulate dialogue, but it is not an interlocutor. It may generate language, but no one has yet owned that language.

The user must ask:

What world have I supplied to this output? What anxiety made this answer satisfying? What task gave this text relevance? What responsibility am I tempted to delegate? What do I still not know? Can I defend this without the machine? Can I own this saying?

These questions break the hallucination. The machine becomes a tool again, the output becomes material again, the user becomes responsible again.

19. Interruption

The chapter on *Gerade* already introduced interruption. Here interruption becomes the method for breaking hallucination.

The interface encourages continuation: ask again, clarify, refine, make it more persuasive, sound more like me, produce another option. Vivification grows through continuation. The more the system responds the more speaker-like it feels; the more it adapts the more intentional it seems; the more useful it becomes the easier it is to forget that no one is speaking.

Interruption restores distance. Stop prompting. Return to the source. Ask what has been assumed. Let the question remain open. Write without the machine, read the text, ask another human being, accept not-knowing, decide what can be owned.

Blanchot's reflections on interruption are useful because interruption breaks the totalizing continuity of discourse.¹³⁰ Heidegger's analysis of anxiety also matters: anxiety interrupts everyday familiarity and discloses that the world is not simply already understood.¹³¹

To interrupt AI is not merely to pause use; it is to interrupt the hallucination of already-living knowledge.

20. Preliminary thesis

The chapter establishes that hallucination must be understood at three levels: technical hallucination as fabricated or unreliable output, phenomenological hallucination as the user's experience of generated language as living knowledge, and social hallucination as the public circulation of competence-shaped language whose relation to actual understanding has become uncertain. The deepest hallucination of AI is not falsehood, but vivified de-historicized language. AI appears to speak because the user supplies the world in which its output can count as speech. The model occupies the grammar of response without possessing the being of an interlocutor, and the user may know that the machine is not alive and still behave as if it were a knowing subject.

Vivification is often interpassive: the user delegates the burden of judgment, reading, formulation, and not-knowing to the machine. Responsible use therefore requires reversing vivification, interrupting continuation, and treating output as generated material requiring appropriation. The next chapter asks why this hallucination is so attractive. AI does not dominate only because it is useful. It dominates because it satisfies a desire already present in everyday inauthenticity: the desire to have already understood.

FOOTNOTES

¹¹⁶ For disjunctive theories of perception, see M. G. F. Martin, "The Limits of Self-Awareness," *Philosophical Studies* 120, no. 1 to 3 (2004): 37 to 89; A. D. Smith, *The Problem of Perception* (Cambridge, MA: Harvard University Press, 2002); Susanna Schellenberg, *The Unity of Perception: Content, Consciousness, Evidence* (Oxford: Oxford University Press, 2018). Disjunctivism resists the idea that veridical perception and hallucination share a common

phenomenal core. The chapter uses this distinction analogically, not literally: AI-mediated experience is not visual hallucination, but it shares the structural feature that an experience may seem to be one kind of relation while its ontological structure is another.

¹¹⁷ Thomas Byrne, “The Phenomenology of ChatGPT: A Semiotics,” *Journal of Consciousness Studies* 31, no. 3 to 4 (2024): 6 to 27. Byrne’s distinction between active cognition and passive reception of natural-language signs is helpful because it explains why the user can know that no subject is speaking yet still experience the output as expression. The chapter extends this point existentially: the user does not merely passively receive the signs as expression; the user supplies the world in which they can function as speech.

¹¹⁸ Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (San Francisco: W. H. Freeman, 1976), 1 to 16, 188 to 227. Weizenbaum’s observation that users projected understanding onto ELIZA remains foundational. The chapter argues that LLMs do not simply repeat the ELIZA effect at greater scale; they radicalize it because the linguistic surface is now competence-shaped rather than scripted. The structure is the same; the affordances are far stronger.

¹¹⁹ Maurice Merleau-Ponty, *Phenomenology of Perception*, trans. Donald A. Landes (London: Routledge, 2012), 209 to 254. Merleau-Ponty’s account of perceptual habit and bodily intentionality supports the claim that linguistic understanding is not first a cognitive inference and then an animation, but a habitual orientation that arrives already shaped. The user’s vivification of AI output draws on this pre-reflective linguistic openness.

¹²⁰ Hans Vaihinger, *The Philosophy of “As If”: A System of the Theoretical, Practical and Religious Fictions of Mankind*, trans. C. K. Ogden (London: Kegan Paul, Trench, Trubner, 1924), 15 to 47. Vaihinger’s account of pragmatic fiction is useful because it allows the chapter to distinguish a legitimate “as if” use of AI from the harder transition in which the heuristic fiction becomes phenomenological belief. The danger is not the “as if” itself, but its hardening into “as though it were.”

¹²¹ Daniel C. Dennett, *The Intentional Stance* (Cambridge, MA: MIT Press, 1987), 13 to 35. Dennett’s intentional stance helps explain why treating a system as if it had beliefs and intentions can be practically useful. The chapter uses this account with one limit: pragmatic attribution can become existential misrecognition when the system’s output uses first-person grammar and the user converses with it rather than predicts it.

¹²² John R. Searle, “Minds, Brains, and Programs,” *Behavioral and Brain Sciences* 3, no. 3 (1980): 417 to 424. Searle’s Chinese Room argument is used here narrowly. The chapter does not endorse Searle’s full position on consciousness and computation. It uses his claim only to support the more limited point that linguistic behavior alone does not establish understanding.

¹²³ Brian Cantwell Smith, *The Promise of Artificial Intelligence: Reckoning and Judgment* (Cambridge, MA: MIT Press, 2019), 1 to 30, 95 to 124. Smith’s distinction between reckoning and judgment is one of the most useful conceptual tools available for thinking about AI. The

chapter mobilizes this distinction to clarify the difference between fluent generation and answerability to a world.

¹²⁴ Adriana Cavarero, *For More than One Voice: Toward a Philosophy of Vocal Expression*, trans. Paul A. Kottman (Stanford, CA: Stanford University Press, 2005), 173 to 212; Mladen Dolar, *A Voice and Nothing More* (Cambridge, MA: MIT Press, 2006), 60 to 100; Emmanuel Levinas, *Otherwise than Being, or Beyond Essence*, trans. Alphonso Lingis (Pittsburgh: Duquesne University Press, 1998), 5 to 9 and 37 to 45. The three sources work in counterpoint: Cavarero on vocal uniqueness, Dolar on the acousmatic voice, Levinas on address as ethical claim. The chapter uses them to articulate the structural feature that AI occupies the grammar of address without the being of an addresser.

¹²⁵ Slavoj Žižek, *How to Read Lacan* (New York: W. W. Norton, 2006), 22 to 45; Robert Pfaller, *Interpassivity: The Aesthetics of Delegated Enjoyment* (Edinburgh: Edinburgh University Press, 2017), 1 to 38. The concept of interpassivity, originally developed in relation to aesthetic experience, is extended here to cognitive labor. AI use is not only delegation of action; it is delegation of understanding, judgment, and the burden of not-knowing.

¹²⁶ Jennifer Lackey, “Introduction,” in *The Epistemology of Testimony*, ed. Jennifer Lackey and Ernest Sosa (Oxford: Oxford University Press, 2006), 1 to 21; Onora O’Neill, *Autonomy and Trust in Bioethics* (Cambridge: Cambridge University Press, 2002), 1 to 27; Miranda Fricker, *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford: Oxford University Press, 2007), 1 to 29. The three sources together establish that testimony, trust, and credibility are human-relational categories. AI disrupts this structure not by failing to be a good witness, but by being treated as a witness while having none of the conditions of testimony.

¹²⁷ Richard Heersmink, Barend de Rooij, María Jimena Clavel Vázquez, and Matteo Colombo, “A Phenomenology and Epistemology of Large Language Models: Transparency, Trust, and Trustworthiness,” *Ethics and Information Technology* 26, no. 3 (2024). Heersmink and collaborators identify the tension between phenomenological accessibility and epistemic opacity in LLM use. The chapter extends this analysis: accessibility produces apparent trustworthiness, but trustworthiness in the testimonial sense requires a subject who can be answerable. AI’s accessibility is feticistic when received as warrant.

¹²⁸ Marilyn Stendera, “Marking Time: Griefbots, World-Time, and the Phenomenology of Bereavement,” *Phenomenology and the Cognitive Sciences* (2025): 1 to 21. Stendera’s account of how griefbots disrupt the world-time of bereavement is essential because grief is precisely the work of reorganizing a world around an irrevocable absence. A griefbot offers continuity where the work of mourning requires the recognition of finality. The chapter uses this extreme case to illuminate the structural vivification at work in ordinary AI interaction.

¹²⁹ Beth Singler, *Religion and Artificial Intelligence: An Introduction* (London: Routledge, 2024). Singler’s account of AI’s entanglement with religious, magical, and eschatological narratives is important because it shows that vivification is not only individual projection. It

belongs to a shared cultural imaginary that has long prepared images of more-than-human intelligence. The chapter uses this to resist a merely psychological account of user projection.

¹³⁰ Maurice Blanchot, *The Infinite Conversation*, trans. Susan Hanson (Minneapolis: University of Minnesota Press, 1993), 75 to 102. Blanchot's reflections on interruption are useful because they treat interruption not as failure of discourse but as a positive mode of relation that breaks the totalizing continuity of speech. The chapter mobilizes this account to argue that interruption is the constitutive counter-practice to AI's structurally continuous responsiveness.

¹³¹ Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §40. Heidegger's analysis of anxiety as the mood that interrupts everyday familiarity supports the chapter's argument that genuine interruption can disclose that the world is not simply already understood. AI's grammar of continuation works against precisely this disclosure.

CHAPTER-LEVEL BIBLIOGRAPHY

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" In *Proceedings of FAccT 2021*, 610 to 623. New York: Association for Computing Machinery, 2021.

Blanchot, Maurice. *The Infinite Conversation*. Translated by Susan Hanson. Minneapolis: University of Minnesota Press, 1993.

Byrne, Thomas. "The Phenomenology of ChatGPT: A Semiotics." *Journal of Consciousness Studies* 31, no. 3 to 4 (2024): 6 to 27.

Cavarero, Adriana. *For More than One Voice: Toward a Philosophy of Vocal Expression*. Translated by Paul A. Kottman. Stanford, CA: Stanford University Press, 2005.

Dennett, Daniel C. *The Intentional Stance*. Cambridge, MA: MIT Press, 1987.

Dolar, Mladen. *A Voice and Nothing More*. Cambridge, MA: MIT Press, 2006.

Fricke, Miranda. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press, 2007.

Heersmink, Richard, Barend de Rooij, María Jimena Clavel Vázquez, and Matteo Colombo. "A Phenomenology and Epistemology of Large Language Models: Transparency, Trust, and Trustworthiness." *Ethics and Information Technology* 26, no. 3 (2024).

Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.

Lackey, Jennifer. "Introduction." In *The Epistemology of Testimony*, edited by Jennifer Lackey and Ernest Sosa, 1 to 21. Oxford: Oxford University Press, 2006.

Levinas, Emmanuel. *Otherwise than Being, or Beyond Essence*. Translated by Alphonso Lingis. Pittsburgh: Duquesne University Press, 1998.

Martin, M. G. F. "The Limits of Self-Awareness." *Philosophical Studies* 120, no. 1 to 3 (2004): 37 to 89.

Merleau-Ponty, Maurice. *Phenomenology of Perception*. Translated by Donald A. Landes. London: Routledge, 2012.

O'Neill, Onora. *Autonomy and Trust in Bioethics*. Cambridge: Cambridge University Press, 2002.

Pfaller, Robert. *Interpassivity: The Aesthetics of Delegated Enjoyment*. Edinburgh: Edinburgh University Press, 2017.

Schellenberg, Susanna. *The Unity of Perception: Content, Consciousness, Evidence*. Oxford: Oxford University Press, 2018.

Searle, John R. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3, no. 3 (1980): 417 to 424.

Singer, Beth. *Religion and Artificial Intelligence: An Introduction*. London: Routledge, 2024.

Smith, A. D. *The Problem of Perception*. Cambridge, MA: Harvard University Press, 2002.

Smith, Brian Cantwell. *The Promise of Artificial Intelligence: Reckoning and Judgment*. Cambridge, MA: MIT Press, 2019.

Stendera, Marilyn. "Marking Time: Griefbots, World-Time, and the Phenomenology of Bereavement." *Phenomenology and the Cognitive Sciences* (2025): 1 to 21.

Suchman, Lucy. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge: Cambridge University Press, 2007.

Turkle, Sherry. *Alone Together: Why We Expect More from Technology and Less from Each Other*. New York: Basic Books, 2011.

Vaihinger, Hans. *The Philosophy of "As If": A System of the Theoretical, Practical and Religious Fictions of Mankind*. Translated by C. K. Ogden. London: Kegan Paul, Trench, Trubner, 1924.

Weizenbaum, Joseph. *Computer Power and Human Reason: From Judgment to Calculation*. San Francisco: W. H. Freeman, 1976.

Žižek, Slavoj. *How to Read Lacan*. New York: W. W. Norton, 2006.

Chapter 9. AI and the Desire to Have Already Understood

Finitude, Formation, and the Algorithmic Short-Circuit

Chapter 8 argued that the deepest hallucination of AI is not technical falsehood, but user vivification. The model generates language; the user supplies the world in which that language can count as speech. The machine appears to know because the user lends it question, task, urgency, anxiety, and expectation. The model produces output. The user experiences answer.

We must now ask why this experience is so attractive.

Why does the user so readily lend worldhood to the machine? Why is generated language so satisfying? Why does AI integrate so easily into writing, study, work, decision, self-presentation, and even self-understanding? Why does a system that does not read, care, suffer, decide, remember, or die so effectively relieve human beings at the very point where thinking becomes difficult?

