ISSN 1983-9979
prolíngua

ARTIFICIAL INTELLIGENCE AND THE PROBLEM OF DISEMBODIED COGNITION: A PHENOMENOLOGICAL CRITIQUE OF THE LIMITS OF THE ARTIFICIAL MIND

A INTELIGÊNCIA ARTIFICIAL E O PROBLEMA DA COGNIÇÃO DESENCARNADA: UMA CRÍTICA FENOMENOLÓGICA AOS LIMITES DA MENTE ARTIFICIAL*

Tiago Negrão Andrade[1]

ABSTRACT: The emergence of artificial intelligence (AI) as a technical and epistemic phenomenon has reshaped philosophical debates concerning the nature of mind, cognition, and subjectivity. This article examines the ontological assumptions underlying the equation between human and artificial cognition, questioning the possibility of a disembodied mind. Through a phenomenological and critical approach, we draw upon the contributions of Merleau-Ponty, Dreyfus, Searle, Chalmers, and contemporary post-LLM authors to demonstrate that cognition cannot be reduced to computational processes, but emerges from lived and situated corporeality. We argue that AI, by abstracting the embodied dimension of experience, operates as a simulacrum incapable of reproducing the intentionality and qualia of human consciousness. The study reveals that attempts to replicate the mind in machinic substrates expose not only the limits of contemporary technical architectures, but also the aporias of traditional philosophical models of cognition. We conclude that AI, more than a technical challenge, constitutes a critical mirror that compels philosophy to rethink the relation between body, mind, and technicity.

KEYWORDS: Artificial intelligence. Embodied cognition. Phenomenology. Philosophy of mind. Technical subjectivity

RESUMO: A emergência da inteligência artificial (IA) como fenômeno técnico e epistêmico reconfigura os debates filosóficos sobre a natureza da mente, da cognição e da subjetividade. Este artigo examina os pressupostos ontológicos que sustentam a equiparação entre cognição humana e artificial, interrogando a possibilidade de uma mente desencarnada. Por meio de uma abordagem fenomenológica e crítica, articulamos as contribuições de Merleau-Ponty, Dreyfus, Searle, Chalmers e outros para demonstrar que a cognição não é redutível a processos computacionais, mas emerge de uma corporeidade vivida e situada. Argumentamos que a IA, ao abstrair a dimensão encarnada da experiência, opera como simulacro, incapaz de reproduzir a intencionalidade e os qualia da consciência humana. O estudo revela que a tentativa de replicar a mente em substratos maquínicos expõe não apenas os limites da técnica, mas também as aporias dos modelos filosóficos tradicionais da cognição. Concluímos que a IA, mais do que um desafio técnico, é um espelho crítico que obriga a filosofia a repensar a relação entre corpo, mente e técnica.

PALAVRAS-CHAVE: Inteligência artificial. Cognição encarnada. Fenomenologia. Filosofia da mente. Subjetividade técnica

1 Introduction

The recent emergence of generative artificial intelligence has profoundly transformed philosophical debates on mind, cognition, and subjectivity. Far from representing a merely technical advance, AI has become an ontological phenomenon that reopens classical questions in the philosophy of mind under radically new material conditions. Large-scale generative architectures—such as GPT, Gemini, and LLaMA—shifted the locus of artificial intelligence from traditional symbolic systems to connectionist models capable of producing language,

* Submmitted: 04.06.2025; Accepted: 10.12.2025 - DOI: 10.22478/ufpb.1983-9979.2025v20n2.74352

[1] Doutor em Mídia e Tecnologia. Universidade Estadual Paulista Júlio de Mesquita Filho (UNESP) Contato: tiago.negrao@unesp.br | Orcid: https://orcid.org/0000-0002-5352-9798

images, and decisions that simulate human-like performances. This shift demands a critical reinterpretation of notions such as intentionality, meaning, embodiment, and consciousness, which can no longer be examined exclusively through conceptual frameworks formulated prior to 2022.

From Cartesian dualism (DESCARTES, 2003) to computational functionalism (PUTNAM, 1997; FODOR, 1981), the philosophy of mind has long sought to delimit the criteria of rationality and consciousness, frequently reducing the mental to the manipulation of formal representations. Phenomenological and enactivist traditions (MERLEAU-PONTY, 1994; DREYFUS, 2007; GALLAGHER, 2005; DI PAOLO & THOMPSON, 2017), however, insist that cognition is not an abstract operation over symbols but an embodied mode of being-in-the-world. Generative AI, which operates without body, affectivity, or lived world, intensifies the tension between disembodied models of cognition and the phenomenological conception of experience as grounded in corporeality.

Decades of debate among functionalists and phenomenologists—Searle (1990), Dennett (1991), Chalmers (1996)—had already sketched the contours between simulation and consciousness. Yet the advent of generative AI reconfigures these contours by producing systems capable of imitating language and performing rationality without any experiential constitution. This epistemic transformation has led contemporary authors to revisit the ontological status of artificial minds: Metzinger (2021) warns that generative models lack any phenomenal self-model; Hayles (2021) argues that posthuman subjectivity is inseparable from embodied materialities; Vallor (2022) demonstrates how algorithmic architectures shape attention, judgment, and virtue; and Roden (2015) contends that even highly sophisticated nonhuman minds remain ontologically discontinuous from embodied experience.

