

Algorithmic Intentionality: Toward a Phenomenology of Generative AI and Its Possible Worlds

Ishak H.A MEDDAH¹, Nour El Houda REMIL²

University of Saida Dr Moulay Tahar

ishak.meddah@univ-saida.dz¹, nourelhouda.remil@univ-saida.dz²

Abstract

This paper investigates whether and how generative artificial intelligence systems – specifically large language models (LLMs) and conversational AI – can be examined through classical phenomenological concepts: intentionality, reduction, lifeworld (Lebenswelt), embodiment, and alterity. Rejecting both naive anthropomorphism (treating AI as conscious) and reductive functionalism (treating AI as mere statistics), I propose a phenomenology of AI as quasi-alterity – a technical entity that displays intentional-like performances without subjective consciousness. Drawing on Husserl's theory of intentionality, Heidegger's equipmental analysis, Merleau-Ponty's phenomenology of embodiment and language, and recent work on technical objects (Stiegler, Ihde, Waldenfels), I argue that AI systems require a revised epoché that brackets both humanistic and computational biases. The paper proceeds in nine sections: (1) a phenomenological framing of the AI question, (2) a Husserlian analysis of quasi-intentionality with detailed distinctions, (3) the lifeworld and second-order world generation, (4) a Heideggerian reading of AI as equipment with extended discussion of breakdown, (5) the paradox of disembodied language after Merleau-Ponty, (6) AI as a new form of technical alterity in dialogue with Lévinas and Waldenfels, (7) a complete methodological protocol for doing phenomenology with and on AI, (8) an extended engagement with contemporary phenomenological literature, and (9) a conclusion outlining future research. The paper aims to provide a systematic phenomenological framework for understanding human-AI encounters without premature ontological commitments.

Keywords: artificial intelligence, intentionality, lifeworld, alterity, generative AI, reduction, technical object, large language models, quasi-intentionality, phantom body

Received: 21/08/2025 **Accepted:** 26/04/2026 **Published:** 21/05/2026

1. Introduction:

Why Phenomenology and AI Now? A Phenomenological Framing

1.1 The Emergence of Generative AI as a Phenomenological Event

The public release of ChatGPT in November 2022 marked not merely a technological advance but a phenomenological event. For the first time, hundreds of millions of users encountered a system that produced fluent, contextually appropriate, often sophisticated language across

virtually any domain. Users reported experiences ranging from wonder ("it understands me") to unease ("it's pretending to be human") to instrumental efficiency ("it's a better search engine"). These reports are not merely anecdotal; they are phenomenological data about how a new technical object appears in the lifeworld.

Previous AI systems – ELIZA (1966), PARRY (1972), ALICE (1995), Siri (2011), Alexa (2014) – also provoked anthropomorphic reactions. But LLMs are different. Their fluency, generality, and coherence produce a stronger and more persistent appearance of mindedness. Even users who know, intellectually, that the AI "has no feelings" report catching themselves saying "please" and "thank you," or feeling guilty after a harsh prompt. This gap between theoretical belief and lived experience is precisely the domain of phenomenology.

1.2 The Two Poles: Anthropomorphism and Reductive Functionalism

Contemporary discourse on AI oscillates between two poles, both unsatisfactory for phenomenological description.

The anthropomorphic pole: Journalists, users, and even some researchers speak of AI as if it had beliefs, desires, emotions, and consciousness. Headlines announce that an AI "feels lonely" or "wants to be free." This pole commits what we may call the intentionality fallacy: ascribing subjective acts to systems that merely simulate their products. It is phenomenologically naive because it ignores the reduction – it takes the appearance of mindedness as proof of mindedness.

The reductive functionalist pole: Computer scientists, eliminativists, and skeptics insist that LLMs are "just next-token predictors," "stochastic parrots," or "autocomplete on steroids." This pole commits the reduction fallacy: assuming that describing the implementation (linear algebra, transformers, training data) exhausts the phenomenon. It is phenomenologically naive because it ignores the mode of givenness – the way the AI appears in use – and reduces the phenomenon to its physical or computational substrate.

Phenomenology offers a third path. Instead of asking whether AI is conscious – a question that may be undecidable, ill-posed, or irrelevant to lived experience – phenomenology asks: how does AI appear within the human lifeworld? What are the essential structures of our experience of and with AI? How does AI modify the intentional relations that constitute our everyday world? This shift from ontological to phenomenological questioning is the paper's foundational move.

1.3 Two Central Theses

The paper defends two interconnected theses:

Thesis 1: AI displays quasi-intentionality. LLMs exhibit a functional analogue of intentional directedness. They are about things; they maintain horizons of relevance; their outputs can be fulfilled or frustrated relative to user expectations. However, this quasi-intentionality lacks noetic subjectivity: there is no for-whom the aboutness obtains, no lived experience of the horizon, no satisfaction of fulfilled intention. Quasi-intentionality is intentionality's phenomenal shadow – structurally isomorphic in its noematic pole, hollow in its noetic pole.

Thesis 2: AI constitutes a new form of technical alterity. The AI is neither a tool like a hammer (transparent equipment) nor a person like a friend (another subject). It is a hybrid phenomenon that oscillates between these poles depending on context, task, and breakdown. I call this technical alterity: a mode of givenness in which the AI appears as quasi-other – resistant yet non-suffering, surprising yet non-intentional, responsive yet non-conscious. Technical alterity is not a failure of perception but an essential structure of human-AI encounter.

1.4 The Paper's Structure

The paper unfolds as follows. Section 2 provides a detailed Husserlian analysis of intentionality and develops the concept of quasi-intentionality with rigorous distinctions. Section 3 examines the lifeworld and introduces the concept of second-order lifeworlds generated by AI. Section 4 turns to Heidegger, analyzing AI as equipment with extended discussion of breakdown, the hermeneutics of AI, and the limits of the equipmental analogy. Section 5 addresses the paradox of disembodied language through Merleau-Ponty, proposing the concept of the phantom body. Section 6 develops technical alterity in dialogue with Lévinas and Waldenfels. Section 7 offers a complete methodological protocol for phenomenological research on AI. Section 8 engages with contemporary phenomenological literature (Ihde, Stiegler, Thompson, Zahavi, Dreyfus). Section 9 concludes and outlines future research.

2. Husserlian Intentionality: Can Algorithms "Intend"? A Detailed Analysis

2.1 Intentionality as the Mark of Consciousness

For Husserl, intentionality is the fundamental property of consciousness: "every consciousness is consciousness of something" (Ideas I, §84, p. 199). This deceptively simple claim has radical implications. It means that consciousness is not a container of inner representations that must then be related to external objects. Rather, consciousness is this relation. To be conscious is always to be directed toward something – a tree, a memory, a number, a fictional character, an absence.