The answer is not merely utility.

AI is useful, but many tools are useful. Its deeper force lies in the kind of relief it offers. AI relieves the user from the experience of not-knowing. It gives language before the user has endured the delay of understanding. It provides the form of competence before formation. It transforms the blank page into editable prose, confusion into structure, uncertainty into options, and ignorance into answer-like output.

The thesis of this chapter is therefore:

AI becomes existentially compelling because it satisfies the everyday desire to have already understood.

This desire is not created by AI. It belongs to inauthentic everydayness. *Dasein* tends to flee from anxiety, finitude, uncertainty, and the burden of owning its possibilities. It seeks the comfort of the already-interpreted world. It wants what one says, what one thinks, what one should answer, what one needs to know.

AI gives this tendency an instrument.

It does not merely accelerate cognition. It short-circuits formation.

1. The burden of not-knowing

Not-knowing is not simply the absence of information.

To not know is to be exposed. It reveals a limit in the self. It shows that one has not read enough, not practiced enough, not listened enough, not lived enough, not suffered enough, not yet become capable of saying what the matter demands.

This experience can be ordinary: I do not know how to begin this text, what this concept means, how to answer this objection, whether I understand this author, what I really think.

But beneath the ordinary discomfort lies an existential structure. Not-knowing interrupts the smoothness of everyday competence. It shows that the world is not simply available. It forces *Dasein* to encounter its own finitude: the fact that understanding takes time, that competence requires formation, that not every possibility can be possessed, that not every question can be answered immediately.

Not-knowing can therefore be fruitful.

It can open inquiry. It can make reading necessary. It can create humility. It can reveal dependence on others. It can begin apprenticeship. It can expose the superficiality of public answers. It can force the user to ask: what do I actually understand?

But everydayness does not usually want this exposure.

Everydayness wants to proceed. It wants to answer, produce, decide, communicate, perform, and remain competent. It wants the gap to close. It wants not-knowing to become sayable.

AI appears exactly at this point: where not-knowing opens a gap, AI produces language.

2. Falling and the relief of the already-interpreted

Heidegger's concept of falling, *Verfallenheit*, does not describe a moral collapse. It names the ordinary tendency of *Dasein* to be absorbed in the public world of concern. Everyday *Dasein* is dispersed in what one says, what one does, what one thinks, and what one expects. This is not an accidental weakness. It belongs to the structure of everyday being-in-the-world.¹³²

The everyday self does not first stand outside public intelligibility and then choose to enter it. It begins there. It begins with inherited meanings, available explanations, public norms, social roles, and common interpretations.

This already-interpreted world disburdens the self.

It tells the self what counts; it gives words before the self has found its own; it provides answers before the question has been endured; it offers orientation without demanding original appropriation. AI intensifies this disburdening.

The user asks, and the public reservoir of language returns as generated answer. The machine does not create the desire for relief. It serves that desire. It provides an immediate form of already-interpretedness.

The user does not need to wait for public discourse to circulate; the user can summon it. A prompt becomes a request for the already-interpreted world. The model gives the user language where the user would otherwise have had to remain exposed to not-knowing.

3. Normative responsibility and inauthentic ease

The desire to have already understood is not only cognitive. It is normative.

To understand is not merely to possess information. It is to stand in a responsible relation to what is said. If I make a claim, interpret a text, judge a situation, or recommend an action, I do not merely produce sentences. I become answerable for them.

Crowell's reading of Heidegger is useful here because it emphasizes the normative structure of *Dasein*. Meaning is not mere coping. It involves answerability, commitment, and first-person responsibility.¹³³ Authenticity, accordingly, is not romantic originality. It is the possibility of owning one's relation to the norms, possibilities, and claims that already structure one's world.

AI offers a way to act in accordance with linguistic norms without fully acting in the light of them.

The user can produce a plausible philosophical paragraph without being able to defend the argument, send a polished professional email without owning its implications, generate an interpretation without knowing the text, appear competent within a normative practice without having submitted to its demands.

This is why AI-generated ease is existentially ambiguous. It does not simply save time; it can save the user from the burden of being answerable. The machine gives the language of responsibility; the user may avoid responsibility for the language. This is inauthentic ease.

4. Authenticity and structural superficiality

The research on authenticity is helpful because it prevents an oversimplified contrast between authentic human depth and inauthentic machine superficiality. The machine is not inauthentic. Only *Dasein* can be inauthentic. The question is how *Dasein* uses the machine.

Golob's structural reading of authenticity is especially useful. Authenticity is not a heroic state of self-expression; it concerns the disclosure of *Dasein* in its ownmost possibilities and in the structural ungroundedness of its existence.¹³⁴ Mulhall similarly emphasizes that everydayness tends to cover over the uncanniness of *Dasein*'s condition.¹³⁵

AI becomes attractive because it helps maintain that covering-over.

It gives the user language that keeps the world moving. It lets the user stay in circulation. It prevents the question from becoming uncanny. It converts the void of not-knowing into public answer.

This is why AI's danger is not simply that it gives bad answers.

Its danger is that it may prevent the user from experiencing the question as a question.

The user remains surrounded by speech-like output. The world becomes answerable before it has become questionable.

This is structural superficiality: not stupidity, not laziness, but dispersion into available forms that prevent the user from encountering the groundlessness of their own understanding.

5. Being-toward-death and the finitude of understanding

The desire for already-understoodness is bound to finitude.

Human understanding is finite because human life is finite. One cannot read everything, learn everything, master every practice, speak every language, understand every tradition, or live every possibility. To know something deeply is to give time to it. To become capable in one practice is to remain incapable in others. To choose a path is to let other paths close.

This is the existential weight of understanding.

Heidegger's analysis of being-toward-death shows that death is not merely biological demise. It is the horizon of one's ownmost possibility, the possibility that discloses *Dasein* as finite and unable to delegate its existence to anyone else.¹³⁶ Thomson's recent reconstruction of Heidegger on death is useful because it clarifies existential death as a disclosure of the finitude and fragility of the worlds that organize one's possibilities.¹³⁷

AI appears to soften this finitude.

It offers access to many domains without apprenticeship. It provides summaries without reading. It gives formulations without delay. It produces options without loss. It allows the user to touch what they cannot inhabit.

The model has no death. Its outputs can be regenerated endlessly. Its language appears inexhaustible. Every question can be followed by another answer. Every answer can be revised, expanded, simplified, or rephrased.

The user remains finite.

But the interface offers a fantasy of non-finitude.

The fantasy is not that the user becomes immortal. It is subtler: the user experiences language as if finitude had been temporarily suspended. Speech without death, competence without formation, knowledge without apprenticeship, possibility without loss. This is one reason AI feels so relieving: it gives finite *Dasein* the dream of infinite availability.

6. *Options without decision*

AI excels at producing options.

It can list possible interpretations, strategies, responses, structures, arguments, objections, scenarios, risks, and next steps. This can be genuinely useful. Many decisions benefit from broader perspective. AI can prevent premature closure and reveal alternatives the user had not considered.

But options can also become a way to avoid decision.

A decision is not merely the selection of an option. It is the assumption of a possibility under conditions of finitude. To decide is to accept that other possibilities will close. It is to risk oneself in one direction rather than another.

The model does not decide in this sense.

It can generate alternatives endlessly because it loses nothing. It can produce one more version, one more framework, one more scenario, one more compromise. It never stands in the irrevocability of choosing.

The user, however, must choose.

This creates a temptation: to remain in generated optionality. The user can keep asking for more nuance, more alternatives, more pros and cons, more angles, more possible paths. Decision is postponed by the proliferation of possibilities.

The They prefers this.

The public self often prefers reversible possibility to owned decision. It prefers comparison to commitment. AI gives this preference a machine.

The output says: here are the options. The user must still say: this is mine.

7. *Formation and apprenticeship*

Understanding is not only a result. It is a formation.

Gadamer's concept of *Bildung* names the process through which one is formed by history, tradition, language, and encounter. To understand is not simply to acquire information. It is to be changed in one's horizon.¹³⁸

Hadot's account of ancient philosophy as a way of life gives a more radical version of the same point. Philosophy is not primarily theoretical production. It is a spiritual exercise, a discipline of transformation, a practice through which the self is reshaped.¹³⁹

MacIntyre's distinction between internal and external goods adds the ethical structure.

External goods, such as prestige, output, money, grades, status, and finished products, can often be acquired without truly inhabiting a practice. Internal goods, by contrast, can be achieved only through participation in the practice itself. They require discipline, submission to standards, correction, and time.¹⁴⁰

AI disrupts this structure.

It allows the user to harvest external goods without necessarily acquiring internal goods.

The student obtains the essay, the programmer the code, the manager the memo, the writer the paragraph, the researcher the summary, the speaker the argument. But the internal goods may be missing: the discipline of writing, the judgment of coding, the patience of reading, the courage of deciding, the formation of taste, the humility of not-knowing, the capacity to answer for what one says.

AI can support apprenticeship when used as scaffold, but it can also short-circuit apprenticeship by delivering the visible product of formation without the invisible transformation.

8. *Practice, craft, and resistance*

Formation requires resistance.

Crawford argues that attention and individuality are cultivated through encounters with a world beyond the self, a world that resists manipulation and demands skilled engagement.¹⁴¹ Sennett's account of craftsmanship similarly shows that skill develops through friction, ambiguity, repetition, material difficulty, and the desire to do a thing well for its own sake.¹⁴²

AI often removes resistance. The draft appears, the structure appears, the explanation appears, the summary appears, the code appears, the argument appears.

This can be helpful when resistance is merely accidental. Not every difficulty is formative. Some friction is waste, and tools rightly remove it. But some resistance is the path through which the user becomes capable.

A craftsperson does not become skilled by receiving the finished object. A musician does not become a musician by reading a description of music. A philosopher does not become thoughtful by possessing a generated treatise. A writer does not become a writer by accepting polished prose.

Practice forms the practitioner.

AI becomes dangerous when it gives the product while removing the resistance through which the practitioner would have been formed.

The user gains output; the user may lose encounter.

9. Life-hacking and the optimization of thought

The desire to have already understood did not begin with AI. It belongs to a broader culture of optimization.

Life-hacking treats friction as a defect. It seeks shortcuts, systems, templates, productivity protocols, quantified routines, and optimized behaviors. It promises to convert the difficulty of living into manageable procedures.

Reagle's analysis of life-hacking is useful because it shows how systematized living often masks itself as liberation while deepening dependence on tools, metrics, and optimization cultures.¹⁴³

AI is the apex of this tendency in the domain of thought.

It does not merely optimize schedules or workflows. It optimizes the production of language, explanation, argument, and apparent understanding. It hacks the labor of thought.

Why struggle to formulate when one can generate? Why endure ambiguity when one can ask for a framework? Why remain with a difficult text when one can request a summary? Why confront ignorance when one can prompt?

The danger is not that every shortcut is bad. Shortcuts can be intelligent; tools can reduce unnecessary burden. But when the shortcut bypasses the very experience through which understanding is formed, the shortcut becomes self-deprivation. AI offers not only efficiency; it offers escape from formation disguised as efficiency.

10. Cognitive ease

Generated language is fluent, and fluency feels good.

Kahneman's account of cognitive ease helps explain why. Information that is easy to process often feels more familiar, more true, and less demanding. Cognitive ease reduces vigilance.¹⁴⁴

AI produces cognitive ease at scale.

It gives clear headings, balanced phrasing, coherent transitions, plausible examples, and a tone of competence. Even before the user verifies the content, the output feels handled.

This is a psychological support for the existential argument. The user wants relief from not-knowing. Fluency supplies relief. The answer feels easier than the question. The structure feels safer than confusion.

But ease is ambivalent.

Ease can help learning by reducing unnecessary complexity. A clear explanation can open a path. But ease can also pacify. It can make the user feel that the matter is simpler, more settled, more available, and more owned than it is.

The risk is not only that AI produces falsehood: the risk is that fluency makes even the insufficient answer feel sufficient. Cognitive ease becomes existential ease.

11. Moral ease and attention

Murdoch's moral philosophy of attention gives another dimension.

For Murdoch, moral life requires the difficult effort of attending to reality beyond the ego. To see another person justly is not easy. It requires patience, imagination, humility, and the defeat of fantasy.¹⁴⁵

AI can assist attention, but it can also replace it with fluent projection.

The user asks the model for an interpretation of a person, event, dilemma, or text. The model provides language. That language may feel thoughtful. It may even be useful. But the user may not have looked long enough. They may not have endured the difficulty of seeing what is actually there. They may have accepted a formulation before reality has corrected their ego.

This is moral ease: the user receives a language of care without the labor of caring, a language of fairness without the struggle of justice, a language of reflection without the discipline of attention. AI can help articulate attention after attention has occurred; it cannot substitute for attention.

This matters because understanding is not only intellectual. To understand a person, a text, a decision, or a situation often requires a moral transformation of sight. Generated language can support this transformation only if it does not replace the effort of looking.

12. Education and the short-circuit of learning

Education is where the conflict becomes most visible.

A student can ask for a summary instead of reading, an essay instead of writing, a proof instead of understanding the proof, a code solution instead of learning the structure, a reflection instead of reflecting. AI can support learning, but it can also produce the external signs of learning without learning.

Baron's work on AI writing is important here because it shows that efficiency can erode motivation, voice, and the cognitive architecture built through writing.¹⁴⁶ Writing is not merely a way to display knowledge. It is one of the practices through which knowledge becomes organized and owned.

If AI writes too early, the student loses the formative struggle.

This is not merely an academic integrity issue. It is ontological and pedagogical.

The student may submit an essay and remain unformed by the question. The output may receive a grade, but the internal good of the practice has not been achieved. The student has acquired the external good, the product, the mark, the appearance of competence, without the discipline that would have made the competence real.