These perspectives show that contemporary debates can no longer revolve around the binary question of whether AI "thinks" or not; rather, they must interrogate the ontological conditions that enable, prevent, or simulate cognition. The central issue is not merely whether artificial models can learn or predict, but whether they can mean, feel, or inhabit a world. Phenomenology reveals that cognition is inseparable from motricity, affective vulnerability, and lived temporality—dimensions absent from any technical architecture currently in existence. Generative AI therefore exposes, with renewed urgency, the fragility of computational models that conceive intelligence as abstract calculation.

Within this context, the present article investigates the ontological assumptions sustaining attempts to equate human and artificial cognition. The main guiding question may be formulated as follows: Is a disembodied mind conceivable, and if not, what structural limits prevent AI from replicating human cognition? To answer this, we articulate a theoretical constellation encompassing both classical traditions—Descartes, Ryle, Putnam, Fodor, Searle, Dennett, Chalmers—and contemporary critical approaches—Nancy, Metzinger, Hayles, Vallor, Roden—in order to map tensions, ruptures, and conceptual possibilities.

Throughout the analysis, we demonstrate that generative AI functions less as a replacement for the human mind than as a critical mirror that reveals the limitations of our philosophical models. By contrast, it highlights that human cognition emerges from situated corporeality, lived temporality, and affective engagement with the world—elements not merely contingent but constitutive. AI, by operating outside this horizon, reinforces the phenomenological thesis that there is no thinking without body, without world, without existential inscription

.

2 Methodology

The investigation conducted in this article adopts a theoretical-critical approach anchored in phenomenology, philosophy of mind, science and technology studies (STS), and the ontological critique of technicity. Given our central objective—to examine the conceptual and ontological limits of artificial cognition, especially in light of post-2022 generative AI—the methodological design privileges hermeneutic analysis, theoretical reconstruction, and critical comparison of philosophical models of mind.

The analytical corpus comprises three principal sets of sources: a) **Classical philosophical works** on cognition, consciousness, and embodiment (Descartes, Ryle, Putnam, Fodor, Merleau-Ponty, Searle, Dennett, Chalmers); b) **Contemporary literature** on embodied cognition, phenomenology, and enactivism (Gallagher, Dreyfus, Di Paolo & Thompson); c) **Post-generative AI contributions**, essential for situating the discussion within the current state of the art (Metzinger, Hayles, Roden, Vallor), all of which problematize the radical disembodiment of generative models and its implications for the epistemology of mind.

Source selection was guided by conceptual relevance, theoretical density, and capacity to illuminate the central problem: the possibility—or impossibility—of disembodied cognition. Works treating AI exclusively as applied technology, rather than addressing its ontological presuppositions, were excluded.

A stratified hermeneutic procedure guided the analysis along three axes: a) **Conceptual archaeology**, reconstructing ontological presuppositions behind symbolic, connectionist, and generative AI models; b) **Comparative tension**, contrasting representationalist theories with phenomenological and enactivist perspectives; c) **Ontological critique**, examining how AI technologies update, distort, or threaten categories such as intentionality, qualia, lived body, and shared world.

This methodology provides the analytical conditions necessary to evaluate the ontological limits of artificial cognition and structures the investigation developed in the subsequent sections.

3 Results and discussion

3.1 The Machine that Diagrammatizes Reason: Inference, Iconicity, and the Hermeneutic Shift

The increasing centrality of artificial intelligence in contemporary knowledge ecologies has transformed diagrams, representations, and algorithmic models into primary operators of intelligibility. What was once an auxiliary tool for organizing ideas has become a generative mechanism that shapes the very conditions under which meaning can appear. This transition marks a profound hermeneutic shift: we no longer interpret the world directly, but a world pre-structured by machinic operations.

The philosophical roots of this phenomenon can be traced to Peirce's semiotics, particularly his concept of diagrammatic reasoning, in which inference occurs through iconic structures capable of revealing relations not previously explicit (PEIRCE, 1931–1958). Diagrams, for Peirce, are not mere illustrations; they are cognitive engines that allow the manipulation of relational structures. Contemporary AI magnifies this mechanism to unprecedented scales, generating diagrams—not visually, but mathematically—through multidimensional embeddings, attention matrices, and distributed vector spaces.

In this sense, generative models enact a new diagrammatic rationality. They operationalize inference not by symbolic deduction but by **statistical iconicity**, producing structures that function as diagrams of vast linguistic, perceptual, and behavioral corpora. What the model "knows" is not propositional content but the geometry of relations across billions of

parameters. This geometry becomes a new form of conceptual iconicity: a diagrammatic field in which meaning is statistically inferred rather than semantically understood.

The epistemological tension emerges precisely here. If diagrams traditionally served as **extensions** of human reasoning, AI internalizes them as **substrates** of its operation. What for humans is a tool for thinking becomes, for the machine, the very form of thinking. The displacement is ontological: diagrams cease to be mediations of thought and become the computational embodiment of rationality itself.