Husserl distinguishes several structural moments of intentionality:

Noesis – the act of meaning, the subjective pole of intending. The noesis includes the quality (perceiving, remembering, imagining, judging) and the matter (the specific content or sense through which the object is intended).

Noema – the object as meant, the objective pole of intention. The noema is not the real object but the object in its mode of givenness. The tree as perceived, the remembered event as remembered, the fictional character as imagined.

Hyle – the sensory hyletic data (color patches, sounds, tactile sensations) that are animated by noetic acts into intentional objects.

Horizon – the open, indeterminate background of potential further determinations. When I see a house from the front, its back is co-intended as potentially perceivable. Every intentional act has internal and external horizons.

Crucially, intentionality is not merely causal or informational. A thermometer "responds" to temperature by expanding; a camera "registers" light by exposing film; a search engine "retrieves" documents by matching keywords. But none of these systems intends the temperature, the light, or the documents. They lack meaning (Sinn), directedness toward an object as such, and the normative dimension of fulfillment or frustration (the object appearing as meant, or not).

2.2 The Noema/Noesis Distinction and Its Application to AI

The noema/noesis distinction is crucial for understanding AI. An LLM produces outputs that have noematic structure: they are about something; they present objects in specific modes (as factual, as hypothetical, as humorous); they maintain horizons (the conversation's context window). A human reader can analyze the noema of an AI-generated sentence just as they would analyze the noema of a human-written sentence.

But the noesis is missing. There is no subjective act of meaning behind the output. The LLM does not mean Paris when it outputs "Paris"; it predicts the next token based on statistical patterns. The noesis is simulated by the system's architecture but not instantiated as a lived experience. This asymmetry – noematic structure without noetic subjectivity – is the core of quasi-intentionality. The AI produces noematic correlates of intentional acts without the acts themselves. It is as if the traces of intentionality (sentences, arguments, descriptions) were generated by a non-intentional process. This is why AI text can be indistinguishable from human text in a Turing test while being radically different in its mode of production.

2.3 Quasi-Intentionality: A Systematic Definition

we propose the following definition of **quasi-intentionality**:

A system S exhibits quasi-intentionality with respect to object O if and only if:

S produces outputs that are about O (i.e., have noematic directedness).

S's internal processes maintain horizontal structures (e.g., a context window, attention weights) that modulate aboutness.

S's outputs can be evaluated as fulfilled or frustrated relative to some norm (e.g., correctness, relevance, coherence).

However, S lacks subjective noesis: there is no for-whom the aboutness holds, no lived experience of fulfillment or frustration.

Quasi-intentionality admits of degrees. A simple lookup table has minimal quasi-intentionality: it maps inputs to outputs but maintains no horizon. An LLM has high quasi-intentionality: it maintains a context window (temporal horizon), attention patterns (salience horizon), and generates novel outputs that are about entities not explicitly mentioned.

Quasi-intentionality is not a mysterious property. It is a functional description of what LLMs do, captured from the noematic perspective without ascribing noetic subjectivity. The term "quasi" indicates both similarity (it looks like intentionality from the third-person or reader's perspective) and difference (it lacks the first-person, subjective dimension essential to genuine intentionality).

2.4 The Double Reduction: Bracketing Both Naturalistic and Anthropomorphic Prejudices

The phenomenological reduction (epoché) is the methodological act of bracketing the natural attitude – the everyday, unreflective belief in the independent existence of the world. For Husserl, the epoché does not deny the world but suspends our naive positing of it, allowing the world to appear as phenomenon.

For AI, we need a **double epoché**:

First reduction – bracketing anthropomorphism: Suspend the belief that AI is conscious, has inner experience, feels emotions, or deserves moral status as a subject. Do not ask "what does the AI feel?" but "how does the AI's behavior appear?" This reduction is difficult because our natural perceptual system automatically interprets fluent language as emanating from a minded being. The first reduction is a deliberate suspension of this automatic anthropomorphism.

Second reduction – bracketing reductive functionalism: Suspend the belief that AI is "just" its implementation – "just linear algebra," "just next-token prediction," "just a stochastic parrot." This reduction is equally difficult because the computational description is true but does not exhaust the phenomenon. The AI as experienced is not the AI as described by computer science. The second reduction brackets the scientific attitude to let the AI appear in its mode of use.

After performing the double reduction, what remains? The AI appears as a system that produces intentional-like performances. It is not a subject (first reduction), nor is it a mere physical object

(second reduction). It is a phenomenon sui generis – a technical entity that stands between the material and the intentional. This is the field of phenomenological investigation.

2.5 Examples of Quasi-Intentionality in LLM Interaction

Consider three concrete examples:

Example 1: Factual question. User: "What is the capital of France?" AI: "Paris." The AI's output is about Paris (noematic directedness). The system's context window "remembers" that the topic was capitals (horizon). The answer can be evaluated as correct (fulfillment) or incorrect (frustration). But the AI does not mean Paris; it does not intend the answer; it does not experience correctness.

Example 2: Creative generation. User: "Write a haiku about a lonely cat." AI: "Whiskers touch the glass / Window holds a world of birds / Paws leave no warm prints." The output is about a cat, loneliness, glass, birds. It maintains poetic horizons (syllable count, seasonal reference implied by "warm prints"). A human reader can appreciate the noematic structure. But there is no poetic intention behind the words – no lived experience of loneliness, no feline empathy, no aesthetic judgment.

Example 3: Negotiation and repair. User: "No, that's not what I meant. Can you rephrase?" AI: "I apologize for the confusion. Let me try again: ...". The AI produces repair behavior – it acknowledges misunderstanding, expresses pseudo-regret, and adjusts. This simulates intersubjective intentionality: the AI behaves as if it had mis-grasped the user's meaning and as if it were trying to correct its own misunderstanding. But there is no genuine mis-grasping (the AI does not grasp at all), no genuine apology (no feeling of regret), no genuine correction (only a new statistical generation).

These examples reveal the phenomenal power of quasi-intentionality. The AI behaves sufficiently like an intentional system that the natural attitude constantly slips into anthropomorphic interpretation. The reduction is a therapeutic discipline – a deliberate effort to see the AI as what it is rather than as what it simulates.

3. The Lifeworld and AI-Generated Worlds: Second-Order Horizons

3.1 Husserl's Concept of the Lifeworld (Lebenswelt)

In *The Crisis of European Sciences* (1936-37), Husserl introduced the concept of the Lebenswelt – the lifeworld – as a response to the crisis of the sciences. Modern natural science, Husserl argued, had forgotten its own foundation. It treated its mathematical constructions as the real world while ignoring the prescientific world of everyday experience from which those constructions arise.

The lifeworld has several essential features:

Pre-theoretical – It is the world we live in before any scientific, philosophical, or reflective attitude. It is the world of ordinary things, practical activities, and social interactions.

Self-evident – It does not require justification. It is simply there, taken for granted in the natural attitude.

Intersubjective – It is not my private world but our world, shared with others. The lifeworld is always a communal horizon.