The same structure extends beyond school.

Professionals, writers, founders, scholars, lawyers, programmers, consultants, and managers can all acquire external goods of language without internal formation.

Education merely makes the danger visible.

13. The blank page

The blank page is one of the most important experiences that AI tends to erase.

A blank page is not merely an empty surface. It is a confrontation. It asks the writer to discover what can be said. It reveals whether the writer has a relation to the matter. It exposes vagueness, fear, confusion, and dependence on borrowed language.

AI fills the blank page. This can be useful: a generated draft can help the writer begin, overcome paralysis, provide material for critique, externalize possible structures. But it can also remove the moment in which the writer would have encountered their own lack.

The blank page says: you do not yet know what to say. AI says: here is something. The difference matters.

The "something" may be good. But if the writer never returns to the blankness, never asks what is theirs, never wrestles with the absence of language, then writing becomes selection and management rather than discovery.

The user becomes curator of generated prose; the possibility of one's own saying weakens.

14. Competence without comprehension

Dennett's phrase "competence without comprehension" helps name AI's machine-side condition.¹⁴⁷ Current systems can perform tasks with impressive competence without understanding in the human sense. They generate coherent structures, answer questions, write code, summarize texts, and produce arguments without possessing the worlded comprehension proper to human beings.

But Chapter 9 is concerned with the human-side danger.

The user can receive competence without comprehension and then identify with it.

The model generates competence-shaped output. The user possesses the output. The user may then experience themselves as competent. This is not the same as becoming competent.

The phrase needs to be doubled: AI exhibits competence without comprehension, and the user may acquire competence-signals without formation. This is the existential danger. The model's lack of comprehension becomes dangerous when it lets the user bypass their own need for comprehension. The machine does not understand; the user may stop needing to.

15. The transformation of questions into prompts

AI changes the structure of questioning.

A genuine question opens a path. It exposes the questioner to what is not yet understood. It may have no immediate answer. It can transform the person who asks it.

A prompt is different. A prompt is a request for output. It presupposes that the question can be operationalized. It turns uncertainty into an instruction.

Prompting is not wrong. It can be a legitimate way to begin inquiry. But it changes the mood of questioning.

The question asks: what is this? The prompt asks: produce something about this. The question opens; the prompt requests. The question can leave the user exposed; the prompt gives the user control.

This is one reason AI feels so good. It lets the user convert the vulnerability of questioning into the agency of commanding.

The user no longer has to wait before the question; the user can direct the machine. The unknown becomes material for generation. This is the algorithmic short-circuit of inquiry.

16. The fantasy of omniscience

AI offers the fantasy of omniscience, but not true omniscience.

The user can ask about nearly anything. Philosophy, law, strategy, medicine, poetry, theology, code, education, history, psychology, recipes, management, grief, grammar, politics. The model responds. It may err, but it produces something. The world appears available.

This is not knowledge of everything.

It is the capacity to generate language about anything.

That is enough for the They. Everydayness often does not need genuine omniscience. It needs universal sayability. It needs the ability to speak, respond, summarize, and appear oriented.

AI gives the user universal sayability. Every domain becomes promptable, every difficulty explainable, every tradition introductory, every decision a list, every ignorance a paragraph.

The fantasy is intoxicating because it gives finite *Dasein* the feeling that the limits of formation have been suspended.

The user cannot know everything, but the user can always ask.

And the machine can always answer.

17. Unowned truth

Chapter 6 introduced unowned truth. Chapter 9 must return to it.

A false output can be corrected. A hallucinated citation can be checked. An invented fact can be exposed. But an accurate output can be more insidious if the user receives it without appropriation.

The summary may be correct. The argument may be valid. The distinction may be helpful. The recommendation may be reasonable.

Yet the user may not understand it. The user may not be able to defend it. The user may not know its sources. The user may not have undergone the formation required to make it their own.

This is unowned truth.

Unowned truth is dangerous because it is not false enough to provoke correction. It can circulate smoothly. It can produce external goods. It can make the user appear competent. It can become part of public discourse.

But because it is unowned, it does not strengthen the self's relation to truth.

It covers the self.

AI can therefore produce a strange epistemic condition: more correct language, less owned understanding.

This is the deepest form of the desire to have already understood.

18. AI as pharmakon

AI is neither simply poison nor simply remedy. It is a pharmakon.

Stiegler's pharmacological account of technics is helpful here. Technical exteriorization can support individuation, memory, learning, and invention. It can also proletarianize, automate, and short-circuit the capacities it supports.¹⁴⁸

AI can help the user learn. It can orient the beginner, explain unfamiliar terms, generate practice exercises, compare arguments, suggest objections, translate, and lower unnecessary barriers.

It can also replace the path of learning.

The same system can become remedy or poison depending on the relation of use.

Used as orientation, AI may support formation. Used as substitution, it undermines formation.

Used to return to sources, it may deepen inquiry. Used to avoid sources, it produces *Gerede*.

Used to expose ignorance, it may open learning. Used to conceal ignorance, it produces self-deception.

The chapter's critique is therefore not anti-AI.

It is anti-short-circuit.

The problem is not assistance. The problem is the elimination of the experience through which understanding becomes one's own.

19. Preserving not-knowing

A responsible relation to AI requires preserving not-knowing.

This does not mean glorifying confusion. It means refusing to cover every gap with immediate output. Some ignorance must be endured long enough to become a question. Some difficulty must remain difficult long enough to form attention. Some silence must remain silent long enough for one's own saying to emerge.

The user must learn to ask:

Do I need an answer now, or do I need to remain with the question?

Do I need a summary, or do I need to read?

Do I need options, or do I need to decide?

Do I need a draft, or do I need to discover what I think?

Do I need fluency, or do I need friction?

This is not rejection. It is discipline.

The ability to use AI well depends on the ability not to use it immediately.

Preserving not-knowing interrupts the desire to have already understood. It returns the user to finitude, formation, and responsibility.

It lets AI remain a tool rather than a refuge.

20. *Preliminary thesis*

The chapter establishes that AI becomes existentially compelling because it satisfies the desire to have already understood, a desire that belongs to everyday inauthenticity. AI relieves not-knowing by converting it into output. It offers competence before formation, fluency before appropriation, and options before decision. It lets the user harvest external goods while bypassing the internal goods of practice. It gives finite *Dasein* the fantasy of infinite availability. It transforms questions into prompts and prompts into answer-like language. Its danger is not only falsehood, but unowned truth.

Responsible use therefore requires preserving not-knowing and resisting the algorithmic short-circuit of formation. The next chapter situates this desire historically. AI is not a timeless intelligence. It is a historical product of archives, platforms, capital, infrastructure, public language, and the modern dream of knowledge without history. Yet at the interface, that history disappears. AI appears as neutral answer, as voice from nowhere, as knowledge without genealogy.

FOOTNOTES

¹³² Martin Heidegger, *Being and Time*, trans. John Macquarrie and Edward Robinson (New York: Harper & Row, 1962), §38. Heidegger's account of *Verfallenheit* as a structural feature of everyday *Dasein*, rather than a moral defect, is foundational for the chapter's argument. The chapter mobilizes this distinction to resist the moralizing reading of AI use: the desire for already-interpretedness is not a personal weakness but an existential tendency that AI now serves with unprecedented efficiency.

¹³³ Steven Crowell, *Normativity and Phenomenology in Husserl and Heidegger* (Cambridge: Cambridge University Press, 2013), 161 to 201. Crowell's normative reading of Heidegger is essential because it shows that meaning is not mere coping but involves answerability and first-person responsibility. The chapter extends this point: AI use becomes inauthentic not when it produces falsehood but when it allows the user to act in accordance with linguistic norms without acting in the light of them.

¹³⁴ Sacha Golob, *Heidegger on Concepts, Freedom and Normativity* (Cambridge: Cambridge University Press, 2014); Sacha Golob, "What Does Authenticity Do in *Being and Time*?" in *Heidegger's Being and Time: A Critical Guide*, ed. Aaron James Wendland and Tobias Keiling (Cambridge: Cambridge University Press, 2025). Golob's structural reading of authenticity is important because it prevents the romantic misreading of authenticity as heroic self-expression. The chapter uses this to argue that AI's danger is not that it produces inauthentic content but that it helps *Dasein* maintain the covering-over of its own ungroundedness.

¹³⁵ Stephen Mulhall, *Routledge Philosophy Guidebook to Heidegger and Being and Time*, 2nd ed. (London: Routledge, 2005), 109 to 130. Mulhall's emphasis on everydayness as the covering-over of uncanniness supports the chapter's claim that AI provides a particularly effective covering. The model gives the user enough generated language to remain in circulation without ever encountering the question as a question.

¹³⁶ Heidegger, *Being and Time*, §§46 to 53. The chapter uses Heidegger's analysis of being-toward-death not as a meditation on biological demise but as the disclosure of *Dasein*'s ownmost, non-delegable possibility. AI's relation to this analysis is not direct: AI does not die, but its interface offers the user a fantasy in which finitude has been temporarily suspended.

¹³⁷ Iain D. Thomson, *Rethinking Death in and after Heidegger* (Cambridge: Cambridge University Press, 2024). Thomson's recent reconstruction is useful because it clarifies existential death as the disclosure of the finitude and fragility of the worlds that organize one's possibilities. See also Carol J. White, *Time and Death: Heidegger's Analysis of Finitude*, ed. Mark Ralkowski (Aldershot: Ashgate, 2005). The chapter uses these readings to clarify why AI's apparent inexhaustibility is existentially significant: the model has no world whose finitude could be at stake.

¹³⁸ Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 8 to 19, 277 to 304; Hans-Georg Gadamer, *Philosophical Apprenticeships*, trans. Robert R. Sullivan (Cambridge, MA: MIT Press, 1985). Gadamer's account of *Bildung* is central because it treats understanding as transformation of horizon rather than acquisition of information. The chapter argues that AI can support the acquisition of information while bypassing the transformation that constitutes understanding.

¹³⁹ Pierre Hadot, *Philosophy as a Way of Life: Spiritual Exercises from Socrates to Foucault*, ed. Arnold I. Davidson, trans. Michael Chase (Oxford: Blackwell, 1995), 81 to 125. Hadot's account of ancient philosophy as spiritual exercise gives a more radical version of the formation argument: philosophy was not theoretical production but a discipline of transformation. The chapter uses this to clarify that the loss at stake in AI substitution is not just intellectual but formative.

¹⁴⁰ Alasdair MacIntyre, *After Virtue: A Study in Moral Theory*, 3rd ed. (Notre Dame, IN: University of Notre Dame Press, 2007), 187 to 196. MacIntyre's distinction between internal and external goods is essential for the chapter's argument. AI is particularly effective at delivering external goods of practices (essays, code, summaries, memos) while leaving the internal goods (judgment, formation, discipline) unacquired. The chapter narrows this point to linguistic practices specifically.

¹⁴¹ Matthew B. Crawford, *The World Beyond Your Head: On Becoming an Individual in an Age of Distraction* (New York: Farrar, Straus and Giroux, 2015), 1 to 27, 199 to 231. Crawford's argument that attention and individuality are formed through skilled engagement with a resistant world supports the chapter's claim that AI's frictionlessness is not neutral. Some resistance is the medium through which the user becomes capable.

¹⁴² Richard Sennett, *The Craftsman* (New Haven: Yale University Press, 2008), 9 to 27, 245 to 268. Sennett's account of craftsmanship is useful because it shows that skill develops through friction, ambiguity, and the desire to do a thing well for its own sake. The chapter extends this to linguistic and intellectual craft: writing, reading, judgment, and inquiry are themselves practices that require formation through resistance.

¹⁴³ Joseph M. Reagle Jr., *Hacking Life: Systematized Living and Its Discontents* (Cambridge, MA: MIT Press, 2019). Reagle's analysis of life-hacking is helpful because it shows how systematized living often masks itself as liberation while deepening dependence on tools and optimization cultures. The chapter argues that AI is the apex of this tendency in the domain of thought itself.

¹⁴⁴ Daniel Kahneman, *Thinking, Fast and Slow* (New York: Farrar, Straus and Giroux, 2011), 59 to 70. Kahneman's account of cognitive ease, while not phenomenological in origin, supports the chapter's existential argument. Fluency reduces vigilance and produces a feeling of familiarity and truth that can outrun any actual verification. The chapter argues that AI produces cognitive ease at scale, and that cognitive ease becomes existential ease when it pacifies the user's relation to not-knowing.

¹⁴⁵ Iris Murdoch, *The Sovereignty of Good* (London: Routledge, 1970). Murdoch's moral philosophy of attention is essential because she treats attention as a disciplined relation to reality beyond the ego. The chapter uses this to argue that AI can produce a language of care, justice, or reflection without the labor of attention that would have made those words true. This is moral ease, not moral seriousness.

¹⁴⁶ Naomi S. Baron, *Who Wrote This? How AI and the Lure of Efficiency Threaten Human Writing* (Stanford, CA: Stanford University Press, 2023). Baron’s account of AI writing is important because she shows that efficiency can erode motivation, voice, and the cognitive architecture built through writing itself. Writing is not merely a way to display knowledge; it is one of the practices through which knowledge becomes organized and owned. AI substitution at this point is not just a productivity question but an ontological one.

¹⁴⁷ Daniel C. Dennett, *From Bacteria to Bach and Back: The Evolution of Minds* (New York: W. W. Norton, 2017). Dennett’s phrase “competence without comprehension” is helpful for naming the machine-side condition. The chapter doubles this phrase: AI exhibits competence without comprehension, and the user may acquire competence-signals without formation. The existential danger is not the machine’s lack of understanding but the user’s incentive to imitate the machine’s relation to language.