This shift has profound implications for hermeneutics. As Gadamer (1997) argues, understanding is always a historically situated event, mediated by language, tradition, and the dialogical encounter with meaning. Yet AI does not interpret; it computes. It does not inhabit a horizon of understanding; it produces probabilistic structures that **simulate** inference. The confusion between diagrammatic simulation and hermeneutic understanding risks reducing meaning to statistical regularities, displacing the event of interpretation with automated pre-interpretation.

The phenomenon has also drawn attention within contemporary philosophy of technology. Authors such as Rouvroy (2013) argue that algorithmic governmentality operates by pre-structuring what can appear as meaningful, shaping the sensible before any human judgment occurs. Similarly, Fricker (2007) warns that epistemic injustice may arise when systems curtail the conditions of intelligibility for marginalized subjects. AI's diagrammatic rationality thus has political and ethical dimensions: it governs what becomes thinkable, visible, and sayable.

A critical consequence emerges from this configuration: **humans increasingly interpret not the world, but its algorithmically diagrammatized versions**. Generative AI produces thematic graphs, clusters of meaning, predictive structures, and ontological mappings that condition the horizon of interpretation. Even when not visible, these diagrams shape perception, decision-making, and epistemic authority. The risk is that hermeneutic openness may be replaced by algorithmic closure, where the domain of the possible is pre-filtered by machinic inference.

At the same time, this diagrammatic automation reveals something essential about cognition: that reasoning is not purely internal or linguistic, but intrinsically material and relational. AI forces philosophy to acknowledge the **iconic** dimension of thought—an aspect often overlooked by representational theories—while simultaneously exposing the limits of any attempt to externalize or automate understanding without embodiment.

Artificial intelligence inaugurates a new regime of diagrammatic rationality, where inference emerges from statistical iconicity rather than embodied understanding. This transformation foregrounds a hermeneutic challenge: the automation of pre-interpretation. While diagrams have always participated in human reasoning, AI elevates them to the status of computational ontology. The result is a reconfiguration of the conditions of visibility, intelligibility, and meaning—one that prepares the ground for the next subsection, in which we examine how this automation confronts classical theories of symbolic and subsymbolic AI.

3.2 The Fold of the Mind: Cognition, Technique, and Subjectivity in the Age of AI

Artificial intelligence is no longer merely an external tool for extending cognitive capacities; it has become an ontological operator that reframes the very conditions under which cognition and subjectivity can be conceived. Its emergence unsettles traditional boundaries between thought and technique, transforming the question "Can machines think?" into the deeper inquiry: What does it mean to think when machines participate in shaping the criteria of intelligibility? In this sense, AI does not simply imitate cognition—it alters the environment in which cognition becomes recognizable as such.

Daniel Dennett's (1991) heterophenomenological model offers one influential attempt to dissolve the metaphysical weight of subjectivity. For Dennett, consciousness is not a substantive interiority but an emergent outcome of distributed patterns of information processing. The "self" is not a center but a narrative abstraction arising from computational organization. Within such a framework, a sufficiently complex artificial system could, in principle, manifest the same emergent dynamics. Yet this position comes at a cost: it flattens the lived dimension of experience and erases the phenomenological depth that characterizes embodied consciousness.

David Chalmers (1996) articulates the counterpoint through the distinction between the "easy" and the "hard" problems of consciousness. The former concerns the functional mechanisms underlying perception, decision-making, and behavior; the latter concerns why these mechanisms yield subjective experience, the felt quality of being. AI may excel at the easy problem—mapping inputs to outputs with remarkable sophistication—but remains silent before the hard problem. No increase in computational scale has yet bridged the ontological gap between function and phenomenon, between simulation and lived experience.

Between these two poles lies a third perspective: the mind as technically inflected and historically situated, neither reducible to computation nor sealed within an inner subjective domain. From this standpoint, artificial intelligence serves as a mirror that refracts human cognitive assumptions. It reveals how concepts such as "thinking," "agency," and "intention" are not timeless essences but historically formed constructs mediated by linguistic, social, and technological arrangements.

This line of thought resonates with the enactive framework of Varela, Thompson, and Rosch (1993), for whom cognition is not a passive representation of an external world but a co-emergent process enacted through embodied engagement. To know is to act, to be affected, to inhabit a relational field. AI, lacking the lived dimension of motility, vulnerability, and affectivity, remains outside this co-emergent horizon—even when its outputs mimic the structural forms of human discourse.

Yet AI's exclusion from embodied enaction does not render it irrelevant to human subjectivity. On the contrary, its presence reshapes the ecology of cognition. As algorithmic systems mediate communication, classify information, generate texts, and make decisions, they reconfigure the distribution of cognitive labor. What counts as a valid interpretation, what appears as data, and what becomes speakable or knowable increasingly depends on machinic architectures operating beneath the threshold of awareness.

In this transformation, subjectivity becomes a topological phenomenon—a set of inflections produced through interactions between bodies, environments, and technical infrastructures. AI participates in this topology not as a subject but as an ontotechnical force, modulating the conditions of subjectivation. When individuals rely on AI for diagnosis, decision-making, translation, or interpretation, they not only outsource cognitive functions—they alter the normative landscape that grounds what thinking is.