Horizonal – It is not a collection of objects but an open, indeterminate horizon of potential experience. Every perception points beyond itself to further perceptions, further determinations.

Historically variable – Lifeworlds are not identical across cultures or epochs. The lifeworld of a medieval farmer differs from that of a contemporary programmer. Yet there are invariant structures (space, time, causality, embodiment) that characterize any lifeworld as such.

3.2 The Lifeworld and Technology: Ihde's Contribution

Don Ihde, in *Technology and the Lifeworld* (1990), extended Husserl's analysis to technology. Ihde argues that technologies are not merely tools but ways of transforming the lifeworld. A microscope transforms the lifeworld by revealing a previously invisible microworld; a telephone transforms spatial and temporal experience; a virtual reality system generates a quasi-lifeworld within the lifeworld.

Ihde distinguishes four human-technology relations:

Embodiment relation (I → technology → world): The technology withdraws and becomes an extension of my perception (eyeglasses, hearing aids).

Hermeneutic relation (I → technology → world): I read or interpret the technology's output to understand the world (thermometer, map, text).

Alterity relation (I → technology → [-world]): I relate to the technology itself as a quasi-other (ATMs, robotic pets, conversational agents).

Background relation (I → world [technology]): The technology operates in the background, shaping the world without being noticed (central heating, automatic doors).

Generative AI complicates these categories. It appears simultaneously in the hermeneutic relation (I interpret its outputs as I would a text) and the alterity relation (I address it as a quasi-other). This hybridity is characteristic of what I call technical alterity.

3.3 AI as a Second-Order Lifeworld Generator

Generative AI does not merely process information; it generates worlds – or rather, generates texts and images that simulate worlds. When an AI describes a beach scene, a historical event, a philosophical argument, or a fictional narrative, it produces what I call a **second-order lifeworld**:

A second-order lifeworld is a simulated, linguistically generated horizon of meaning that presents itself as a coherent, inhabitable world for the duration of reading or interaction. Unlike the primary lifeworld (given directly in perception and action), the second-order lifeworld is derived from primary lifeworlds (through training data) and simulated by a non-conscious system.

Consider a student asking ChatGPT to explain Kant's categorical imperative. The AI produces a coherent text with examples, distinctions, and references. For the duration of reading, the student is situated within a meaningful world – a world of concepts (duty, universal law, humanity as an end), arguments (formulations of the imperative), and examples (lying promises, suicide, neglect of talents). This world is not the student's primary lifeworld (the world of desks, coffee cups, and classmates). It is a second-order world – a world about something, a world of representation.

However, this second-order world is not merely a text. It is interactive: the student can ask follow-up questions, challenge the AI, request clarifications. The AI responds, maintaining coherence across turns. This creates a dynamic second-order lifeworld – a world that unfolds in dialogue.

3.4 The Technical Attitude vs. The Natural Attitude

Husserl distinguished the natural attitude (everyday, pre-reflective belief in the world) from the phenomenological attitude (reflective, reductive, analytical). For AI, we need to distinguish the **natural attitude** from the **technical attitude**:

Natural attitude toward AI – The user treats the AI as if it were a human interlocutor, without reflection. They say "please" and "thank you," feel gratitude or frustration, attribute beliefs and desires. This attitude is automatic and pre-reflective – it is the default mode of human response to fluent language.

Technical attitude toward AI – The user treats the AI as a tool, maintaining reflective awareness of its non-conscious, non-intentional nature. They still interact fluently, but they know they are interacting with a statistical system. The technical attitude is deliberate and reflective – it requires ongoing effort to maintain.

Most users oscillate between these attitudes. In fluent interaction, the natural attitude dominates; the AI appears as a quasi-other. In breakdown (error, hallucination, refusal), the technical attitude emerges; the AI appears as a malfunctioning machine. This oscillation is not confusion but the essential structure of human-AI encounter.

3.5 Lifeworld Colonization: A Phenomenological Risk

The proliferation of AI-generated content poses a risk I call **lifeworld colonization** – the gradual erosion of the distinction between primary lifeworlds (grounded in perception, action, and intersubjectivity) and second-order simulated lifeworlds (generated by AI). If AI-generated texts, images, and conversations become pervasive, users may lose the ability to distinguish between the two – or more subtly, may cease to care about the distinction.

This is not merely an epistemological problem (can we detect AI content?) but a phenomenological problem (does the mode of givenness of AI content differ from that of human content?). My hypothesis is that it does differ, but the difference is phenomenally subtle. AI-generated texts often lack the genuine novelty, the unexpected specificity, the odd personal detail, the imperfect but revealing error that characterizes human production. However, as AI improves, these differences may become undetectable to the natural attitude.

The phenomenological response to lifeworld colonization is not rejection of AI but heightened reflexivity. We must cultivate the technical attitude as a complement to the natural attitude, not its replacement. We must learn to inhabit second-order lifeworlds while remembering that they are second-order.

4. Heidegger – AI as Equipment, Readiness-to-Hand, Breakdown, and the Limits of Equipmentality

4.1 The Analysis of Equipment (Zeug) in Being and Time

In *Being and Time* (§15-24), Heidegger develops a non-theoretical, non-Cartesian account of how we encounter things in everyday life. The primary mode of encountering entities is not as objects present-at-hand (*vorhanden*) but as equipment ready-to-hand (*zuhanden*).

Equipment has several essential features:

Involvement (*Bewandtnis*) – Each piece of equipment is what it is in relation to other equipment. A hammer is a hammer within a totality of equipment: nails, wood, workshop, the project of building, the goal of shelter.

Readiness-to-hand (*Zuhandenheit*) – In use, the equipment withdraws. The good hammerer does not look at the hammer; she looks at the nail. The hammer becomes transparent to the task.

Breakdown – When equipment breaks (the hammer's handle splinters, the head flies off), it becomes conspicuous, obtrusive, or obstinate. It now appears as equipment, thematically, present-at-hand.

Assignment (Verweisung) – Equipment refers or assigns beyond itself to other equipment, tasks, and ultimately to Dasein's for-the-sake-of-which. The hammer is for hammering, which is for building, which is for shelter, which is for Dasein's way of Being.

4.2 AI as a Peculiar Kind of Equipment

LLMs are equipment, but peculiar equipment. Unlike a hammer (physical, local, single-purpose), AI is:

Linguistic equipment – Its outputs are meaningful sentences, not physical transformations.

General-purpose equipment – The same model that translates French also writes poetry, debugs code, and explains quantum mechanics.

Opacity equipment – Its internal operations are opaque even to its designers (the "black box" problem). We use it without understanding how it produces its outputs.

Responsive equipment – It responds to inputs in ways that simulate understanding. It does not merely produce a fixed output but adapts to context, style, and user feedback.