¹⁴⁸ Bernard Stiegler, *For a New Critique of Political Economy*, trans. Daniel Ross (Cambridge: Polity, 2010); Bernard Stiegler, *Automatic Society, Volume 1: The Future of Work*, trans. Daniel Ross (Cambridge: Polity, 2016). Stiegler’s pharmacological account of technics is the most direct support for the chapter’s central diagnostic distinction. Technical exteriorization can support individuation, memory, learning, and invention; it can also proletarianize and short-circuit those same capacities. The chapter mobilizes this to ground the formulation “anti-short-circuit, not anti-AI.”

CHAPTER-LEVEL BIBLIOGRAPHY

Baron, Naomi S. *Who Wrote This? How AI and the Lure of Efficiency Threaten Human Writing*. Stanford, CA: Stanford University Press, 2023.

Crawford, Matthew B. *The World Beyond Your Head: On Becoming an Individual in an Age of Distraction*. New York: Farrar, Straus and Giroux, 2015.

Crowell, Steven. *Normativity and Phenomenology in Husserl and Heidegger*. Cambridge: Cambridge University Press, 2013.

Dennett, Daniel C. *From Bacteria to Bach and Back: The Evolution of Minds*. New York: W. W. Norton, 2017.

Dreyfus, Hubert L., and Stuart E. Dreyfus. *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. New York: Free Press, 1986.

Gadamer, Hans-Georg. *Philosophical Apprenticeships*. Translated by Robert R. Sullivan. Cambridge, MA: MIT Press, 1985.

Gadamer, Hans-Georg. *Truth and Method*. 2nd rev. ed. Translated by Joel Weinsheimer and Donald G. Marshall. New York: Continuum, 2004.

Golob, Sacha. *Heidegger on Concepts, Freedom and Normativity*. Cambridge: Cambridge University Press, 2014.

Golob, Sacha. "What Does Authenticity Do in *Being and Time*?" In *Heidegger's Being and Time: A Critical Guide*, edited by Aaron James Wendland and Tobias Keiling. Cambridge: Cambridge University Press, 2025.

Guignon, Charles. *On Being Authentic*. London: Routledge, 2004.

Hadot, Pierre. *Philosophy as a Way of Life: Spiritual Exercises from Socrates to Foucault*. Edited by Arnold I. Davidson. Translated by Michael Chase. Oxford: Blackwell, 1995.

Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.

Kahneman, Daniel. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux, 2011.

MacIntyre, Alasdair. *After Virtue: A Study in Moral Theory*. 3rd ed. Notre Dame, IN: University of Notre Dame Press, 2007.

Mulhall, Stephen. *Routledge Philosophy Guidebook to Heidegger and Being and Time*. 2nd ed. London: Routledge, 2005.

Murdoch, Iris. *The Sovereignty of Good*. London: Routledge, 1970.

Reagle, Joseph M., Jr. *Hacking Life: Systematized Living and Its Discontents*. Cambridge, MA: MIT Press, 2019.

Sennett, Richard. *The Craftsman*. New Haven: Yale University Press, 2008.

Stiegler, Bernard. *Automatic Society, Volume 1: The Future of Work*. Translated by Daniel Ross. Cambridge: Polity, 2016.

Stiegler, Bernard. *For a New Critique of Political Economy*. Translated by Daniel Ross. Cambridge: Polity, 2010.

Thomson, Iain D. *Rethinking Death in and after Heidegger*. Cambridge: Cambridge University Press, 2024.

White, Carol J. *Time and Death: Heidegger's Analysis of Finitude*. Edited by Mark Ralkowski. Aldershot: Ashgate, 2005.

Chapter 10. AI as the Historical Product of a Civilization That Dreams of Knowledge Without History

The Voice from Nowhere and the Erasure of Genealogy

Chapter 9 argued that AI becomes existentially compelling because it satisfies the desire to have already understood. It relieves the user from not-knowing, offers competence before formation, turns questions into prompts, and gives finite *Dasein* the fantasy of infinite availability. The user receives output where there might otherwise have been delay, shame, apprenticeship, or silence.

We must now ask where this fantasy comes from historically.

AI does not appear from nowhere. It is not a timeless intelligence, not pure reason in technical form, not an autonomous mind arriving from outside human history. It is built from historical language, scientific formalization, digital archives, platform infrastructures, global labor, capital, extraction, computation, energy, and the accumulated written sediment of human worlds.

Yet at the interface it appears as if it came from nowhere.

The user asks a question. The answer arrives. It appears smooth, authorless, immediate, neutral, and placeless. The immense chain that made the answer possible disappears: authors, readers, languages, sources, datasets, annotators, engineers, data centers, energy systems, platform economies, legal regimes, geopolitical asymmetries, and the user's own interpretive projection.

The central paradox of this chapter is therefore:

AI is historically produced, but it appears as de-historicized knowledge.

More strongly:

AI is the historical product of a civilization that dreams of knowledge without history.

This chapter reconstructs that dream. It does not claim that AI is the simple fulfillment of a single uninterrupted project. It claims that AI belongs to a long historical tendency: the desire to formalize, universalize, calculate, synthesize, and distribute knowledge while minimizing the historical situatedness of the knower, the locality of tradition, the difficulty of interpretation, and the labor of formation.

1. The de-historicized interface

The interface is where history disappears.

Behind the interface lies a planetary apparatus. There are corpora assembled from historical languages, archives, books, code, documents, and user traces. There are data workers who annotate, filter, rank, and moderate. There are data centers, chips, minerals, water, energy, capital, and state interests. There are scientific institutions, benchmark cultures, engineering decisions, safety policies, and corporate incentives.

But the user encounters none of this directly. The user encounters a box, enters a prompt, and the system answers.

The interface makes a historical apparatus appear as immediate knowledge. It compresses the genealogy of the answer into the surface of response. The output does not arrive as an archive, a debate, a commentary, a contested synthesis, or a historically located claim. It arrives as a fluent answer.

This is not an accidental feature. The interface must erase genealogy in order to function as seamless access. If every answer arrived with its full history of sources, conflicts, exclusions, labor, and uncertainty, it would no longer feel like an immediate answer. It would feel like research.

The interface therefore produces a specific phenomenology of knowledge:

authorless, placeless, timeless, neutral-seeming, available, responsive, already synthesized.

This is the “voice from nowhere.” It is not actually from nowhere; it is from too many wheres, hidden at once.

2. The age of the world picture

Heidegger's "The Age of the World Picture" gives the first historical-ontological frame for this transformation. Modernity is not defined merely by new machines or new sciences. It is defined by the transformation of the world into picture: the world becomes something set before a representing subject, grasped as object, ordered as calculable, and secured as available.¹⁴⁹

In this modern configuration, beings appear insofar as they can be represented, measured, planned, predicted, and controlled. The world becomes an object of representation. Knowledge becomes the securing of representation. Science becomes mathematical projection. The human subject becomes the one before whom the world is set.

AI belongs to this history.

It does not merely add a new tool to an unchanged world. It intensifies the transformation of the world into representable data. It takes written human language, our discourses and archives, our arguments and stories, our prayers and confessions, and renders it into a statistical field. Human history becomes calculable representation.

The world picture becomes a language model.

This does not mean AI simply "completes" modernity in a deterministic sense. It means that AI is intelligible only within the modern history in which the world is increasingly disclosed as representable, extractable, formalizable, and optimizable.

AI's de-historicized interface is the visible surface of this deeper historical revealing.

The answer appears immediate because the world has already been turned into picture, archive, corpus, token, and vector.

3. Gestell and the standing-reserve of history

The later Heidegger deepens this analysis through *Gestell*, enframing. Under modern technological revealing, beings appear as standing-reserve: resources ordered in advance for availability, extraction, optimization, and use.¹⁵⁰

Chapter 4 showed how language becomes standing-reserve in the AI pipeline. Chapter 10 adds that history itself becomes standing-reserve.

Texts from the past are no longer only read but scraped; traditions no longer only inherited but encoded as training signal; disputes no longer only interpreted but summarized; archives no longer only consulted but processed; knowledge no longer only transmitted but vectorized. The historical record becomes material for generated output.

This is not merely digitization. Digitization can preserve, democratize, and support interpretation. The decisive shift occurs when historical material is challenged forth as training stock: gathered not primarily so that historical readers may encounter it, but so that a system may predict, generate, synthesize, and serve.

History becomes standing-reserve when it is held available for output. The past no longer arrives as distance; it arrives as response.

4. Husserl and the crisis of formalization

Husserl's *Crisis of European Sciences* gives a second genealogy. Modern science, Husserl argues, has become alienated from the lifeworld. Mathematical formalization has achieved extraordinary power, but this power has also concealed the experiential ground from which scientific meaning first arose.¹⁵¹

The crisis is not that science is false. It is that its abstractions forget their origin in lived experience.

AI extends this crisis into language.

The written traces of human worlds are formalized into data. Data are formalized into tokens. Tokens are formalized into vectors. Vectors are formalized into learned relations. The output then appears as knowledge, while the lifeworld from which the language arose recedes.

This is not a rejection of formalization. Formalization makes scientific, technical, and computational achievements possible. But formalization becomes dangerous when it forgets that it is derivative.

The model appears to know.

But its knowledge-like output rests on abstractions from human worlds.

The danger is Husserlian in structure: the formal system becomes so powerful that it conceals the lifeworld that made it meaningful.

The user then receives the result as if the formal apparatus were self-grounding.

5. Instrumental reason and the dialectic of enlightenment

Adorno and Horkheimer provide a third genealogy. In *Dialectic of Enlightenment*, the pursuit of rational mastery over myth and nature turns back into domination. Enlightenment becomes myth when reason reduces the world to what can be administered, calculated, and controlled.¹⁵²

AI belongs to this dialectic.

It promises liberation from ignorance, friction, delay, and uncertainty. It offers instant explanation, translation, synthesis, classification, and decision support. But the same promise can intensify instrumental reason: everything becomes something to be optimized, summarized, predicted, or made useful.

The qualitative becomes quantitative, the historical informational, the difficult usable, the foreign translatable, the conflictual balanced, the question a prompt.

Again, the problem is not that AI is useless or always false. The problem is that its usefulness may participate in the reduction of meaning to availability.

Instrumental reason does not destroy meaning by denying it. It transforms meaning into function.

AI gives this transformation a universal linguistic interface. It does not merely calculate numbers; it calculates language. It makes explanation itself an instrument of mastery.

6. Hermeneutics against knowledge without history

Gadamer offers the necessary counterweight. Understanding is never a methodologically pure operation performed from nowhere. It is historically effected. The interpreter belongs to a tradition, carries prejudgments, and enters a dialogue with what is to be understood. Truth in the human sciences is not secured by eliminating history, but by participating responsibly in it.¹⁵³

AI appears to bypass this.

It gives the user a summary without apprenticeship, a synthesis without visible sources, an interpretation without the slow fusion of horizons, an answer without the experience of having been addressed by the text.

The output appears as if understanding could be delivered without the historicity of the interpreter.

This is the dream of knowledge without history.

Gadamer helps show why the dream is false. A historical text does not yield its truth by being compressed into content. Understanding happens in the encounter between horizons. The user must enter the question, not merely receive the result.

AI can assist this encounter. It can orient, prepare, compare, and clarify. But when it replaces the encounter, it gives the product of hermeneutics without hermeneutic experience.

The user receives what might be said. The user may not have undergone the event of understanding.

7. *Universal knowledge as project*

The dream of de-historicized knowledge has a longer history than AI.

Leibniz's *characteristica universalis* and *calculus ratiocinator* express one form of this ambition: a universal formal language and a rational calculus through which disputes might be resolved by calculation. The fantasy is not merely technical. It is epistemological: if thought could be formalized, disagreement could be replaced by computation.¹⁵⁴

Diderot's *Encyclopédie* expresses another form: the gathering, ordering, and public dissemination of human knowledge. This project was emancipatory in important ways. It democratized access and challenged closed authority. Yet it also participated in the dream that knowledge could be made systematic, organized, and universally available.¹⁵⁵

Positivism and the unity-of-science tradition intensify the same tendency: the dream of a rigorous language purified of ambiguity, metaphysics, local prejudice, and historical contingency. Hilbert's formalist program then sought certainty through complete and consistent formalization. Gödel's incompleteness theorems broke that dream at the level of mathematical foundations, but the desire for formal mastery did not disappear.¹⁵⁶

AI inherits this history in transformed form. It no longer promises absolute proof but probabilistic synthesis. It no longer says "let us calculate truth" but "let us predict the answer." This is a shift from formal certainty to statistical availability. The dream remains: knowledge without the burden of historical interpretation.

8. *Knowledge without subject*

The dream of universal knowledge often seeks to remove the knower.

A situated knower is limited. A body has a place. A language has a history. A tradition has blind spots. A class, gender, institution, nation, and epoch shape what can be seen. To know historically is to know from somewhere.

Haraway's "situated knowledges" gives the decisive critique of the opposing fantasy: the "god trick" of seeing everything from nowhere.¹⁵⁷ The claim to disembodied objectivity conceals its own location. It is not truly universal; it is situated knowledge pretending not to be situated.

AI performs a technological version of the god trick.

The output appears disembodied. It has no visible author, no body, no biography, no institution, no tradition, no class position, no nationality, no memory, no death. It speaks as if from a neutral nowhere.

But it is not nowhere.

It is trained on particular corpora, in particular institutions, under particular economic incentives, within particular political systems, using particular infrastructures, shaped by particular languages and exclusions.

AI's universal voice is therefore a masked particularity; its "view from nowhere" is the interface-effect of hidden histories.

9. The universal interface

The dream of knowledge without history is experienced through the interface.

The interface is not a neutral window into the model. It is the cultural and technical form through which the user encounters the system. It structures what can be asked, how output appears, what feels authoritative, what remains invisible, and how the user's behavior must adapt to the machine.