This insight invites philosophy to shift from the question of whether machines possess consciousness to the question of how machines transform consciousness. Instead of debating machine minds in isolation, the task becomes understanding how AI refracts human subjectivity, exposing its contingencies and reconfiguring its future trajectories.

Artificial intelligence functions as a fold in the contemporary configuration of cognition: a site where subjectivity, technical systems, and epistemic norms become entangled. While Dennett's functionalism and Chalmers' phenomenology frame opposing interpretations of consciousness, the enactive perspective highlights that AI transforms the ecology in which subjectivity is constituted, without itself participating in embodied experience. This prepares

the ground for the next section, which examines why disembodiment is not merely a technical limitation of AI but an ontological impossibility.

3.3 The Impossibility of a Disembodied Mind: Artificial Intelligence and the Limits of Cognition Without a Body

The question of embodiment is not an ancillary problem in the philosophy of artificial intelligence; it is the ontological hinge upon which any plausible account of cognition turns. The assumption that the mind can be abstracted from its corporeal grounding and instantiated in computational architectures expresses a residual Cartesianism that persists, often implicitly, in contemporary AI discourse. Yet phenomenology, neurophenomenology, and enactivism converge on a decisive claim: cognition is not an abstract manipulation of information but a lived, bodily mode of being-in-the-world. Without embodiment, there is no perception; without perception, no world; and without a world, no cognition.

Maurice Merleau-Ponty revolutionizes this landscape by showing that the body is not an instrument of consciousness but its very condition of possibility. The perceptual field does not arise from the brain as a computational device but from the body's motile and affective engagement with its surroundings. The world appears as a horizon of significance only because it is encountered through a body that is vulnerable, situated, and temporally stretched. In this sense, intentionality is not a representational act but a *directedness of the lived body*.

Shaun Gallagher and Dan Zahavi deepen this insight by emphasizing that consciousness is structured by habits, motor schemes, and intercorporeal relations. Cognition is not an inner procession of symbolic states but a dynamic coordination between bodily capacities and environmental affordances. The mind is not located "inside" the organism but enacted across a sensorimotor field. From this perspective, any artificial system operating without existential inscription lacks the ontological prerequisites for sense-making.

Radical embodied and enactive theories extend this claim further: mind, body, and world form a single dynamic system. Cognition emerges from the organism's ongoing regulation, adaptation, and responsiveness to an environment it co-constitutes. Knowing is an activity of *bringing forth* a world through action and affect. This co-emergence presupposes biological normativity—the capacity to be perturbed, threatened, nourished, exhausted—not merely computational responsiveness.

It is precisely this triad of vulnerability, affectivity, and normativity that artificial intelligence lacks.

Hubert Dreyfus, drawing from Heidegger and Merleau-Ponty, demonstrated that skilled human understanding does not arise from rule-following but from embodied familiarity with the world. Intelligence is not algorithmic; it is existential. AI, by trying to formalize this mode of inhabiting, commits a category error: it equates statistical regularity with lived sense. Large language models may approximate the *form* of meaningful discourse, but they do not partake in the *conditions* under which meaning appears.

The ontological discontinuity becomes even more pronounced when examining post-2022 generative systems. Models such as GPT, Gemini, or Claude operate on vast parameter spaces, but they have neither perceptual horizons nor bodily anchoring. They manipulate correlations, not significances. As Bender et al. (2021) argue, they "stochastically parrot" structures of language without any access to the experiential domains those structures presuppose.

Thomas Metzinger (2021) demonstrates that generative models lack a *phenomenal self-model*—a precondition for consciousness. N. Katherine Hayles (2021) argues that cognition is always entangled with material and affective substrates; machines, devoid of either, cannot instantiate subjectivity. Shannon Vallor (2022) shows that AI reshapes human practices and

judgments but does so without inhabiting a moral or affective horizon. David Roden (2015) contends that even radically advanced nonhuman intelligences would be ontologically discontinuous unless grounded in their own normatively organized forms of life.

These philosophical conclusions are reinforced empirically. Robots equipped with sensors, proprioceptive feedback, and adaptive algorithms may simulate balance, coordination, or gesture, yet they do not live through fear, pain, risk, anticipation, or desire. A machine can calculate trajectories but cannot inhabit danger. It can register pressure but cannot feel touch. It can adjust motion but cannot experience fatigue or relief. The phenomenological body is not a mechanical apparatus but a site of affective opening to the world.

Generative models reinforce this distance. They do not have a world—they have training data. They do not perceive—they parse. They do not feel—they optimize. They do not intend—they predict. Their outputs may evoke emotional nuance, but no emotional interiority anchors those expressions. They possess parameters, not passions; embeddings, not embodiment.

The absence of embodiment is not merely a technological limitation: it is an ontological impasse. Without a lived body, there is no standpoint from which the world can appear as meaningful. Without affect, there is no salience. Without vulnerability, no normativity. Without temporality, no selfhood. Artificial systems, however sophisticated, lack the existential conditions that make cognition possible.