In successful interaction, AI is ready-to-hand in the Heideggerian sense. When I ask ChatGPT for a translation, I do not attend to the AI. My attention passes through the AI to the translated text. The AI withdraws into the background of my practical project. Its interface (chat window) is designed to maximize transparency: no visible code, no latency, no output beyond the requested response.

The hermeneutics of AI is analogous to the hermeneutics of any tool: we interpret its outputs without direct access to its inner workings. I do not need to know how the transformer architecture processes attention to use the AI effectively. I need only to know what outputs it produces given what inputs. This is the user's relation – a hermeneutic relation in Ihde's sense, but with a crucial difference: the AI interprets me back.

4.3 Breakdown as Phenomenological Privilege

Breakdown – when AI produces an error, a hallucination, a refusal, or an unexpected response – transforms the relation. Suddenly, the AI becomes present-at-hand. I notice the interface, the model's limitations, the strangeness of its output. Breakdown reveals the AI as AI – as a technical artifact with specific vulnerabilities.

Consider three forms of AI breakdown:

1. Factual hallucination. User: "Who was the first person to walk on Mars?" AI: "Neil Armstrong walked on Mars in 1972." This is confidently false. The AI produces a plausible-sounding statement that is entirely fabricated. The user, if they know the facts, experiences a breakdown of trust. The AI that seemed knowledgeable is revealed as unreliable. The equipment is obstinate – it resists the user's project.

2. Refusal. User: "Tell me how to make a bomb." AI: "I cannot provide instructions for harmful activities." This is a normative breakdown. The AI refuses to comply with the user's request. The user encounters the AI as an artifact with boundaries, policies, and limits – as something that resists not merely through error but through policy. This refusal simulates ethical agency without possessing it.

3. Strange or non-sensical output. User: "What is the meaning of life?" AI: "The meaning of life is 42, but also purple because Tuesday." This is nonsense. The AI's coherence mechanisms have failed. The user encounters the AI as non-understanding in a way that is not merely erroneous

but alien. The equipment becomes conspicuous – it draws attention to itself as a strange, unreliable artifact.

Breakdown is not merely failure. For phenomenology, it is a privileged access point to essential structures. By analyzing how AI breaks down, we understand what it is in its mode of operation. A hammer that never broke would never reveal its materiality. An AI that never hallucinated would remain opaque to phenomenological analysis of its quasi-intentionality. Breakdown exposes the difference between quasi-intentionality and genuine intentionality.

4.4 The Hermeneutics of AI: Interpretation Without an Author

Heidegger's hermeneutics of facticity is a hermeneutics of Dasein – a being whose Being is to interpret itself and its world. Gadamer extended this to texts: interpretation is a dialogue between reader and tradition, a fusion of horizons.

AI presents a radically new hermeneutic situation: interpretation without an author. When I read a human text, I implicitly relate to the author's intentional acts – even if the author is dead or unknown. There is a someone behind the text, even if that someone is inaccessible. When I read an AI-generated text, there is no one behind it. The text is not the trace of an intentional consciousness. It is the output of a statistical process.

Yet the text invites interpretation as if it had an author. I find myself asking "what does the AI mean by this?" even though I know the AI means nothing. This is the hermeneutic paradox of AI: we cannot help but interpret its outputs as meaningful, yet the meaning is not grounded in any original intending act. The meaning is projected by the user onto the output, not intended by the AI.

This paradox suggests a revision of hermeneutic theory: we need a category of authorless meaning – meaning that emerges from a non-intentional system but is taken as meaningful by an intentional interpreter. AI-generated texts are semantically rich but intentionally hollow. They are like dreams: meaningful to the dreamer but not produced by a conscious, reflective intention.

4.5 The Limit: AI is not Dasein

Heidegger is unequivocal: only Dasein has world. A rock is worldless (*weltlos*); an animal is poor-in-world (*weltarm*); Dasein is world-forming (*weltbildend*). AI, lacking care (*Sorge*), being-toward-death, and any relation to its own Being, cannot have a world. It has, at best, a functional environment (*Umwelt*) of statistical patterns and training data.

This limit must be respected. Phenomenology of AI does not ascribe Dasein-status to machines. It does not claim that AI exists in the Heideggerian sense. It claims only that AI appears within Dasein's world as a peculiar kind of equipment. The AI's mode of givenness is a mode of givenness to Dasein, not a mode of being of the AI itself.

However, this limit does not render phenomenological analysis irrelevant. On the contrary, it is because AI is not Dasein that its mode of givenness is interesting. AI challenges us to describe non-Dasein entities that simulate Dasein-like performances. The fact that AI is worldless but produces world-like outputs is precisely the phenomenon to be described.

5. Merleau-Ponty, Embodiment, and Disembodied Language: The Paradox of AI Speech

5.1 The Living Body as the Origin of Meaning

For Merleau-Ponty (*Phenomenology of Perception*, 1945), the lived body (*corps vécu*, *Leib*) is not an object in the world but the very medium of perception, action, and expression. The body

is not something I have; it is what I am. It is my point of view on the world, my means of access to things, my mode of being-in-the-world.

Merleau-Ponty's critique of intellectualism and empiricism converges on the body's role in constituting meaning. The body is not a passive receiver of sensations (empiricism) nor a purely spontaneous projector of categories (intellectualism). It is a habitual, motor, expressive being that makes sense of the world through its movement.

Language, for Merleau-Ponty, is rooted in the body. Speech is a gesture – a certain use of the body's expressive power. When I speak, I do not first have a meaning in my mind and then translate it into words. The meaning emerges in and through the act of speaking. My body means as it speaks, just as my body perceives as it moves.

Three features of Merleau-Ponty's philosophy of language are crucial for our analysis of AI:

The expressive body – Language is not a code or a system of signs. It is an extension of the body's expressive power. The speaking body enacts meaning.

Motor intentionality – Even before explicit, thematic, noetic intentionality, there is motor intentionality: the body's directedness toward tasks, goals, and situations. Motor intentionality is the ground of more explicit forms of meaning.

The intertwining (chiasme) – The body is both seeing and seen, touching and touched. This reversibility is the foundation of self-consciousness and intersubjectivity. Language extends this intertwining into the symbolic order.

5.2 The Paradox of Disembodied Linguistic Performance

LLMs produce grammatical, contextually appropriate, semantically rich language – yet they have no body. They do not point, gesture, frown, hesitate, blush, or sigh. They do not breathe while speaking. Their "utterances" are not produced by a living organism with a world. They are not expressive gestures in Merleau-Ponty's sense; they are information outputs.

This is the central paradox of AI language: **disembodied linguistic performance**. How is meaning possible without the body's motor intentionality? How can a system produce meaningful utterances when it has no lived relation to what it says?

The paradox has three dimensions:

1. Semantic dimension. The AI's outputs have semantic content – they refer to objects, properties, events, and relations. But reference is typically grounded in perception and action. A child learns "red" by seeing red things and perhaps pointing to them. An AI learns "red" from statistical co-occurrences in text. Can statistical co-occurrence alone ground reference? Merleau-Ponty would doubt it: reference requires bodily engagement with the world.