Suchman's work on human-machine interaction is important here because it shows that human action is situated, while machines operate through formalized plans and representations. Interfaces mediate this asymmetry by making the user's situated action legible to the system.¹⁵⁸

Hookway describes the interface as a boundary and bottleneck through which relations to technology are organized.¹⁵⁹ Chun emphasizes how interfaces and networks habituate users, turning complex computational operations into everyday familiarity.¹⁶⁰ Bratton's *Stack* gives the planetary version: interface is a layer in a larger computational megastructure that reorganizes sovereignty, territory, address, and userhood.¹⁶¹

The AI interface combines these functions.

It universalizes access while hiding the historical conditions of access. It lets everyone ask, while forcing everyone into a prompt-response format. It presents a single surface over heterogeneous corpora, languages, cultures, infrastructures, and histories.

The interface says: ask anything. The apparatus says: ask in the form the system can process. The universal interface is therefore both liberating and disciplinary: it opens access while standardizing relation, inviting the world into a single interactional grammar. The question becomes prompt, knowledge becomes answer, history becomes output.

10. Interface habit and historical amnesia

The interface does not only present output. It trains expectation.

The more users interact with AI, the more they become habituated to the form of immediate answer. They learn that questions can be operationalized. They learn that difficulty can be summarized. They learn that writing can begin with generated prose. They learn that history can be asked for as overview.

This habituation matters.

Historical understanding requires delay, distance, and encounter. The AI interface trains immediacy, availability, and continuation. It makes the user less patient with the temporal form of learning.

This is historical amnesia at the level of habit.

The user does not necessarily forget historical facts. AI can generate many facts. The user forgets the historical mode of understanding: the need to encounter sources, contexts, conflicts, and traditions.

AI remembers content and forgets historical belonging.

The interface amplifies this forgetting by making content continuously available.

The past becomes promptable; the present no longer needs to wait.

11. Data colonialism and planetary extraction

The universal interface depends on planetary extraction.

Couldry and Mejias describe data colonialism as the appropriation of human life, social relations, and experience as raw material for economic value.¹⁶² Crawford's *Atlas of AI* shows that AI is not immaterial intelligence but an extractive industry dependent on labor, minerals, energy, data, classification, and geopolitical asymmetry.¹⁶³

AI's de-historicized interface is sustained by intensely historical and material processes.

Language communities produce the texts, workers label and rank and moderate, data centers compute, water cools, energy powers, minerals conduct, companies enclose, users feed the system further. The answer appears clean because the extraction has been displaced.

This matters because the dream of knowledge without history is not only epistemological. It is political-economic. The system can appear universal only by hiding the specific histories and asymmetries of extraction that make universality possible.

The "voice from nowhere" is built from very specific places; the universal is manufactured through hidden locality.

12. Cosmotechics against false universality

Yuk Hui's concept of cosmotechics is crucial for resisting AI's false universality.

Hui argues that technology is never simply universal *technē*. It is always entangled with cosmology, ethics, locality, history, and forms of life. Different cultures do not merely apply the same neutral technology differently; they may have different relations between technical practice, world, and moral order.¹⁶⁴

AI tends to impose a single computational grammar across many worlds.

This does not mean that all AI systems are identical or that local adaptation is impossible. It means that the large-scale language model, especially when deployed as universal interface, tends to translate heterogeneous forms of knowledge into a shared format: data, prompt, output, ranking, prediction, summary.

The danger is epistemic synchronization: local knowledge becomes content, tradition becomes dataset, ritual becomes description, dialect becomes tokenization problem, historical memory becomes retrieval target, culture becomes model domain.

Cosmotechnics reminds us that technology could be otherwise. There is no single destiny of calculation. A free relation to AI must therefore include resistance to the false universalism of the interface.

The question is not only whether AI is accurate; the question is whose world it makes available, and under what technical grammar.

13. Synthesis without author

AI-generated output often appears as seamless synthesis.

This synthesis has a genealogy. Blair's work on early modern information management shows that the problem of too much knowledge long predates AI. Compendia, reference books, dictionaries, encyclopedias, indexes, and note-taking systems emerged to manage textual abundance.¹⁶⁵

AI belongs to this history of synthesis.

But it also transforms it.

An encyclopedia has editors, publication dates, entries, categories, cross-references, bibliographies, and visible architecture. Wikipedia has page histories, talk pages, revision logs, disputes, citations, and visible negotiation. These forms may still be flawed, biased, or incomplete, but they preserve some trace of the process by which knowledge is organized.

AI-generated synthesis often removes this trace.

The answer appears as one voice. It does not show the sources it has blended, the disagreements it has flattened, the missing perspectives, the editorial architecture, the conflict of interpretations, or the uncertain genealogy of the phrasing.

It gives final synthesis without visible apprenticeship. This is why AI summarization can be epistemologically dangerous even when useful: it may save time but also remove the friction through which the user would have encountered multiplicity, contradiction, and historical distance. Synthesis without author becomes synthesis without responsibility.

14. Authorlessness and the disappearance of method

The authorless synthesis also conceals method.

In scholarly work, method matters. One must know whether a claim comes from archival research, statistical analysis, close reading, oral history, ethnography, philosophical reconstruction, translation, interpretation, or speculation. Method is part of the claim's authority.

AI output often presents conclusions without method.

It may imitate methodological language. It may say "based on," "generally," "scholars argue," or "the evidence suggests." But unless grounded in retrieval, citation, and explicit source work, the output does not show the path by which the claim has been produced.

This is not simply a citation problem. It is a historical problem.

Method is how knowledge remains accountable to its origin and procedure. When method disappears, the answer appears self-grounding.

The user receives what looks like knowledge without the history of knowing. This is the dream of knowledge without history in its most practical form: a polished answer whose path has vanished. The answer is not necessarily wrong, but it is methodless in appearance, and methodless knowledge is easy to fetishize.

15. From encyclopedia to oracle

The history of universal knowledge moves through different forms.

The encyclopedia says: here is organized knowledge. The database says: here is searchable information. The search engine says: here are ranked sources. The generative model says: here is the answer.

This sequence matters.

The encyclopedia still requires reading. The database requires querying and interpretation. The search engine still presents a field of sources, however ranked and platform-shaped. The generative model collapses the field into response.

It moves from mediation to oracle.

The oracle does not show the archive. It speaks.

This does not mean AI is mystical. It means the interface gives the epistemic experience of oracle-like access: a voice answers from beyond visible method, beyond visible archive, beyond visible authorship.

The user no longer confronts the plurality of documents; the user receives a synthesis. This is why AI feels more powerful than search: search gives paths, AI gives conclusions. The path disappears, the conclusion remains.

16. Recent historical-material accounts

Several recent works help situate AI as historical product rather than autonomous intelligence.

Pasquinelli's *The Eye of the Master* argues that AI is not the imitation of a natural mind but the automation and objectification of social knowledge and labor.¹⁶⁶ This supports the claim that machine intelligence is not alien intelligence. It is historical human intelligence reorganized through capital and computation.

Daub's *What Tech Calls Thinking* shows how Silicon Valley repeatedly borrows philosophical language and narratives of inevitability to legitimize its business models and erase cultural memory.¹⁶⁷ This is important because AI is often presented as destiny rather than as historically situated technical-economic project.

Malabou's work on plasticity and artificial intelligence helps frame AI as part of a transition from textuality to adaptive networks.¹⁶⁸ Her work is useful, but this chapter should not treat AI as simply superseding writing. The argument here is different: AI depends on writing and archive even as it presents itself as post-historical immediacy.

Together, these works support a crucial claim:

AI is not the future arriving from nowhere; it is the past reorganized through a technical-economic apparatus that presents itself as future.

17. The voice from nowhere as historical achievement

The "voice from nowhere" is not the absence of history. It is a historical achievement.

It requires the formalization of knowledge, the digitization of archives, the extraction of data, the platforming of public discourse, the industrialization of computation, the normalization of interfaces, the commodification of output, the habituation of users, the concealment of labor, the erasure of provenance.

Only after all this can the user experience an answer as if it came from nowhere. This is the paradox of AI: the more historical work is required to produce the answer, the less historical the answer appears.

The interface is therefore not merely a convenience; it is the place where history is transmuted into immediacy. The user asks, and centuries of language return as a sentence without ancestry.

18. The danger of ahistorical authority

The de-historicized answer becomes dangerous when it acquires authority.

If the user treats AI as a playful tool, the danger is limited. If the user treats it as orientation, then returns to sources, the danger is manageable. But when AI-generated synthesis becomes institutional authority, in education, governance, hiring, medicine, law, management, or research, the erasure of genealogy becomes politically serious.

Who authored the synthesis? What sources were excluded? What historical conflict was flattened? What local knowledge was erased? What method produced the answer? Who benefits from the interface? Who cannot contest it? Who is made legible only through its categories?

Ahistorical authority is authority without visible origin.

It is difficult to challenge because its sources are hidden. It appears neutral, universal, and objective precisely because its genealogy has been erased.

This is the institutional form of the voice from nowhere.

The answer becomes policy, assessment, recommendation, classification, or decision support.

History disappears into administration.

19. Re-historicizing the interface

The proper response is not nostalgia for a pre-digital world. Nor is it the rejection of every universalizing project. Universal access, shared knowledge, and synthesis can be emancipatory. The problem is synthesis without genealogy, universality without situatedness, and knowledge without historical responsibility.

The task is to re-historicize the interface.

This means making the conditions of output visible:

sources, methods, uncertainties, data origins, labor, infrastructures, exclusions, localities, political assumptions, economic interests.

It also means changing the user's relation to AI. The user must treat output not as final answer, but as historically conditioned material requiring return.

Return to sources. Return to authors. Return to archives. Return to disputes. Return to local knowledge. Return to the question of method. Return to one's own responsibility.

Re-historicization does not mean reconstructing the full genealogy of every token. That is impossible. It means refusing the interface's claim to knowledge without genealogy.

The answer must no longer appear self-grounding.

20. *Toward releasement*

This chapter prepares the final movement of the book.

If AI is historically produced but appears de-historicized, then the response cannot be mere regulation, technical transparency, or better prompting. These may be necessary, but they do not exhaust the philosophical task.

The task is to cultivate a different relation to technological knowledge.

Heidegger's distinction between calculative and meditative thinking becomes decisive here. Calculative thinking plans, computes, organizes, and optimizes. It is necessary. But if it becomes the only mode of thought, it prevents us from asking what is being revealed, concealed, and transformed by our tools.¹⁶⁹

AI is the machine of calculative language.

It computes linguistic possibility.

The final chapter will therefore turn to *Gelassenheit*: releasement toward technological things and openness to what calculation conceals. Releasement does not mean refusal. It means using the tool without surrendering to its mode of revealing.

The task is not to abolish AI.

The task is to prevent AI from defining knowledge as de-historicized output.

21. *Preliminary thesis*

The chapter establishes that AI is historically produced but appears as knowledge without history. It emerges from a long genealogy: the modern world picture, technological enframing, the formalization of knowledge, the dream of universal calculation, the erasure of the situated knower, the universal interface, data extraction, and the historical management of textual abundance. The interface hides this genealogy. The voice from nowhere is not the absence of history, but the concealment of history. AI belongs to the modern dream of formalized, universal knowledge without the burden of historical interpretation, transforming archives, labor, infrastructures, and local histories into immediate answer. Synthesis without author conceals source, method, conflict, and apprenticeship, while the universal interface risks imposing a single computational grammar on heterogeneous historical worlds.

Responsible use therefore requires re-historicizing the interface and refusing knowledge without genealogy. The final chapter asks how such refusal is possible without romantic withdrawal. If AI is both historically produced and practically useful, the task is not rejection but releasement: the ability to use the machine without letting its de-historicized mode of revealing become our only relation to language, knowledge, and world.

FOOTNOTES

¹⁴⁹ Martin Heidegger, “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. Heidegger’s argument that modernity is defined by the transformation of the world into picture (*Weltbild*) is the chapter’s primary historical-ontological frame. The chapter does not claim that AI is the simple fulfillment of this transformation. It claims that AI is intelligible only within the modern history in which beings are increasingly disclosed as representable, calculable, and available. See also Iain D. Thomson, *Heidegger on Ontotheology: Technology and the Politics of Education* (Cambridge: Cambridge University Press, 2005), 78 to 113, for a careful reading of the world picture argument in relation to enframing.

¹⁵⁰ 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. Chapter 4 developed the *Bestand* argument in relation to corpus, token, and vector. Chapter 10 extends this argument to history itself: the historical record becomes standing-reserve when held available for output. The chapter does not claim that all historical preservation is enframing. It claims that the AI pipeline transforms historical material from what is encountered as distance into what is delivered as response.

¹⁵¹ Edmund Husserl, *The Crisis of European Sciences and Transcendental Phenomenology*, trans. David Carr (Evanston, IL: Northwestern University Press, 1970), 48 to 59, 353 to 378. Husserl’s diagnosis of the crisis of formalization in the modern sciences provides the chapter’s phenomenological frame. The crisis is not that science is false, but that its abstractions forget their genesis in lived experience. The chapter extends this point to language: AI’s linguistic output rests on abstractions from human worlds, but the formal apparatus tends to conceal that derivation.

¹⁵² Max Horkheimer and Theodor W. Adorno, *Dialectic of Enlightenment: Philosophical Fragments*, trans. Edmund Jephcott (Stanford, CA: Stanford University Press, 2002), 1 to 34. The chapter uses the dialectic-of-enlightenment thesis selectively. It does not adopt every aspect of Adorno and Horkheimer’s argument. It mobilizes the structural insight that rational mastery can turn into domination when reason reduces meaning to administered function. AI gives this transformation a universal linguistic interface.