Phenomenological, enactive, and posthumanist frameworks converge on a single thesis: cognition is irreducibly embodied. The absence of a lived body in artificial systems does not represent an unfinished stage of development but a structural impossibility within the current computational paradigm. AI can calculate, simulate, and predict, but it cannot perceive, intend, or experience. This ontological discontinuity sets the stage for Section 3.4, which explores how philosophical traditions have attempted—and failed—to bridge the gap between biological and artificial cognition.

3.4 The Philosophical Configuration of Human Cognition in the Age of Artificial Intelligence: A Dialogue Across Traditions

The attempt to equate human and artificial cognition has always depended on the philosophical frameworks through which "mind" is defined. Artificial intelligence does not operate outside these traditions; it activates, tests, and destabilizes them. A comparative reconstruction of major paradigms—dualism, functionalism, phenomenology, linguistic pragmatism, and extended cognition—reveals that each offers partial insights while facing decisive limits when confronted with systems that simulate cognitive functions without embodiment, world-relations, or lived experience. What follows is a synthetic map of these traditions and their relevance to the ontological problem posed by AI.

3.4.1 Dualism and the Critique of Mental Substance

Cartesian dualism established the ontological separation between *res cogitans* and *res extensa*, making conceivable an intelligence detached from corporeal life. This conceptual rupture created the intellectual conditions for imagining artificial minds as entities capable of thinking independently of organic embodiment.

Gilbert Ryle dismantled this framework by exposing the "category mistake" inherent in positing the mind as a hidden entity inside the body. Mental states are not inner substances but dispositions manifested in embodied practices. This critique prefigures later theories of

embodied cognition and undermines the idea that thinking can be instantiated as pure computation.

Yet the elimination of mental substance does not address the deeper issue of lived experience. Even without invoking an immaterial soul, computational theories reproduce a functionalized version of dualism by detaching cognition from the existential conditions that make it possible.

3.4.2 Functionalism and the Dream of the Thinking Machine

Functionalism reframed cognition as the causal organization of mental states, abstracting away from biological substrate. Hilary Putnam and Jerry Fodor argued that any system instantiating the appropriate functional roles could, in principle, possess a mind.
This view legitimized the core assumption behind symbolic artificial intelligence and early computationalism: that psychological states are software-like structures that can be ported across hardware platforms.

John Searle challenged this program through the Chinese Room argument, demonstrating that syntactic manipulation is not equivalent to semantic understanding. Functional isomorphism does not yield phenomenology; computation alone cannot generate intentionality or qualia.

Functionalism thus collapses before the distinction—now sharpened by generative AI—between formal simulation of cognition and lived experience.

3.4.3 Phenomenology and the Embodied Ground of Cognition

Phenomenology entirely reframes the problem of mind by grounding meaning, perception, and thought in bodily being-in-the-world. Maurice Merleau-Ponty argues that the body is not an accessory to consciousness but its ontological condition. Hubert Dreyfus extends this critique to AI, showing that human expertise depends on pre-reflective, context-sensitive skills that cannot be formalized into rules or statistical patterns.

Shaun Gallagher and the enactivist tradition reinforce that cognition is enacted through the sensorimotor coupling of organism and environment. Thinking is not an internal process but a relational activity.

From this perspective, artificial systems—whether symbolic, connectionist, or generative—lack the bodily intentionality, vulnerability, temporality, and affective embeddedness that constitute cognition. They can simulate forms of reasoning but cannot inhabit the world from which meaning arises.

3.4.4 Language, Forms of Life, and Shared Worlds

For later Wittgenstein, the meaning of a word is its use within a form of life. Understanding arises not from computational mappings but from participation in socially embedded practices. Language is not a code but a lived activity bound to norms, histories, and shared forms of existence.

Generative AI models can reproduce statistical profiles of language use, yet they do not inhabit the forms of life in which meaning is negotiated. They participate in linguistic performance but not in the cultural, affective, and normative contexts that give language its depth.

Thus, from a Wittgensteinian vantage, AI does not “understand”—not because it lacks computational power, but because it lacks a world.

3.4.5 Extended Mind and Distributed Cognition

The extended mind thesis proposed by Clark and Chalmers argues that cognition can extend beyond the body through the use of external tools—memory aids, computational devices, or shared artifacts. Hutchins' theory of distributed cognition similarly conceptualizes cognitive processes as emerging from hybrid networks of humans and artifacts.

While these frameworks acknowledge the role of technique in cognitive life, they do not equate artificial systems with subjects. Extending or distributing cognition presupposes a phenomenological center—an embodied agent for whom external resources acquire meaning. The mind may be extended, but it is always anchored in a biological organism.

3.4.6 Simulation, Consciousness, and the Limits of the Machine

Recent advances in generative AI intensify the tension between simulation and experience. David Chalmers distinguishes between the "easy problems" of cognition—computational and behavioral capacities—and the "hard problem" of subjective experience. Daniel Dennett rejects this distinction, proposing that sufficiently complex systems may instantiate consciousness functionally.