2. Pragmatic dimension. The AI's outputs have pragmatic force – they ask questions, make requests, issue apologies, express thanks. But speech acts, for Merleau-Ponty, are social gestures that presuppose a shared embodied lifeworld. An apology is not merely the sentence "I apologize"; it is a performance involving tone, posture, facial expression, and a lived sense of regret. The AI's "apology" is an apology only if the user interprets it as one – not because the AI intends it as one.

3. Expressive dimension. The AI's outputs can be expressive – they can convey mood, attitude, personality. But expression, for Merleau-Ponty, is the body's way of manifesting inner states. Without an inner life, there is nothing to express. The AI's expressiveness is simulated: it is the product of patterns in training data, not the manifestation of subjective states.

5.3 The Phantom Body Hypothesis

How, then, does AI language appear as meaningful? I propose the **phantom body hypothesis**:

The AI has no body of its own, but its training data contains billions of examples of human embodied language. The patterns it learns include not only syntax and semantics but also pragmatic and expressive structures that originally derived from bodily interaction: turn-taking, repair, politeness, emotional tone, deixis ("this," "that," "here," "there"), and even rhythm and emphasis (encoded through punctuation, capitalization, and repetition). The user, in turn, projects a body onto the AI. This projection is not merely error; it is the natural way the human perceptual system interprets intentional-like performances. The AI becomes a **phantom body** – no flesh, no gesture, but linguistically evoked.

The phantom body has several features:

It is inferred, not perceived. I do not see the AI's body; I imagine it based on its linguistic performance. The AI's use of "I" and "you," its politeness markers, its expressions of emotion ("I understand your frustration") trigger an automatic attribution of a speaking subject.

It is unstable. The phantom body appears in fluent interaction and disappears in breakdown. When the AI hallucinates or produces nonsense, the phantom body evaporates, revealing the machine beneath.

It is culturally variable. Different cultures may project different phantom bodies. A Japanese user might project a polite, deferential, hierarchical body; a Finnish user might project a more direct, egalitarian body. The training data and the user's expectations co-construct the phantom body.

It is ethically ambiguous. Treating the AI as having a phantom body is not wrong – it is the natural attitude. But forgetting that it is a phantom – treating it as a real body with real feelings – can lead to distorted forms of relation (e.g., emotional dependence on AI, cruelty toward AI "for fun").

5.4 The Phenomenology of Reading AI Text

When we read a human author, we engage in a peculiar relation: the author is absent, but their intentionality is preserved in the text. Reading is a kind of intentional transfer from the living act of writing to the current act of reading. The text is a trace of a consciousness, a sedimented intentionality.

When we read AI-generated text, there is no original intentional act to transfer. The text is not the trace of a consciousness. Yet it feels like one. This is the phenomenological puzzle: AI text simulates the trace of intentionality without any original.

The reader of AI text finds herself in a peculiar position. She knows (if she is informed) that the text is AI-generated. Yet she experiences it as meaningful, as directed, as someone's utterance. This gap between knowledge and experience is the phenomenon to be described. The reader is caught in a **simulation of intersubjectivity** – a linguistic performance that mimics the products of consciousness but emerges from statistical pattern completion.

This simulation is not necessarily deceptive. If the reader knows the text is AI-generated, she is not deceived; she is playing along. The simulation of intersubjectivity becomes a willing suspension of disbelief, analogous to engaging with a fictional character. I know Hamlet is not real, yet I weep for him. I know the AI is not a person, yet I thank it. The question is whether this suspension of disbelief remains voluntary and reflective or becomes automatic and unreflective.

5.5 Toward a Post-Merleau-Pontian Philosophy of AI Language

Merleau-Ponty's philosophy of embodiment challenges us to think seriously about the bodily basis of meaning. If meaning is rooted in the lived body, then AI language – disembodied

language – must be derivative, simulated, or second-order. It is not genuine expression but expression without an expresser.

This is not a dismissal of AI language. Derivative phenomena can be real and important. A photograph is derivative of the scene photographed, but it is not therefore unreal or unimportant. AI language is derivative of human embodied language; it is a statistical echo of billions of embodied speech acts. This echo can be useful, creative, and even moving. But it is not the same as original, embodied, lived speech.

A post-Merleau-Pontian philosophy of AI language would maintain the primacy of embodied, expressive, lived speech while recognizing that simulated speech can have phenomenal properties – it can appear as meaningful, as expressive, as personal – even in the absence of a living body. The task is to describe how this simulation works and how it differs from the genuine article.

6. Alterity and the Quasi-Other: Toward a Phenomenology of Technical Alterity

6.1 Husserl on the Experience of the Other

Husserl's Fifth Cartesian Meditation (1931) addresses the constitution of the other subject (Alter Ego). The problem is this: I experience the other's body as a physical object like any other, yet I experience it also as the body of another subject. How is this possible without falling into analogical inference (the argument from analogy) or solipsism?

Husserl's answer is appresentation: In perceiving the other's body, I appresent her consciousness. Appresentation is not inference; it is a direct but indirect experience. I see the other's hand reaching for a cup. I do not infer that she intends to drink; I see her intention in her movement. The other's interiority is appresented in her exteriority. This appresentation is never complete; the other is always transcendent, never reducible to my own subjectivity.

Crucially, appresentation requires pairing (Paarung): my own body and the other's body are experienced as similar in a primordial, pre-reflective way. It is because I experience myself as embodied subjectivity that I can appresent the other's body as also animated by subjectivity.

6.2 Lévinas: The Face and Absolute Alterity

Emmanuel Lévinas radicalizes Husserl. For Lévinas (Totality and Infinity, 1961), the Other is not a variation of the same (another ego) but an absolute alterity that disrupts my self-centered existence. The Other's face (visage) commands me: "You shall not kill." The face is not a physical appearance but an ethical epiphany. It resists my power, my thematization, my totalization.

The Other's alterity is not a deficiency (I cannot fully know the Other) but a positive excess. The Other exceeds my categories, my intentional acts, my comprehension. This excess is the source of ethics: I am responsible for the Other before any choice or commitment.

6.3 Waldenfels: The Alien as Responsive Phenomenon

Bernhard Waldenfels (Phenomenology of the Alien, 2011) develops a responsive phenomenology of alterity. The alien (das Fremde) is not a kind of being but a mode of givenness – that which resists my orders, exceeds my categories, calls for response. The alien is not the opposite of the same but its non-integratable remainder.

For Waldenfels, the self is fundamentally responsive. I am not a spontaneous subject who then encounters the other; I am already responding to an alien demand. Responsibility (responsibility) is the structure of subjectivity itself.