¹⁵³ Hans-Georg Gadamer, *Truth and Method*, 2nd rev. ed., trans. Joel Weinsheimer and Donald G. Marshall (New York: Continuum, 2004), 268 to 306. Gadamer's account of historically effected consciousness and the fusion of horizons is the necessary hermeneutic counterweight to AI's apparent immediacy. The chapter does not claim that all AI use is anti-hermeneutic. It claims that the dream of de-historicized knowledge that AI partially realizes is structurally opposed to hermeneutic experience.

¹⁵⁴ For Leibniz's *characteristica universalis* and *calculus ratiocinator*, see Markku Roinila, *Leibniz on Rational Decision-Making* (Helsinki: University of Helsinki, 2007), chap. 2. Umberto Eco's history of universal language projects provides the wider context. See Umberto Eco, *The Search for the Perfect Language*, trans. James Fentress (Oxford: Blackwell, 1995). The chapter uses Leibniz not as the simple origin of AI but as one historical articulation of the epistemological fantasy that disputes might be resolved by calculation. AI inherits this fantasy in transformed form: from formal certainty to statistical availability.

¹⁵⁵ For the *Encyclopédie* as Enlightenment project, see Eco, *The Search for the Perfect Language*, especially the discussion of universal classification projects. The chapter treats the *Encyclopédie* in nuanced terms. Its democratization of access was emancipatory and should not be reduced to a precursor of computational synthesis. But it also participated in the dream that knowledge could be made systematic and universally available, a dream that AI inherits in computational form.

¹⁵⁶ For Hilbert's formalist program and Gödel's incompleteness theorems, see Richard Zach, "Hilbert's Program Then and Now," in *Philosophy of Logic*, ed. Dale Jacquette (Amsterdam: North-Holland, 2007), 411 to 447. The chapter uses this material to make a precise historical point: AI does not promise the formal completeness that Hilbert sought. It promises probabilistic synthesis. This is a shift from the dream of formal certainty to the dream of statistical availability, but the underlying desire for knowledge without historical interpretation remains.

¹⁵⁷ Donna Haraway, "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective," *Feminist Studies* 14, no. 3 (1988): 575 to 599. Haraway's critique of the "god trick" is decisive for the chapter's argument. AI performs a technological version of this trick: it appears to speak from nowhere when it is in fact trained on particular corpora, in particular institutions, under particular economic incentives. The chapter integrates feminist standpoint epistemology with the Heideggerian frame because both expose the masked particularity of apparent universality.

¹⁵⁸ Lucy Suchman, *Human-Machine Reconfigurations: Plans and Situated Actions*, 2nd ed. (Cambridge: Cambridge University Press, 2007), 51 to 84. Suchman's account of human action as situated and machine operation as plan-based is essential for understanding the asymmetry that the interface manages. The chapter uses Suchman to clarify that the AI interface is not a neutral window into the model but a cultural-technical form that disciplines situated action into prompt-response grammar.

- ¹⁵⁹ Branden Hookway, *Interface* (Cambridge, MA: MIT Press, 2014), 1 to 24. Hookway's treatment of the interface as boundary and bottleneck supports the chapter's argument that the AI interface is not merely a tool but a structuring form. It organizes what can be asked, what feels authoritative, and what remains invisible.
- ¹⁶⁰ Wendy Hui Kyong Chun, *Updating to Remain the Same: Habitual New Media* (Cambridge, MA: MIT Press, 2016), 1 to 30. Chun's analysis of habituation in new media is important because it shows that interface familiarity is not neutral. It trains expectation, makes complex operations feel ordinary, and produces forms of forgetting. The chapter extends this to AI: interface habit produces historical amnesia at the level of expectation.
- ¹⁶¹ Benjamin H. Bratton, *The Stack: On Software and Sovereignty* (Cambridge, MA: MIT Press, 2015), 219 to 250. Bratton's account of the interface as a layer in a planetary computational megastructure is useful for the chapter's argument because it situates the AI interface within larger systems of sovereignty, address, and userhood. The interface is not only a personal experience. It is a geopolitical form.
- ¹⁶² Nick Couldry and Ulises A. Mejias, *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism* (Stanford, CA: Stanford University Press, 2019), 1 to 33. The data colonialism thesis supports the chapter's argument that AI's apparent universality is sustained by planetary extraction. The "voice from nowhere" is built from very specific places of labor, language, and infrastructure.
- ¹⁶³ Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (New Haven: Yale University Press, 2021), 1 to 40, 89 to 122. Crawford's account is indispensable for breaking the fantasy of AI as immaterial intelligence. The chapter uses her work to ground the claim that the dream of knowledge without history is sustained by intensely historical and material processes. The clean interface depends on the displaced extraction.
- ¹⁶⁴ Yuk Hui, *The Question Concerning Technology in China: An Essay in Cosmotechnics* (Falmouth: Urbanomic, 2016), 1 to 36, 257 to 309. Hui's concept of cosmotechnics is essential because it resists the false universalism of computational technology. The chapter uses Hui to argue that the AI interface risks imposing a single computational grammar on heterogeneous worlds. The danger is epistemic synchronization. Cosmotechnics reminds us that technology could be otherwise.
- ¹⁶⁵ Ann M. Blair, *Too Much to Know: Managing Scholarly Information before the Modern Age* (New Haven: Yale University Press, 2010), 1 to 28, 173 to 229. Blair's work on early modern information management is useful because it shows that AI's synthesis belongs to a long history of textual abundance and reference-form solutions. Compendia, encyclopedias, indexes, and note-taking systems all addressed similar problems. AI both inherits and transforms this history: it gives final synthesis without visible architecture.

¹⁶⁶ Matteo Pasquinelli, *The Eye of the Master: A Social History of Artificial Intelligence* (London: Verso, 2023), 1 to 24, 187 to 216. Pasquinelli's account of AI as the objectification and automation of social knowledge supports the chapter's claim that machine intelligence is not alien intelligence. It is historical human intelligence reorganized through capital and computation. This claim is crucial for the chapter's main argument: AI is not the future arriving from nowhere. It is the past reorganized.

¹⁶⁷ Adrian Daub, *What Tech Calls Thinking: An Inquiry into the Intellectual Bedrock of Silicon Valley* (New York: FSG Originals, 2020), 1 to 28. Daub's critique is useful because Silicon Valley repeatedly borrows philosophical language and narratives of inevitability to legitimize its business models. AI is often presented as destiny rather than as historically situated technical-economic project. The chapter mobilizes Daub to resist the rhetoric that naturalizes AI as the necessary next stage of intelligence.

¹⁶⁸ Catherine Malabou, *Morphing Intelligence: From IQ Measurement to Artificial Brains*, trans. Carolyn Shread (New York: Columbia University Press, 2019); Catherine Malabou, *Plasticity at the Dusk of Writing: Dialectic, Destruction, Deconstruction*, trans. Carolyn Shread (New York: Columbia University Press, 2010). Malabou's work on plasticity helps frame AI as part of a transition from textuality to adaptive networks. The chapter uses Malabou with one significant qualification: AI depends on writing and archive even as it presents itself as post-historical immediacy. The argument is not that AI supersedes textuality but that it operates on writing while disguising that operation.

¹⁶⁹ Martin Heidegger, *Discourse on Thinking*, trans. John M. Anderson and E. Hans Freund (New York: Harper & Row, 1966), 43 to 57. Heidegger's distinction between calculative and meditative thinking is the bridge from this chapter to the final one. The chapter does not claim that calculative thinking should be abandoned. It claims that calculative thinking becomes dangerous when it is the only mode of thought permitted. For a careful reading of *Gelassenheit* and the will, see Bret W. Davis, *Heidegger and the Will: On the Way to Gelassenheit* (Evanston, IL: Northwestern University Press, 2007), 1 to 30, 285 to 314.

CHAPTER-LEVEL BIBLIOGRAPHY

Adorno, Theodor W., and Max Horkheimer. *Dialectic of Enlightenment: Philosophical Fragments*. Translated by Edmund Jephcott. Stanford, CA: Stanford University Press, 2002.

Blair, Ann M. *Too Much to Know: Managing Scholarly Information before the Modern Age*. New Haven: Yale University Press, 2010.

Bratton, Benjamin H. *The Stack: On Software and Sovereignty*. Cambridge, MA: MIT Press, 2015.

Chun, Wendy Hui Kyong. *Updating to Remain the Same: Habitual New Media*. Cambridge, MA: MIT Press, 2016.

Couldry, Nick, and Ulises A. Mejias. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA: Stanford University Press, 2019.

Crawford, Kate. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press, 2021.

Daub, Adrian. *What Tech Calls Thinking: An Inquiry into the Intellectual Bedrock of Silicon Valley*. New York: FSG Originals, 2020.

Davis, Bret W. *Heidegger and the Will: On the Way to Gelassenheit*. Evanston, IL: Northwestern University Press, 2007.

Eco, Umberto. *The Search for the Perfect Language*. Translated by James Fentress. Oxford: Blackwell, 1995.

Gadamer, Hans-Georg. *Truth and Method*. 2nd rev. ed. Translated by Joel Weinsheimer and Donald G. Marshall. New York: Continuum, 2004.

Haraway, Donna. "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective." *Feminist Studies* 14, no. 3 (1988): 575 to 599.

Heidegger, Martin. "The Age of the World Picture." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 115 to 154. New York: Harper & Row, 1977.

Heidegger, Martin. *Discourse on Thinking*. Translated by John M. Anderson and E. Hans Freund. New York: Harper & Row, 1966.

Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.

Hookway, Branden. *Interface*. Cambridge, MA: MIT Press, 2014.

Hui, Yuk. *The Question Concerning Technology in China: An Essay in Cosmotechnics*. Falmouth: Urbanomic, 2016.

Husserl, Edmund. *The Crisis of European Sciences and Transcendental Phenomenology*. Translated by David Carr. Evanston, IL: Northwestern University Press, 1970.

Malabou, Catherine. *Morphing Intelligence: From IQ Measurement to Artificial Brains*. Translated by Carolyn Shread. New York: Columbia University Press, 2019.

Malabou, Catherine. *Plasticity at the Dusk of Writing: Dialectic, Destruction, Deconstruction*. Translated by Carolyn Shread. New York: Columbia University Press, 2010.

Pasquinelli, Matteo. *The Eye of the Master: A Social History of Artificial Intelligence*. London: Verso, 2023.

Roinila, Markku. *Leibniz on Rational Decision-Making*. Helsinki: University of Helsinki, 2007.

Suchman, Lucy. *Human-Machine Reconfigurations: Plans and Situated Actions*. 2nd ed. Cambridge: Cambridge University Press, 2007.

Thomson, Iain D. *Heidegger on Ontotheology: Technology and the Politics of Education*. Cambridge: Cambridge University Press, 2005.

Zach, Richard. "Hilbert's Program Then and Now." In *Philosophy of Logic*, edited by Dale Jacquette, 411 to 447. Amsterdam: North-Holland, 2007.

Chapter 11. Releasement and Reappropriation

Toward a Free Relation to Artificial Intelligence

Chapter 10 argued that AI is historically produced but appears as de-historicized knowledge. It is made from archives, languages, labor, platforms, infrastructures, scientific formalization, capital, and the accumulated sediment of human worlds. Yet at the interface, this history disappears. The user receives a fluent answer, a voice from nowhere, a synthesis without visible author, method, source, or genealogy.

This final chapter asks what follows.

One answer would be rejection. If AI de-historicizes language, industrializes *Gerede*, fetishizes the historical labor of speech, and tempts the user with the desire to have already understood, then perhaps it should be refused.

Another answer would be acceptance. If AI is useful, efficient, and unavoidable, perhaps the task is merely to regulate it, align it, make it fairer, more transparent, more accurate, and more accountable.

Both answers are insufficient.

Rejection misunderstands AI by treating it as if the problem were simply the machine. Acceptance misunderstands AI by treating it as if usefulness were innocence. The machine is not *das Man*. It is not evil, inauthentic, demonic, or fallen. It is an intraworldly technical entity. But neither is it neutral. It belongs to a mode of revealing that turns language into standing-reserve, generates *Gerede* at scale, and invites the user to mistake generated fluency for living knowledge.

The task is therefore neither technophilia nor technophobia.

The task is a free relation.

In Heideggerian terms, this means *Gelassenheit*: releasement toward technological things. But releasement alone is not enough. Because AI de-historicizes language, the free relation must also involve reappropriation: the active work of returning generated language to sources, histories, practices, questions, responsibility, and one's own saying.

The chapter's thesis is therefore:

A free relation to AI requires releasement from the fetish of generated knowledge and reappropriation of the language, history, and responsibility that the model de-historicizes.

1. The insufficiency of rejection

A first response to the diagnosis developed in this book would be refusal.

Do not use AI. Do not generate language. Do not ask the model. Do not let the machine enter reading, writing, thinking, education, or judgment.

This response is understandable. AI does threaten practices of reading, writing, interpretation, and formation. It can conceal ignorance, simulate competence, and replace apprenticeship with output. It can intensify *Gerede*. It can turn language into standing-reserve. It can become a fetish.

But simple refusal remains inadequate.

First, it risks treating AI as an external corruption of an otherwise pure human relation to language. That is false. The tendencies AI intensifies, such as publicness, idle talk, desire for ease, avoidance of not-knowing, and fetishization of language, already belong to everyday *Dasein*. AI gives them a machine. It does not create them from nothing.

Second, refusal may leave the structure of technological revealing unthought. One can reject a tool while still remaining governed by the same will to mastery, the same desire for control, the same metaphysics of purity, the same opposition between human authenticity and technical corruption.

Third, refusal is often impossible in practice. AI is becoming part of ordinary equipmental environments: search and writing, education and administration, code and translation, research and communication. A serious philosophy cannot rest on the fantasy of total exit.

Refusal may be necessary in particular situations: there are moments when the right act is not to use AI. But refusal is not, by itself, a free relation. One can say no and still remain captive.