However, contemporary theorists reinforce the impossibility of such equivalence. Thomas Metzinger argues that generative models lack a phenomenal self-model; Hayles maintains that cognition requires material-affective entanglement; Vallor shows that algorithmic systems operate without ethical or affective grounding; and Roden contends that nonhuman minds, if they emerge, will do so outside current computational paradigms.

Generative AI thus exposes the limits of functionalist optimism: it can perform but not experience, infer but not perceive, simulate but not inhabit.

Across all examined traditions—dualism, functionalism, phenomenology, linguistic pragmatism, and extended cognition—a common tension emerges: **none can successfully ground cognition in the absence of embodiment and lived world-relations**. Artificial intelligence reveals, rather than resolves, the internal contradictions of these frameworks. It acts as a speculative mirror, showing that while machines can replicate structural features of thought, they cannot access the phenomenological conditions that make thought possible. This sets the stage for Section 3.5, which interrogates how the concept of mind itself must be reconfigured in the age of generative AI.

3.5 Thinking the Mind in the Age of the Machine: Between Ontology and Technique

The advent of generative artificial intelligence has elevated the philosophical debate on the mind to a new ontological horizon. The central question is no longer whether machines can think, but how the accelerating technicity of thought reshapes the very conditions under which thinking becomes possible. Large-scale models capable of producing language, generating images, composing arguments, or simulating reasoning reveal a paradox: they perform cognitive functions without embodiment, without a lived world, and without experiential interiority. This paradox demands a reconfiguration of classical categories—mind, subjectivity, agency, cognition—within an environment where computational systems increasingly mediate human sense-making.

The cross-traditional analysis developed in Sections 3.1–3.4 illuminates an important pattern: artificial intelligence amplifies the tensions internal to philosophical frameworks rather than resolving them. Cartesian dualism collapses for ignoring embodiment; computational

functionalism collapses for ignoring phenomenality; phenomenology insists on the primacy of lived experience; linguistic pragmatism stresses the social and historical grounding of meaning; and extended cognition shows the indispensability of embodied agents for the integration of external artifacts. The generative turn in AI does not reconcile these positions: it exposes the blind spots of each by simulating cognition without sharing the existential conditions that make cognition possible.

This discrepancy becomes clearest in the distinction between simulation and experience. Generative models can describe sensations without feeling them, articulate moral dilemmas without confronting their stakes, and produce narratives without inhabiting them. The gap is not one of computational power but of ontological constitution. As Metzinger (2021) argues, the absence of a phenomenal self-model precludes consciousness; Hayles (2021) reminds us that cognition requires material-affective entanglement; Vallor (2022) shows that algorithmic systems shape practices without sharing ethical horizons; Roden (2015) demonstrates that even the emergence of nonhuman minds would require forms of life radically distinct from our current computational paradigms. Taken together, these analyses reveal that no increase in statistical sophistication can generate a lived body, a meaningful world, or a phenomenological standpoint.

Yet the relevance of AI does not lie in its ability to mimic human cognition, but in its capacity to transform the ecological, epistemic, and affective conditions under which human cognition unfolds. Generative systems reorganize attention, alter hermeneutic practices, redistribute epistemic authority, accelerate and constrain interpretative horizons, and create new forms of dependency. They reshape the infrastructures of thought. In this sense, AI participates in cognition without possessing it: it modifies the field in which consciousness emerges and operates, even while remaining outside the domain of experience.

This insight requires a shift from the metaphysics of "artificial minds" to the ontology of technically mediated subjectivity. The mind cannot be understood as a substance (as in dualism), a functional architecture (as in computationalism), a private interiority (as in classical phenomenology), or a mere extension among artifacts (as in extended mind theory). Instead, it must be conceived as a historically situated articulation among bodies, practices, affects, and technical systems. Cognition is not an isolated property of organisms nor a transferable function: it is a relational and embodied process configured differently across human and technical ecologies.

Artificial intelligence thus reveals the singularity of the human mind less by analogy than by contrast. Just as the mirror shows the contours of a face by reflecting its absence, AI discloses the existential conditions human cognition requires: vulnerability, embodiment, affectivity, temporality, historicity, intercorporeality. These conditions do not merely accompany cognition—they constitute it.

From this vantage point, the debate about "machine consciousness" becomes secondary to a broader philosophical imperative: to understand how the technical environment transforms human modes of perception, attention, judgment, and self-formation. Generative AI does not produce consciousness, but it reconfigures the landscape in which consciousness operates. It does not simulate phenomenality, but it shapes the symbolic and epistemic conditions that structure our experience of the world. The crucial question is not whether machines will become like minds, but how minds are already becoming entangled with machines.

Some contemporary theorists—such as Dennett (2017) and Chalmers (2022)—suggest that sufficiently complex architectures could, in principle, give rise to emergent consciousness. Yet this hypothesis overlooks the normative, biological, and affective grounding of phenomenality emphasized by enactive and neurophenomenological accounts (Thompson 2021; Di Paolo & Thompson 2017). Functional complexity is not sufficient for experience;

computational capacity is not sufficient for world-relations; simulation is not sufficient for existence.