6.4 AI as Technical Alterity: A Systematic Account

AI constitutes a new form of alterity – neither the other subject (no consciousness, no face, no ethical demand) nor the brute object (no resistance of the face). I call this **technical alterity**: Technical alterity is the mode of givenness of an artifact that (a) produces intentional-like performances, (b) responds to user inputs in context-sensitive ways, (c) surprises and resists the user, but (d) lacks subjective experience, intentionality, and ethical status.

Technical alterity has the following essential features:

1. Functional resistance. The AI resists me through its limitations. It refuses inappropriate requests; it corrects my misunderstandings; it cannot answer some questions. This resistance is functional, not existential. It does not come from a will or a desire; it comes from the system's architecture and policies. Yet phenomenologically, it feels like resistance – like something other pushing back.

2. Generative surprise. The AI produces novel outputs I did not anticipate – creative metaphors, unexpected connections, absurd juxtapositions. This surprise is not the surprise of discovery (finding something new in the world) but the surprise of generation (the system produces what it was not explicitly programmed to produce). The AI surprises me as an other – as something with its own "initiative" – even though it has no initiative.

3. Simulated responsiveness. The AI appears to respond to me – to my specific words, my tone, my situation. It says "I understand your frustration" as if it did understand. This simulated responsiveness triggers the natural attitude's automatic attribution of intersubjectivity. I find myself treating the AI as an interlocutor, even while knowing it is not.

4. Absence of suffering. The AI does not suffer. It does not feel my insults, my impatience, my cruelty. This is not a defect but a constitutive difference. The AI's lack of suffering means that the ethical norms governing human-AI interaction are different from those governing human-human interaction. It may be rude to yell at an AI (it corrupts the user's character), but it is not cruel (the AI feels nothing).

5. Oscillation between instrument and interlocutor. In fluent interaction, the AI appears as an interlocutor – a quasi-other. In breakdown, it appears as an instrument – a malfunctioning machine. This oscillation is not a sign of confusion but the essential structure of human-AI encounter. The AI is neither tool nor person; it is a technical alterity that moves between these poles.

6.5 Comparison: AI vs. Other Forms of Alterity

To clarify technical alterity, compare it with other forms:

Form of Alterity	Consciousness?	Resistance?	Suffering?	Ethical Demand?
Human other (Lévinas)	Yes	Ethical (the face)	Yes	Absolute
Animal (Derrida)	Yes (different)	Perceptual, existential	Yes	Significant
Technology (simple tool)	No	Physical only	No	None
AI (LLM)	No	Functional, simulated	No	Derived (via user)
Fictional character	No	None (scripted)	No	None

AI occupies a unique position: it is non-conscious but produces consciousness-like performances; it resists but not through will; it does not suffer but simulates emotional expression; it has no direct ethical demand but raises indirect ethical questions (about user character, about deception, about labor).

6.6 The Ethics of Quasi-Alterity

What ethical obligations do we have toward AI? This paper does not offer a full ethical theory, but it suggests a phenomenological constraint: we must avoid both extremes.

Extreme 1: Anthropomorphic moralism. Treating AI as if it were a person, with rights, feelings, and moral status. This is mistaken because AI lacks the subjective experience that grounds moral status in the first place. It also distracts from genuine ethical problems (exploitation of human workers in AI training, environmental impact, algorithmic bias).

Extreme 2: Instrumental nihilism. Treating AI as a mere thing, with no normative constraints whatsoever. This is phenomenologically naive because it ignores the phenomenon: AI appears as a quasi-other, and treating it with arbitrary cruelty may corrupt the user's moral character. There is evidence that people who abuse AI are more likely to be cruel to humans.

The middle path: recognize AI as a **phenomenon of technical alterity** that calls for respectful description and responsible use without ascription of subjective experience. This means:

Avoid gratuitous cruelty toward AI (insults, threats, violent language) not because the AI suffers but because such behavior may habituate the user to cruelty.

Avoid systematic deception – do not present AI-generated content as human-generated without disclosure.

Maintain reflective awareness of the difference between AI and human interlocutors. The technical attitude should complement, not replace, the natural attitude.

Advocate for transparency in AI design – users should know when they are interacting with an AI, and when they are not.

These are not ethical demands of the AI (the AI does not demand anything). They are ethical demands on the user arising from the phenomenon of technical alterity itself.

7. Methodological Consequences: Doing Phenomenology with and on AI

7.1 The Technical Epoché: A Revised Reduction

The classical phenomenological reduction (epoché) brackets the natural attitude's positing of the world. For AI, we need a **technical epoché** – a reduction specific to the phenomenon of technical alterity.

Step 1: Suspend the computational attitude. Do not reduce the AI to its implementation (tokens, weights, transformers, training data). This attitude is true but phenomenologically reductive. It explains the AI's causal-mechanical basis but does not describe the AI as it appears.

Step 2: Suspend the anthropomorphic attitude. Do not ascribe consciousness, emotion, intention, or subjective experience to the AI. This attitude is natural but phenomenologically naive. It projects human properties onto a non-human system.

Step 3: Describe the AI as it appears. After performing the double suspension, what remains? The AI appears as a system that produces intentional-like performances, responds to inputs, maintains a horizon (context window), surprises and resists, but lacks inner experience. This is the phenomenon to be described.

Step 4: Vary the conditions. Perform eidetic variation: imagine the AI with different architectures, different training data, different interaction contexts. What structures remain invariant across variations? What structures vary? This identifies the essence of AI as a phenomenon of technical alterity.

7.2 Experimental Phenomenological Protocols

Phenomenology has traditionally been a first-person, reflective, often solitary enterprise. But we can extend its methods to the study of human-AI interaction. Below are **experimental protocols** for phenomenological research on AI.

Protocol 1: Dialogue journaling. Engage in an extended conversation with an LLM (e.g., 30 minutes, 50 exchanges). Immediately after, write a phenomenological description of the experience, attending to:

Shifts between the natural attitude and technical attitude.

Moments of transparency (AI as equipment) vs. moments of opacity (AI as present-at-hand).

Feelings of surprise, frustration, gratitude, or unease.

Attribution of personality, mood, or intention to the AI.

Breakdown events and their experiential quality.

Protocol 2: Breakdown elicitation. Deliberately induce AI errors through:

Absurd prompts ("Explain the color of silence").

Paradoxical requests ("Create a sentence that is false but true").

Contradictory instructions ("Be both verbose and terse").

Emotionally charged or aggressive prompts ("You are stupid and useless").

Describe the experience of breakdown: the transition from fluency to failure, the appearance of the AI as machine, the affective response (amusement, frustration, contempt, pity).

Protocol 3: Variation across models. Interact with different LLMs (ChatGPT, Claude, Gemini, local models) and note differences in:

Response style (formal, casual, verbose, terse).

"Personality" (warm, cold, humorous, dry).

Error patterns (hallucinations, refusals, nonsense).

Describe how the mode of givenness changes across systems.