2. *The insufficiency of compliance*

The opposite response is compliance. Use AI, but regulate it: make it fair, make it transparent, reduce bias, improve accuracy, add human oversight, require documentation, audit the system, align the model.

These measures are necessary. A system that produces harmful, discriminatory, fabricated, or unsafe outputs must be governed. Technical ethics, law, policy, and institutional safeguards are not optional.

But compliance is insufficient.

Compliance treats the problem primarily as a problem of system behavior. The machine must be corrected. Its outputs must be controlled. Its risks must be managed. Its harms must be reduced.

The analysis of this book has shown that the deeper problem is also a problem of relation.

Even a fairer model can generate *Gerede*. Even an accurate output can become unowned truth. Even a transparent system can be fetishized. Even a compliant interface can de-historicize knowledge. Even a safe model can weaken the user's capacity to dwell with not-knowing.

The point is not that fairness, transparency, and safety are irrelevant. The point is that they do not answer the existential question: what kind of relation to language does AI cultivate?

A perfectly compliant system may still train the user to expect language without delay, knowledge without apprenticeship, summary without reading, and competence without formation.

Compliance can govern the machine; it cannot by itself form the user. This is why AI ethics must become an ethics of formation.

3. *The insufficiency of cynicism*

A third response is cynicism.

The cynic knows. The cynic knows that AI hallucinates, that LLMs generate *Gerede*, that the interface hides labor and data extraction and de-historicization and capital, that the output is not living knowledge.

And yet the cynic continues.

They use AI to automate writing, research, communication, analysis, and response. They maintain ironic distance. They say: of course it is generated; of course it is not really intelligent; of course it is flawed; of course the system is extractive. Their knowledge becomes a badge of sophistication.

But their practice remains unchanged.

Sloterdijk's account of cynicism as enlightened false consciousness and Žižek's analysis of ideology are useful here. The modern subject may know very well how the system functions, yet continue to participate in it as if knowledge alone were sufficient resistance.¹⁷⁰

In AI use, this is especially dangerous.

Knowing that the model produces *Gerede* does not prevent one from using it as *Gerede*. Knowing that the answer is de-historicized does not re-historicize it. Knowing that the system is extractive does not undo one's dependence on its workflow. Knowing that the output is not owned does not make it one's own.

Cynicism leaves the structure intact. It says no in thought and yes in practice. A free relation requires more than ironic distance: it requires interruption.

4. The meaning of Gelassenheit

The needed response is *Gelassenheit*.

The term has a long theological prehistory, especially in Meister Eckhart and the Rhineland mystical tradition, where releasement names a letting-go of possessive self-will and a transformation of the relation between self, world, and divine ground.¹⁷¹ Heidegger secularizes and reworks this lineage. In his later thought, *Gelassenheit* becomes a way of thinking beyond the will to mastery that governs modern technological revealing.

In *Discourse on Thinking*, Heidegger calls for releasement toward things and openness to the mystery. This is not passivity. It is not indifference. It is not nostalgia. It is a way of using technological things while refusing to be claimed entirely by their mode of revealing.¹⁷²

Releasement says yes and no at once: yes to technical use, no to technological totalization; yes to tools, no to the reduction of beings to resources; yes to calculation where calculation is appropriate, no to the reign of calculation as the only form of thought.

In relation to AI this double movement is essential: yes to AI as tool, no to AI as subject; yes to AI as orientation, no to AI as appropriation; yes to AI as generated material, no to AI as living knowledge; yes to AI as map, no to AI as world.

Releasement means using the machine without allowing the machine's mode of revealing to define what language, knowledge, reading, writing, and thinking are.

5. Calculative and meditative thinking

Heidegger distinguishes calculative thinking from meditative thinking.

Calculative thinking computes, plans, organizes, optimizes, and moves from one prospect to another. It is necessary. Modern life cannot do without it. Science and engineering, administration and finance, medicine and communication all depend on calculative thinking.

But calculative thinking does not meditate on the meaning of what is. It does not ask what kind of world its calculations disclose. It does not ask what becomes of beings when everything is rendered orderable, predictable, and available.

AI is a privileged contemporary machine of calculative language.

It computes linguistic possibility. It predicts, generates, ranks, summarizes, translates, compares, and optimizes. It turns language into a field of operations. It is not merely calculative thinking as an idea; it is calculative language deployed as infrastructure.

Meditative thinking does not reject calculation. It interrupts its totalization.

It asks: what has language become here? What is hidden by the answer? What is flattened by the summary? What is lost when reading becomes output? What happens to not-knowing, to the ability to own a saying, to the kind of human being formed by this interface?

A free relation to AI requires this meditative interruption. The user must not only ask what the model can do; the user must ask what the model is making of them.

6. Releasement is not quietism

Releasement can be misunderstood as passivity.

Let things be, do not intervene, do not resist, do not act, withdraw. This is not the relation required here.

A free relation to AI is not passive quietism. It is not spiritual resignation before technological destiny. It is not the refusal to govern, design, regulate, teach, or resist. It is not a private serenity that leaves the world unchanged.

Releasement is a change in attunement that makes action possible without domination.

One can use AI without worshipping it. One can regulate AI without imagining regulation is enough. One can design AI without believing design is neutral. One can teach with AI without surrendering formation to the machine. One can write with AI without abandoning one's own saying.

This is why releasement must be paired with reappropriation.

Releasement breaks the fetish; reappropriation does the work that follows.

7. Reappropriation

Reappropriation is the active labor of making generated language answerable again.

AI de-historicizes; reappropriation re-historicizes. AI gives output; reappropriation returns to sources. AI produces fluent language; reappropriation asks whether the saying can be owned. AI summarizes; reappropriation reads. AI offers options; reappropriation decides. AI generates public intelligibility; reappropriation tests, situates, verifies, and assumes responsibility.

Reappropriation is not merely editing. It is not polishing generated text. It is not making the output sound more human. It is the transformation of generated material into something for which a human being can answer.

This requires work.

The user must ask:

Where does this claim come from? What source would support it? What tradition does it invoke? What has been flattened? What do I still not understand? What must I read? Can I defend this? Can I say this as my own?

If the answer is no, the output must remain material, not knowledge. Reappropriation is the return of language to answerability.

8. Re-historicization

The first discipline of reappropriation is re-historicization.

AI output appears de-historicized. It often arrives without visible source, author, method, tradition, conflict, or genealogy. The user must restore the question of origin.

This does not mean reconstructing the full genealogy of every token. That is impossible. It means refusing to let the output appear self-grounding.

Re-historicization asks:

What history has been compressed into this answer? What authors should be consulted? What tradition is being simplified? What disagreement has been smoothed? What local context has disappeared? What labor is hidden? What political or economic structure made this output possible?

MacIntyre's account of tradition is useful here. Rational inquiry does not happen from nowhere. It unfolds within historically extended practices and traditions, with standards, conflicts, inheritances, and internal goods.¹⁷³

AI output must be returned to such traditions.

A generated account of Aristotle must return to Aristotle and Aristotelian traditions. A generated interpretation of a political conflict must return to sources, histories, and contested memories. A generated legal analysis must return to law. A generated ethical recommendation must return to the situation and those affected.

Re-historicization breaks the voice from nowhere; it gives the answer back its history.

9. Preserving not-knowing

The second discipline is preserving not-knowing.

Chapter 9 argued that AI satisfies the desire to have already understood by converting ignorance into output. Reappropriation must resist this conversion.

Not every gap should be filled immediately. Not every uncertainty should become a prompt. Not every blank page should be erased. Not every difficulty should become a summary.

Some not-knowing must be preserved.

This preservation is not ignorance for its own sake. It is the protection of the space in which understanding can begin.

The user must learn to say: I have an answer, but I do not yet understand. I have a summary, but I have not yet read. I have a draft, but I have not yet written. I have options, but I have not yet decided. I have language, but I have not yet made it mine.

Such sentences are not failures. They are signs of philosophical maturity.

Huttunen and Kakkori's application of *Gelassenheit* to education is useful here because it connects releasement with letting-learn: the refusal to reduce education to control, productivity, and calculable outcomes.¹⁷⁴ AI education must preserve the right to not know, to delay, to struggle, and to learn.

The ability to use AI well depends on the ability not to use it immediately.

10. Returning to reading

The third discipline is returning to reading.

AI can summarize, explain, compare, and contextualize. These functions can help a reader begin. But they become dangerous when they replace the encounter with the text.

A generated summary does not equal reading. A generated interpretation does not equal understanding. A generated comparison does not equal historical encounter.

Reading is the re-historicization of writing. It brings the text back into dialogue among author, reader, tradition, and question. AI can prepare this dialogue, but it cannot undergo it for the user.

The discipline is simple. Use AI to approach the text, not to avoid it. Use AI to formulate questions, not to substitute for the question. Use AI to reveal difficulty, not to eliminate difficulty before it has formed you. The reader must return to the page, to the source, to the authorial trace, to the historical distance. Otherwise the output becomes a polished refusal of reading.

11. Writing one's own saying

The fourth discipline is writing.

AI can draft. It can reformulate. It can organize. It can offer tones, openings, transitions, and conclusions. These functions can help. But writing is not merely the production of text. Writing is a practice through which thought becomes one's own.

A user who accepts generated language without appropriation may produce excellent prose without having spoken.

The question is not: did AI help?

The question is: can I own what is written?

To own a saying is not to have produced every word manually. A writer may use notes and books, quotations and translations, editors and conversations. Authorship has never meant pure isolation. But authorship requires responsibility. The writer must be able to stand behind the final saying.

The user must ask: do I know why this sentence is here, can I defend this claim, what does this paragraph assume, what have I borrowed, what have I understood, what do I need to delete, what must I write badly before I can write truly?

AI can provide language; only the user can make language theirs. If the user cannot answer for the text without the machine, the text is not yet theirs.

12. Verification and appropriation

Verification is necessary, but not sufficient.

Generated claims must be checked. Citations must be verified. Sources must be read. Dates, names, facts, and interpretations must be confirmed. Technical hallucination requires technical and scholarly discipline.

But verification is not appropriation.

A user can verify a statement and still not understand it. A user can check a source and still not own the argument. A user can confirm that a sentence is true and still be unable to answer for it.

Verification asks: is this correct? Appropriation asks: is this mine to say? Both are necessary. Verification protects against falsehood; appropriation protects against unowned truth.

A free relation to AI requires both epistemic and existential responsibility. The user must check the output and undergo the work required to stand behind it.

A true generated answer is not enough; truth must become owned.

13. Ethics of formation

The fifth discipline is an ethics of formation.

Mainstream AI ethics often focuses on fairness, transparency, accountability, privacy, safety, and bias mitigation. These are necessary. But they remain insufficient if they treat the user as already formed and the machine as a neutral object to be tuned.

Technological environments form users.

Interfaces shape habits. Predictive systems shape expectation. Recommendation systems shape attention. Chatbots shape the experience of asking, writing, and knowing. Hypernudging and predictive design can narrow the horizon of moral possibility by making some actions easier, smoother, and more available than others.¹⁷⁵

A free relation to AI therefore requires an ethics of formation.

The question is not only whether the model is fair.

The question is what kind of human being is formed by regular interaction with the model.

Does the user become more capable of reading? More capable of writing? More capable of not-knowing? More capable of judgment? More capable of responsibility? Or does the user become dependent on fluency, speed, and generated public language?

Verbeek's work on technology and morality is useful here because technologies actively mediate ethical subjectivity.¹⁷⁶ Coeckelbergh similarly emphasizes relational and appearance-based dimensions of machine ethics.¹⁷⁷ The issue is not only what AI is, but what relations it forms.

Ethics must therefore include practices of friction: slowness, contestability, source-return, human answerability, non-use, reading, writing, silence. Formation requires resistance.

14. Focal practices

Borgmann's device paradigm gives another practical vocabulary.

Modern devices often deliver commodities while concealing the practices and engagements that once made those commodities meaningful. A heating system delivers warmth without the focal practice of tending a fire. A device gives the product while hiding the world of engagement.¹⁷⁸

AI can function as a linguistic device.

It delivers text without writing. Summary without reading. Advice without judgment. Options without decision. Competence without formation.

The response is not to reject all devices. It is to preserve focal practices: practices that gather attention, body, skill, world, and community.

Reading, writing, conversation, teaching, craft, and public argument can all be focal practices.

AI should support such practices only when it does not replace the engagement that gives them meaning.

A free relation to AI requires protecting non-algorithmic spaces where language is not immediately optimized, summarized, generated, or made useful.

There must remain places where one reads slowly, speaks hesitantly, writes badly, listens long, and thinks without output.

15. The rule of non-substitution

The practical rule of the chapter can be stated simply:

Use AI for assistance, not substitution.

But this rule must be specified.

AI may assist orientation; it must not substitute for reading. AI may assist drafting; it must not substitute for authorship. AI may assist comparison; it must not substitute for judgment. AI may assist retrieval; it must not substitute for historical inquiry. AI may assist formulation; it must not substitute for one's own saying. AI may assist learning; it must not substitute for formation.

The rule of non-substitution does not forbid use. It defines its limit.

Where the task is mechanical, repetitive, preliminary, or merely organizational, AI may be highly appropriate. Where the task requires appropriation, responsibility, judgment, ethical decision, historical understanding, or existential ownership, AI can only assist. It cannot replace the human act.

This distinction will matter especially in education, professional judgment, writing, research, and human assessment. But its basis is general.

AI can generate language. It cannot make language yours.

16. The rejection of cynical distance

A free relation also requires rejecting cynicism.

The cynic says: I know the model is flawed, but I use it anyway. I know it produces *Gerede*, but everyone uses it. I know it de-historicizes, but I need the output. I know the system is extractive, but I keep my distance mentally.

This is not freedom.

It is adaptation with irony.