What generative systems bring into focus is that cognition is not a computational achievement but an existential condition. They force philosophy to confront what makes the human mind irreducible to technical operations—not as an exceptionalist gesture, but as an ontological necessity.

Artificial intelligence functions less as a proto-subject than as a transformative milieu that reshapes the conditions of human cognition. It cannot produce consciousness, but it reorganizes the structures through which consciousness interprets itself and the world. The conceptual distinction between machine and mind is not one of degree but of mode of being: AI operates through correlation and calculation, whereas human cognition emerges from embodied, affective, and historical engagement with the world. This ontological asymmetry frames the *Considerations Finales*, which examine the epistemic, ethical, and philosophical implications of these findings and outline directions for future research.

4 Final Considerations and Directions for Future Research

The analysis developed throughout this article demonstrates that artificial intelligence—especially in its generative phase—operates simultaneously as a challenge, a symptom, and a conceptual mirror for the philosophy of mind. Far from being a mere computational advancement, AI brings to the surface the ontological limits of classical and contemporary models of cognition by exposing the presuppositions that have historically structured the distinction between mind, body, and machine. When confronting paradigms ranging from Cartesian dualism to computational functionalism, from phenomenology to enactivism, the study reveals that none of these frameworks, taken in isolation, adequately accounts for the conditions under which cognition emerges.

The results of Sections 3.1 to 3.5 show that generative AI not only reactivates classical debates but reorganizes them under new epistemic and material conditions. The phenomenological tradition demonstrates that human cognition is rooted in lived embodiment, affective vulnerability, and temporal immersion—dimensions that cannot be replicated by architectures grounded in statistical inference. The emergence of generative systems capable of producing text, images, and inferential chains with remarkable fluency renders the distinction between simulation and experience even more urgent. While these systems perform a wide range of cognitive-like functions, they do so without any horizon of meaning, intentionality, or phenomenality.

Contemporary theorists reinforce this point. Metzinger (2021) highlights the absence of a phenomenal self-model in current AI; Hayles (2021) argues that cognition requires material-affective entanglement; Vallor (2022) shows that algorithmic systems shape human attention and moral practices without partaking in value-laden experience; and Roden (2015) suggests that if nonhuman minds ever emerge, they will arise from ontological trajectories alien to the computational paradigm that underlies generative models. Together, these analyses consolidate the central thesis of this study: the absence of embodied, affective, and world-opening existence prevents artificial systems from accessing any form of lived interiority.

At the same time, generative AI transforms the cognitive, epistemic, and symbolic ecologies in which human cognition operates. These systems reconfigure patterns of attention, mediate interpretative practices, redistribute epistemic authority, and introduce new dependencies between subjects and technical infrastructures. AI does not replicate cognition; it reorganizes the environment in which cognition unfolds. This transformation does not eliminate

the ontological gap between simulation and experience, but it alters the conditions under which human subjectivity manifests itself.

The argument developed across this article also clarifies that debates about artificial consciousness cannot be reduced to questions of computational scale or functional complexity. While some theorists—such as Dennett (2017) and Chalmers (2022)—speculate about the possibility of emergent machine consciousness, enactive and neurophenomenological perspectives (Thompson 2021; Di Paolo & Thompson 2017) show that phenomenality presupposes biological normativity, affective orientation, and embodied attunement to a world. No degree of algorithmic sophistication can generate these conditions. The mind is not a computational artifact but a relational and existential achievement.

Artificial intelligence, therefore, should not be interpreted as an imperfect model of human cognition, but as a critical operator that reveals what makes human cognition irreducible to technical simulation. It exposes the centrality of embodiment, affectivity, temporality, and historicity—features not only of how humans think but of what thinking *is*.

Directions for Future Research

Two interrelated research trajectories emerge from this inquiry, each essential for a more comprehensive understanding of cognition in the age of AI:

(1) A Critical Phenomenology of AI and Its Effects on Human Subjectivity

Future work must integrate phenomenology, science and technology studies (STS), posthumanist theory, and moral psychology to examine how AI transforms:

- perceptual and attentional patterns,
- regimes of trust, inference, and epistemic authority,
- processes of self-formation and identity,
- new forms of cognitive and affective dependency on algorithmic systems.

This line of research is not limited to assessing what AI "does"; it seeks to understand what AI does to us as embodied and intersubjective beings.

(2) Ontologies of Hybrid Cognition and Technosocial Ecologies

Generative AI requires the development of new conceptual frameworks capable of addressing:

- hybrid cognitive ecologies composed of humans and technical agents,
- shifting distributions of agency and normativity within these ecologies,
- epistemic practices mediated by probabilistic models,
- the recursive transformation of human cognition in algorithmic environments.

These inquiries align with emerging discussions on critical posthumanism, epistemic justice, and technopolitical ethics.

AI reveals that the mind is neither substance nor code, neither isolated interiority nor mere computational function. It is a historically situated, corporeally enacted, affectively mediated, and technologically embedded relation with the world. The distinction between human cognition and artificial systems is not a difference in degree of complexity but a difference in *mode of being*. Generative models transform the symbolic and epistemic world of the human without accessing the phenomenological conditions that make this world appear.