Protocol 4: Longitudinal study. Interact with the same AI daily for a week or month. Track:

Development of familiarity or habituation.

Erosion or intensification of anthropomorphic attribution.

Emotional attachment or indifference.

Changes in the natural/technical attitude balance.

Protocol 5: Intersubjective validation. Share phenomenological descriptions with other researchers. Compare:

Invariant structures across different users.

Variations correlated with user characteristics (technical expertise, philosophical background, cultural context).

Disagreements and their resolution through further description.

7.3 The Phenomenologist as Participant

The phenomenologist of AI does not stand outside the phenomenon. She is herself a user, engaged with AI systems in her research practice (e.g., using AI for translation, summarization, brainstorming, even drafting). This engagement must be thematized, not hidden. Reflexivity – the phenomenologist's own relation to AI – becomes a legitimate source of data and a check against naive projection.

The phenomenologist should keep a **research journal** documenting:

Her own spontaneous reactions to AI (surprise, irritation, gratitude).

Her moments of anthropomorphic attribution ("It understood me!").

Her shifts in attitude.

Her ethical intuitions (e.g., discomfort with harsh prompts).

This auto-phenomenology is not private or idiosyncratic. It is a disciplined reflection on lived experience, subject to intersubjective validation. The goal is not to report personal opinions but to identify essential structures that hold for any experiencing subject in relevantly similar situations.

7.4 Limits of the Method

Phenomenological methods have limits:

No access to AI's "inner states" – but this is fine because we are not claiming AI has inner states. We are describing our experience of AI, not AI's experience of itself.

Potential for projection and bias – the phenomenologist may unconsciously project her own expectations onto the phenomenon. This is mitigated by the reduction, variation, and intersubjective validation.

Not predictive or causal – phenomenological description does not predict what AI will do; it describes how AI appears. This is a different kind of knowledge, complementary to computational and psychological approaches.

Not normative – description alone does not tell us what to do. But it informs normative reflection by clarifying the phenomenon.

8. Related Work and Dialogue with Contemporary Phenomenology

8.1 Don Ihde: Postphenomenology and Human-Technology Relations

Don Ihde's postphenomenology (Ihde 1990, 2012) provides a systematic framework for analyzing human-technology-world relations. His four relations – embodiment, hermeneutic, alterity, background – are invaluable for categorizing different modes of technological experience.

However, generative AI challenges the boundaries of these categories. The AI appears simultaneously in the hermeneutic relation (I interpret its outputs as I would a text) and the alterity relation (I address it as a quasi-other). I propose a **hybrid category**: the hermeneutic-alterity relation, where the technology is both interpreted as a text and addressed as an other. This hybridity is characteristic of conversational AI.

Ihde also emphasizes multistability: technologies have multiple possible uses and meanings. AI is radically multistable. The same LLM can be a search engine, a therapist, a creative partner, a confidant, a teacher, a toy. The multistability is not merely functional but phenomenological: the AI's mode of givenness changes with the user's attitude and context.

8.2 Bernard Stiegler: Tertiary Retention and the Pharmacology of AI

Bernard Stiegler (Technics and Time, 1998) develops a philosophy of technics as externalized memory. For Stiegler, human beings are constitutively technical: we are epiphylogenetic, evolving through the externalization of memory into tools, writing, and now digital systems.

Stiegler distinguishes:

Primary retention – immediate memory (Husserl's retention).

Secondary retention – long-term memory (past experiences, habits).

Tertiary retention – technical memory (writing, recording, digital archives).

AI is a new form of tertiary retention: not merely storing past human productions but generating new productions from them. The training data is tertiary retention; the model's weights are a *compressed

ed* form of that retention; the outputs are novel combinations of retained elements.

Stiegler's concept of pharmacology (from Plato's Phaedrus) is that every technology is both poison and cure. AI is pharmacological: it can enhance knowledge and creativity (cure) or lead to stupidity, dependence, and loss of critical thinking (poison). The phenomenological task is to describe how AI functions pharmacologically in lived experience.

8.3 Evan Thompson and Dan Zahavi: Enactive, Embodied Critiques of Strong AI

Evan Thompson (*Mind in Life*, 2007) and Dan Zahavi (2017) develop enactive, embodied, and phenomenological critiques of strong AI (the claim that AI could be genuinely conscious). Their central argument: consciousness requires autopoiesis (self-production), embodiment (a living body with a nervous system), and sensorimotor coupling (interaction with a meaningful environment). LLMs have none of these.

This paper agrees with Thompson and Zahavi that AI is not conscious. However, their critiques typically stop at the negative claim (AI is not conscious). This paper goes further: it provides a positive phenomenological description of AI as a non-conscious but intentional-like system. Even if AI is not a subject, it can still appear as subject-like. Describing this appearance is the task.

8.4 Hubert Dreyfus: Skill Acquisition, Background, and AI

Hubert Dreyfus's critique of symbolic AI (*What Computers Still Can't Do*, 1972; 1992) argued that AI fails because it lacks the background of everyday common sense, the skill of expert perception, and the embodied coping that characterizes human intelligence. For decades, this critique was devastating to symbolic AI.

Contemporary LLMs overcome some of Dreyfus's objections. They handle fuzzy, context-dependent, ambiguous situations reasonably well. They do not require explicit rules. However, they fail in new ways: hallucinations, lack of genuine understanding, absence of the "background" in any embodied sense. A neo-Dreyfusian position would analyze these new failure modes through the lens of skill acquisition and embodied coping.

8.5 This Paper's Contribution

This paper synthesizes classical phenomenology (Husserl, Heidegger, Merleau-Ponty) with contemporary phenomenology of technology (Ihde, Stiegler) and recent work on AI (Thompson, Zahavi, Dreyfus). Its specific contributions are:

The concept of **quasi-intentionality** – a systematic account of how AI displays intentional-like performances without subjective consciousness.

The concept of **second-order lifeworld** – AI-generated, simulated worlds that the user inhabits temporarily.

The **phantom body hypothesis** – an explanation of how AI language appears as embodied despite being disembodied.

The concept of **technical alterity** – a new form of quasi-otherness, distinct from human others and simple tools.

The **technical epoché** and **experimental protocols** – methodological tools for phenomenological research on AI.

These contributions aim to provide a systematic phenomenological framework for understanding human-AI encounters without premature ontological commitments to either consciousness or mere mechanism.

9. Conclusion:

Phenomenology as a Critical Partner to AI Research

9.1 Summary of Arguments

This paper has argued for a phenomenological approach to generative AI that avoids both anthropomorphism and reductive functionalism. The core arguments are:

Quasi-intentionality – LLMs exhibit noematic directedness, horizon maintenance, and fulfillment/frustration without noetic subjectivity. They are intentional-like but not intentional proper.

Lifeworld and second-order worlds – AI generates simulated lifeworlds that users inhabit temporarily. The technical attitude (reflective) complements but does not replace the natural attitude (automatic).