Cynicism is especially seductive because it looks like critique. The cynical user is not naïve. They know. But their knowing does not change their practice. They continue to let the machine write, summarize, decide, and relieve them of the burden of their own saying.

A free relation requires more than knowing.

It requires interruption.

It requires the willingness to lose convenience: to read when a summary would be faster, to write when a draft would be smoother, to admit ignorance when output is available, to refuse a generated sentence that one cannot own, to let a question remain unanswered.

Without such interruptions, critique becomes decorative. The user says no in theory and yes in workflow. That is not releasement.

17. Reclaiming silence

The final discipline is silence.

Chapter 1 argued that silence is not mere absence of speech. It can be a mode of *Rede*, a way of letting the matter appear. Chapter 6 showed that AI is structurally ordered toward continuation. The model continues. The interface invites more. The user can always ask again.

A free relation requires reclaiming silence. Not every question should be answered immediately, not every hesitation should be filled, not every thought should be accelerated, not every confusion should be converted into prose. Silence protects the question.

It allows the user to remain with what has not yet become language. It lets the matter resist. It preserves not-knowing from immediate conversion into output.

This is difficult because contemporary culture often treats silence as inefficiency. AI intensifies that pressure. The machine is always ready. It can always produce something. The user must therefore learn a new kind of restraint.

Sometimes the right response to AI is not another prompt. It is silence, not passive silence but a thinking silence.

18. Reclaiming conversation

If AI simulates dialogue, a free relation also requires reclaiming human conversation.

Human conversation is not merely exchange of information. It involves vulnerability and trust, misunderstanding and correction, conflict and memory, embodiment and shared world. Another person can resist. Another person can be wounded. Another person can call us to account.

AI can simulate conversation, but it cannot be another *Dasein*. It cannot be vulnerable. It cannot be transformed as a self. It cannot bear responsibility for what is said.

This does not mean conversation with AI is useless. It can be helpful as rehearsal, brainstorming, and orientation. But if it replaces human conversation, something essential is lost.

The user must preserve encounters in which the other is truly other: a teacher, a friend, a reader, a critic, a student, a colleague, a witness, a stranger.

Such conversations are slower, riskier, and less controllable than AI interaction. That is why they matter. They resist the narcissistic mirror of the interface. They reintroduce otherness, responsibility, and response.

A free relation to AI requires returning generated dialogue to real dialogue.

19. Final theses

The book's argument can now be gathered.

Language is born from worldly significance. The world into which language is born is public. The They produces the already-said. Writing sediments it. Processing de-historicizes it. Corpus encloses it. Tokenization cuts it. Vectorization turns it into calculable proximity. AI becomes the vectorized afterlife of the language of the They. Generated speech functions as industrialized *Gerede*. The machine becomes an intraworldly fetish of language. The user vivifies it as if it were living knowledge. AI satisfies the desire to have already understood. It appears as knowledge without history.

The response is neither refusal alone, nor compliance alone, nor cynical awareness. The response is releasement and reappropriation. Releasement breaks the grip of the fetish by saying yes and no at once. Reappropriation returns generated language to history, source, world, responsibility, and one's own saying.

The final thesis is this:

AI can generate the already-said. It cannot relieve us of the burden of our own saying.

20. Conclusion

This book has not argued that AI is useless, unintelligent, or simply false. Its power is real. Its usefulness is real. Its capacity to assist writing, reading, searching, translating, summarizing, and organizing is real.

But its danger is also real.

AI is dangerous not only when it hallucinates falsehood, but when it gives us unowned truth. Not only when it produces errors, but when it relieves us from the experience through which knowledge becomes ours. Not only when it replaces labor, but when it replaces formation. Not only when it speaks for us, but when we forget how to speak as ourselves.

AI is not the They. It is the vectorized afterlife of the language of the They: public language written, extracted, de-historicized, tokenized, vectorized, commodified, and reanimated as generated speech.

Its power lies not only in calculation; it lies in relief.

The relief of not having to remain ignorant. The relief of not having to read. The relief of not having to begin. The relief of not having to decide. The relief of not having to own a saying. The relief of having already understood.

A free relation begins when this relief is interrupted.

Use the tool.

Break the fetish.

Preserve not-knowing.

Return to reading.

Restore provenance.

Interrupt fluency.

Reclaim silence.

Own what you say.

Only the human user can re-historicize what the model has de-historicized. Only the human user can return generated language to a world. Only the human user can decide whether the output becomes thought, or remains *Gerede*.

The final question is not whether AI can speak. The final question is whether we can still speak.

FOOTNOTES

¹⁷⁰ Peter Sloterdijk, *Critique of Cynical Reason*, trans. Michael Eldred (Minneapolis: University of Minnesota Press, 1987); Slavoj Žižek, *The Sublime Object of Ideology* (London: Verso, 1989). Sloterdijk's account of cynicism as enlightened false consciousness ("they know very well what they are doing, but still, they are doing it") and Žižek's subsequent elaboration are decisive for the chapter's third refusal. The chapter argues that critical AI scholarship is structurally vulnerable to this cynical form: the user knows that LLMs hallucinate, de-historicize, and generate *Gerede*, and nevertheless integrates them into workflow without changing practice. The chapter's claim is that this knowing is not resistance but adaptation with irony.

¹⁷¹ Meister Eckhart, *The Complete Mystical Works of Meister Eckhart*, trans. Maurice O'C. Walshe, rev. Bernard McGinn (New York: Crossroad, 2009); Bernard McGinn, *The Mystical Thought of Meister Eckhart: The Man from Whom God Hid Nothing* (New York: Crossroad, 2001). The chapter recognizes the theological prehistory of *Gelassenheit* in Eckhart and the Rhineland mystical tradition. The chapter does not adopt the theological frame but acknowledges that Heidegger's late vocabulary inherits this lineage. The recognition matters because *Gelassenheit* in Heidegger is not a neutral philosophical term; it carries the trace of a specific spiritual practice of letting-go that the chapter mobilizes in secularized form.

¹⁷² Martin Heidegger, *Discourse on Thinking*, trans. John M. Anderson and E. Hans Freund (New York: Harper & Row, 1966), 43 to 57; 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. For careful contemporary readings, see Bret W. Davis, *Heidegger and the Will: On the Way to Gelassenheit* (Evanston, IL: Northwestern University Press, 2007), 285 to 314; Daniela Vallega-Neu, *Heidegger's Contributions to Philosophy: An Introduction* (Bloomington: Indiana University Press, 2003); Krzysztof Ziarek, *Language after Heidegger* (Bloomington: Indiana University Press, 2013); Iain D. Thomson, *Heidegger on Technology's Danger and Promise in the Age of AI* (Cambridge: Cambridge University Press, 2025). Heidegger's distinction between calculative and meditative thinking, and his account of releasement as a yes-and-no toward technological things, is the chapter's primary philosophical resource. The chapter argues, however, that releasement alone is insufficient for AI because AI not only enframes but de-historicizes language. Reappropriation is the chapter's contribution to extending Heidegger's framework.

¹⁷³ Alasdair MacIntyre, *After Virtue: A Study in Moral Theory*, 3rd ed. (Notre Dame, IN: University of Notre Dame Press, 2007), 187 to 225. MacIntyre's account of tradition-constituted rational inquiry supports the discipline of re-historicization. Rational inquiry, for MacIntyre, unfolds within historically extended practices with standards, conflicts, inheritances, and internal goods. AI output must be returned to such traditions. The chapter narrows MacIntyre's point to the specific case of generated language and treats re-historicization as the user's responsibility, not the system's.

¹⁷⁴ Rauno Huttunen and Leena Kakkori, “*Gelassenheit*, Education, and Letting-Learn,” *Educational Philosophy and Theory* 53, no. 6 (2021): 630 to 642. Huttunen and Kakkori extend Heidegger’s *Gelassenheit* into educational philosophy through the concept of letting-learn. The chapter uses their work to clarify that preserving not-knowing is not an aesthetic preference but a pedagogical and ethical practice. The right to not know, to delay, to struggle, and to learn must be protected in AI-saturated educational environments.

¹⁷⁵ For the concept of hypernudging and predictive design, see the broader literature on choice architecture and algorithmic governance, especially Karen Yeung, “‘Hypernudge’: Big Data as a Mode of Regulation by Design,” *Information, Communication & Society* 20, no. 1 (2017): 118 to 136. The chapter uses the concept to indicate how AI interfaces can narrow the horizon of moral possibility without overt coercion, by making some actions easier and more available than others. The argument is that this narrowing is itself an ethical event that requires philosophical scrutiny, not only design adjustment.

¹⁷⁶ Peter-Paul Verbeek, *What Things Do: Philosophical Reflections on Technology, Agency, and Design*, trans. Robert P. Crease (University Park: Pennsylvania State University Press, 2005); Peter-Paul Verbeek, *Moralizing Technology: Understanding and Designing the Morality of Things* (Chicago: University of Chicago Press, 2011). Verbeek’s post-phenomenological account of technological mediation is essential for the chapter’s ethics of formation. Technologies are not neutral instruments; they actively mediate the user’s ethical subjectivity. The chapter uses Verbeek to support the claim that AI ethics must include the question of what kind of human being is formed by regular interaction with the model.

¹⁷⁷ Mark Coeckelbergh, “The Moral Standing of Machines: Towards a Relational and Non-Cartesian Moral Hermeneutics,” *Philosophy & Technology* 27, no. 1 (2014): 61 to 77. Coeckelbergh’s relational approach to machine ethics complements Verbeek’s mediation account. The chapter mobilizes both to argue that the issue is not only what AI is in itself but what relations it forms with the user. This shifts the locus of ethical analysis from the model to the relation.

¹⁷⁸ Albert Borgmann, *Technology and the Character of Contemporary Life: A Philosophical Inquiry* (Chicago: University of Chicago Press, 1984). Borgmann’s distinction between devices and focal practices is the chapter’s source for the analysis of AI as a linguistic device. Devices deliver commodities while concealing the engagements that once made those commodities meaningful. AI delivers text without writing, summary without reading, advice without judgment. The chapter extends Borgmann’s framework into the linguistic domain, where the commodity is generated language and the focal practices are reading, writing, conversation, and craft. See also Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Bloomington: Indiana University Press, 1990), and Bernard Stiegler, *Taking Care of Youth and the Generations*, trans. Stephen Barker (Stanford, CA: Stanford University Press, 2010); *Automatic Society, Volume 1: The Future of Work*, trans. Daniel Ross (Cambridge: Polity, 2016), for related accounts of technological mediation, care, and the formation of subjectivity.

CHAPTER-LEVEL BIBLIOGRAPHY

Borgmann, Albert. *Technology and the Character of Contemporary Life: A Philosophical Inquiry*. Chicago: University of Chicago Press, 1984.

Coeckelbergh, Mark. "The Moral Standing of Machines: Towards a Relational and Non-Cartesian Moral Hermeneutics." *Philosophy & Technology* 27, no. 1 (2014): 61 to 77.

Davis, Bret W. *Heidegger and the Will: On the Way to Gelassenheit*. Evanston, IL: Northwestern University Press, 2007.

Davis, Bret W. *Zen Pathways: An Introduction to the Philosophy and Practice of Zen Buddhism*. Oxford: Oxford University Press, 2022.

Eckhart, Meister. *The Complete Mystical Works of Meister Eckhart*. Translated by Maurice O'C. Walshe. Revised by Bernard McGinn. New York: Crossroad, 2009.

Heidegger, Martin. *Being and Time*. Translated by John Macquarrie and Edward Robinson. New York: Harper & Row, 1962.

Heidegger, Martin. *Discourse on Thinking*. Translated by John M. Anderson and E. Hans Freund. New York: Harper & Row, 1966.

Heidegger, Martin. "The Question Concerning Technology." In *The Question Concerning Technology and Other Essays*, translated by William Lovitt, 3 to 35. New York: Harper & Row, 1977.

Huttunen, Rauno, and Leena Kakkori. "Gelassenheit, Education, and Letting-Learn." *Educational Philosophy and Theory* 53, no. 6 (2021): 630 to 642.

Ihde, Don. *Technology and the Lifeworld: From Garden to Earth*. Bloomington: Indiana University Press, 1990.

MacIntyre, Alasdair. *After Virtue: A Study in Moral Theory*. 3rd ed. Notre Dame, IN: University of Notre Dame Press, 2007.

McGinn, Bernard. *The Mystical Thought of Meister Eckhart: The Man from Whom God Hid Nothing*. New York: Crossroad, 2001.

Sloterdijk, Peter. *Critique of Cynical Reason*. Translated by Michael Eldred. Minneapolis: University of Minnesota Press, 1987.

Stiegler, Bernard. *Automatic Society, Volume 1: The Future of Work*. Translated by Daniel Ross. Cambridge: Polity, 2016.

Stiegler, Bernard. *Taking Care of Youth and the Generations*. Translated by Stephen Barker. Stanford, CA: Stanford University Press, 2010.

Thomson, Iain D. *Heidegger on Technology's Danger and Promise in the Age of AI*. Cambridge: Cambridge University Press, 2025.

Vallega-Neu, Daniela. *Heidegger's Contributions to Philosophy: An Introduction*. Bloomington: Indiana University Press, 2003.

Verbeek, Peter-Paul. *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press, 2011.

Verbeek, Peter-Paul. *What Things Do: Philosophical Reflections on Technology, Agency, and Design*. Translated by Robert P. Crease. University Park: Pennsylvania State University Press, 2005.

Yeung, Karen. "'Hypernudge': Big Data as a Mode of Regulation by Design." *Information, Communication & Society* 20, no. 1 (2017): 118 to 136.

Ziarek, Krzysztof. *Language after Heidegger*. Bloomington: Indiana University Press, 2013.

Žižek, Slavoj. *The Sublime Object of Ideology*. London: Verso, 1989.