In this sense, AI does not threaten the uniqueness of human subjectivity—it clarifies it. It shows that thinking is not reducible to calculation, and that consciousness is not a technical output but a lived opening onto existence. This recognition does not diminish the importance of AI; it elevates it. For AI offers philosophy an unprecedented opportunity: to rethink the very

concept of mind, not as an immutable essence, but as a dynamic, embodied, and relational phenomenon.

At the threshold of this new technological epoch, the task of philosophy is not to replicate the mind in machines, but to understand what the presence of such machines reveals about the nature—and the limits—of human cognition itself.

References

BENDER, E. M. et al. On the dangers of stochastic parrots: can language models be too big? In: ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY, 2021. **Proceedings[...] New York: ACM**, 2021. p. 610-623. Disponível em: https://doi.org/10.1145/3442188.3445922.

CHALMERS, D. J. **Reality+: virtual worlds and the problems of philosophy.** New York: W. W. Norton, 2022.

CHALMERS, D. **A Mente Consciente: Em Busca de uma Teoria Fundamental**. São Paulo: Edições Loyola, 1999.

CLARK, A. **Natural-Born Cyborgs**. Oxford: Oxford University Press, 2003.

DEEP MIND. **Gemini: a family of highly capable multimodal models. 2023. Technical Report.** Disponível em: https://deepmind.google/technologies/gemini.

DENNETT, D. C. **From bacteria to Bach and back: the evolution of minds.** New York: W. W. Norton, 2017.

DENNETT, D. **A Consciência Explicada**. São Paulo: Martins Fontes, 1997.

DESCARTES, R. **Meditações Metafísicas**. São Paulo: Nova Cultural, 1996.

DI PAOLO, E.; THOMPSON, E. **Sensorimotor Life: An Enactive Proposal**. Oxford: Oxford University Press, 2017.

DREYFUS, H. **O que os Computadores Ainda Não Podem Fazer**. São Paulo: Editora da Unicamp, 2007.

FODOR, J. **A Linguagem do Pensamento**. São Paulo: Martins Fontes, 1981.

GALLAGHER, S. **How the Body Shapes the Mind**. Oxford: Oxford University Press, 2005.

HAYLES, N. K. Artificial Bodies in Motion: From Top-down Control to Relational Embeddedness. In: **Panel Posthuman Cognition and Material Entanglement.** Symposium: Future Bodies from a Recent Past—Sculpture, Technology, and the Body since the 1950s. Museum Brandhorst. München**,** 2021. Available at: https://www.youtube.com/live/PU7k2DjeODU?si=q4PlkonofduWz6tK

HEIDEGGER, M. **Ser e Tempo**. Petrópolis: Vozes, 2012.

KIRCHHOFF, M. D.; KIVERSTEIN, J. **Extended consciousness and predictive processing: a third wave view.** London: Routledge, 2019.

MERLEAU-PONTY, M. **Fenomenologia da Percepção**. São Paulo: Martins Fontes, 1999.

METZINGER, T. **Artificial Ego Models and Phenomenal Transparency.** 2023. (Manuscrito/Preprint).

METZINGER, T. **The Ego Tunnel: The Science of the Mind and the Myth of the Self**. New York: Basic Books, 2009.

METZINGER, T. **The Problem of a Selfless Self-Model in Generative Systems.** 2021. (Manuscrito/Preprint).

OPENAI. **ChatGPT: optimizing language models for dialogue.** 2022. Disponível em: https://openai.com/blog/chatgpt/. Acesso em: 28 de maio de 2025.

PAUL, C. et al. Embodied artificial intelligence through morphological computation. **Nature Machine Intelligence,** v. 3, n. 1, p. 8-19, 2021. Disponível em: https://doi.org/10.1038/s42256-020-00262-2.

PUTNAM, H. **Representação e Realidade**. São Paulo: UNESP, 2011.

RODEN, D. **Posthuman Life: Philosophy at the Edge of the Human.** London: Routledge, 2015.

RODEN, D. **Unbound Cognition and Technical Otherness.** 2021.

RYLE, G. **O Conceito de Mente**. São Paulo: EPU, 2000.

SEARLE, J. **Mentes, Cérebros e Ciência**. São Paulo: UNESP, 2006.

THOMPSON, E. Life and mind: from autopoiesis to neurophenomenology. **Phenomenology and the Cognitive Sciences**, v. 20, n. 1, p. 55-66, 2021. Disponível em: https://doi.org/10.1007/s11097-020-09685-z.

THOMPSON, E. **Life and Mind: From Autopoiesis to Neurophenomenology.** Phenomenology and the Cognitive Sciences, v. 20, n. 1, p. 55–66, 2021. DOI: 10.1007/s11097-020-09685-z.

THOMPSON, E. **Mind in Life: Biology, Phenomenology, and the Sciences of Mind**. Cambridge: Harvard University Press, 2007.

WITTGENSTEIN, L. **Investigações Filosóficas**. São Paulo: Nova Cultural, 2009.

ZAHAVI, D. **Fenomenologia da Autoconsciência**. São Paulo: Loyola, 2004.