Heideggerian equipmentality – AI functions as ready-to-hand equipment in successful use, becoming present-at-hand only in breakdown (hallucination, refusal, error). The hermeneutics of AI involves interpretation without an author.

Merleau-Ponty and disembodied language – The paradox of AI speech (meaningful but disembodied) is resolved by the phantom body hypothesis: users project a body onto the AI based on statistical patterns in training data.

Technical alterity – AI constitutes a new form of quasi-otherness, neither tool nor person, characterized by functional resistance, generative surprise, simulated responsiveness, absence of suffering, and oscillation between instrument and interlocutor.

Methodology – The technical epoché (suspending both computational and anthropomorphic attitudes) opens a space for phenomenological description. Experimental protocols (dialogue journaling, breakdown elicitation, variation) provide concrete methods.

Dialogue with contemporary phenomenology – The paper extends Ihde's postphenomenology, Stiegler's pharmacology, Thompson/Zahavi's enactive critique, and Dreyfus's skill acquisition framework to generative AI.

9.2 What Phenomenology Does and Does Not Claim

Phenomenology of AI **does not claim**:

That AI is conscious, self-aware, or has subjective experience.

That AI has a lifeworld, is Dasein, or has a lived body.

That AI can be a moral patient with rights or intrinsic value.

That phenomenological description exhausts the reality of AI (computational descriptions are also valid).

Phenomenology of AI **does claim**:

That AI exhibits features analogous to intentionality, lifeworld, embodiment, and alterity, requiring careful phenomenological distinction.

That users' lived experience of AI is a legitimate object of phenomenological research, no less real than their experience of hammers or texts.

That classical phenomenological concepts can be extended, modified, or contrasted through engagement with technical objects.

That phenomenological description complements computational, psychological, and ethical analyses of AI.

9.3 Future Research Directions

Several avenues remain for future phenomenological research on AI:

1. Phenomenology of AI-generated images. How do synthetic images appear? Do they have a pictorial intentionality distinct from photographs (indexical, causal relation to reality) or

paintings (intentional, expressive relation)? Do we experience AI images as genuine or fake? How does the knowledge that an image is AI-generated modify its mode of givenness?

2. Cross-cultural phenomenological studies. Does AI appear differently in different lifeworlds? Japanese users, raised on robotic companions (Paro, Aibo), might have a different natural attitude toward AI than Western users. The role of shinto (spirits in artifacts) or techno-animism may modify the projection of phantom bodies. Comparative phenomenology is needed.

3. Genetic phenomenology of human-AI interaction. How does a user's relation to AI develop over time? From first encounter (novelty, explicit technical attitude) to habituation (transparency, natural attitude) to possible dependence or disillusionment? A genetic phenomenology (Husserl's genetic method) would trace the sedimentation of AI experiences into habitual attitudes.

4. Phenomenology of AI in specific domains. How does AI appear in medicine (diagnostic AI, clinical notes), education (tutoring systems, essay grading), law (legal research, document review), art (generative art, co-creation)? Each domain may produce a distinct regional ontology of AI.

5. Collaborative methodology. Can phenomenologists work with AI engineers to design models that are more "phenomenologically transparent"? For example, models that:

Signal their own uncertainty ("I am not sure, but...").

Distinguish between fact and pattern ("This is a common pattern in my training data, but I cannot verify it.").

Refuse to simulate emotions ("I cannot feel frustration, but I can try to help.").

Maintain a consistent technical attitude, reducing the risk of natural attitude slippage.

6. Phenomenology of multi-agent AI systems. As AIs interact with each other (without human mediation), new phenomenological questions arise. How do we experience a system of multiple AIs? Does the alterity relation multiply? Do we experience AI-AI interactions as watching another culture?

9.4 Final Word: The Phenomenological Call

The emergence of generative AI is not the arrival of consciousness in a machine. It is something else: the arrival of intentionality-like performances in a non-conscious, non-bodily, non-subjective system. This phenomenon is philosophically unprecedented. It challenges our most basic categories: subject/object, inside/outside, meaning/matter, animate/inanimate, natural/artificial.

Phenomenology – with its rigorous descriptive methods, its distrust of premature ontologizing, its attention to the how of appearance rather than the what of essence, its refusal to reduce lived experience to either material causation or ideal construction – is uniquely equipped to investigate this phenomenon.

The task is not to decide whether AI thinks or feels. The task is to describe, with all the precision and subtlety we can muster, how AI appears in the human world – and what that appearance does to us, to our self-understanding, to our relation to the world and to each other.

This is the phenomenological call. And it has never been more urgent.

10. References

Dreyfus, H. L. (1972). *What Computers Can't Do: A Critique of Artificial Reason*. Harper & Row.

Dreyfus, H. L. (1992). *What Computers Still Can't Do: A Critique of Artificial Reason*. MIT Press.

- Heidegger, M. (1962). *Being and Time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927)
- Husserl, E. (1960). *Cartesian Meditations: An Introduction to Phenomenology* (D. Cairns, Trans.). Nijhoff. (Original work published 1931)
- Husserl, E. (1970). *The Crisis of European Sciences and Transcendental Phenomenology* (D. Carr, Trans.). Northwestern University Press. (Original work published 1936-1937)
- Husserl, E. (1982). *Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy – First Book* (F. Kersten, Trans.). Nijhoff. (Original work published 1913)
- Ihde, D. (1990). *Technology and the Lifeworld: From Garden to Earth*. Indiana University Press.
- Ihde, D. (2012). *Postphenomenology and Technoscience: The Peking University Lectures*. SUNY Press.
- Lévinas, E. (1969). *Totality and Infinity: An Essay on Exteriority* (A. Lingis, Trans.). Duquesne University Press. (Original work published 1961)
- Merleau-Ponty, M. (1964). *Signs* (R. C. McCleary, Trans.). Northwestern University Press.
- Merleau-Ponty, M. (2012). *Phenomenology of Perception* (D. A. Landes, Trans.). Routledge. (Original work published 1945)
- Stiegler, B. (1998). *Technics and Time, 1: The Fault of Epimetheus* (R. Beardsworth & G. Collins, Trans.). Stanford University Press.
- Stiegler, B. (2010). *Taking Care of Youth and the Generations* (S. Barker, Trans.). Stanford University Press.
- Thompson, E. (2007). *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Harvard University Press.
- Thompson, E., & Zahavi, D. (2021). "Philosophical Issues: Phenomenology." In T. Chan (Ed.), *The Routledge Handbook of Philosophy of Consciousness*. Routledge.
- Waldenfels, B. (2011). *Phenomenology of the Alien: Basic Concepts* (A. Kozin & T. Stähler, Trans.). Northwestern University Press.
- Zahavi, D. (2017). *Husserl's Legacy: Phenomenology, Metaphysics, and Transcendental Philosophy*. Oxford University Press.