

# A Transcendental Philosophy of Large Language Models

M. Beatrice Fazi 

**Abstract:** In this article, M. Beatrice Fazi responds to Shane Denson’s commentary on her paper “The Computational Search for Unity: Synthesis in Generative AI,” published in the *Journal of Continental Philosophy* in 2024. The article develops Fazi’s transcendental argument about large language models (LLMs). While Denson raises questions about conceptual relativism through Donald Davidson’s critique of conceptual schemes, Fazi maintains her position that LLMs construct “a representational world within” rather than referring to “the world.” Responding to Denson’s proposal for a model of synthesis based on the phenomenology of Jean-Paul Sartre, Fazi argues that a structuralist reinterpretation of Kantian synthesis better accounts for the operations of LLMs, where unity is that of a structure, not a self, and representation is technically central to synthetic activity. These distinctions preserve the functional aspects of Kantian synthesis without anthropomorphizing artificial intelligence, thus strengthening Fazi’s original position that LLMs create their own internal representational reality.

**Keywords:** artificial intelligence; Immanuel Kant; large language models; representation; synthesis; transcendental philosophy

I would like to thank the editors and guest editors of *Philosophy & Digitality* for the invitation to respond to Shane Denson’s article featured in this issue ([Denson 2025](#)). I have greatly valued Denson’s commentary on my essay “The Computational Search for Unity: Synthesis in Generative AI” ([Fazi 2024](#)), which was originally published in the *Journal of Continental Philosophy*. Engaging in philosophical dialogue is crucial for the development of the field of digital studies. This exchange seeks to uphold the tradition of scholarly debates that have always characterized philosophy journals. Denson raises important points about my philosophical account of synthesis in generative artificial intelligence (AI). He proposes an alternative framework based on a phenomenological approach, and that reframing presents intriguing questions about some implications of my own position. I maintain the validity of my original thesis on synthesis as a search for unity that is fundamental to the making of a representational reality.

The perspectives offered by Denson, however, have revealed dimensions of that hypothesis that warrant attention. What follows here, then, is not simply a defense of my original claims but an attempt to engage substantively with Denson's remarks. I will forgo recapitulating the principal contentions of my 2024 article, as Denson already provides an overview of them. He correctly identifies my proposal as a *transcendental argument*. In this response, I will elaborate on that transcendental stance while also addressing other points in Denson's commentary.

## Transcendental Structures

- 2 Scholastic philosophy used the qualification "transcendental" to denote universal properties or "the most general predicates of things" (Aertsen 2012, 14). However, in the second half of the eighteenth century, Immanuel Kant introduced a novel perspective in modern thought. In Kant, "transcendental" refers to the fundamental conditions that make our experiences possible: minds possess inherent structures (such as the pure forms of intuition, like space and time, and the categories of understanding) that influence how reality is constructed. Today, the concept of the transcendental continues to be of crucial importance in addressing the boundaries of knowledge, understanding, and reason. It serves as a reminder that when attempting to comprehend the world, we must recognize the impact of mental processes on experiences and how reality is constituted.
- 3 It is important to note that Kant's transcendental philosophy is not a psychological project. On the contrary, it is an argument against psychologism.<sup>1</sup> It posits that certain mental structures are intrinsic prerequisites for any experience to be possible, yet these structures are not psychological states discovered through introspection or empirical observation but the very conditions that define how we can have experiences at all, including any coherent psychological state or response. At the end of the nineteenth century, philosophers like Frege (Frege 1960) and Husserl (Husserl 1970) expanded on Kant's criticism of psychologism (see (Mohanty 1999)). While an in-depth study of these positions lies beyond my present scope, here we should consider that views against psychologism take on renewed significance when addressed alongside the operational principles of contemporary generative AI, particularly large language models (LLMs).
- 4 Within a Kantian framework, subjectivity is not psychological but structural. All experiences are shaped by necessary transcendental structures, which are universal but also subjective because they form the mind's architecture rather than the objects themselves. In my 2024 article on synthesis in generative AI, I wrote that *mind is a structure that structures*. That assertion emphasized structures not as fixed arrangements but as activities with a function. Although the

---

1. Psychologism reduces logical laws, mathematical truth, and philosophical concepts to psychological processes. Kant's distinction between "transcendental logic" and "empirical psychology" predated but influenced debates on psychologism, which today overlap with concerns about philosophical naturalism.

context of that sentence was Kantian, the direction of my argument was structuralism. I was pointing toward an understanding of structures as dynamic and transformational wholes. Historically, Kant's transcendental philosophy provided a foundation for several structuralist concepts, which emerged in part through a reinterpretation of transcendental philosophy's themes. While structuralism diverged from some fundamental aspects of the Kantian framework (undoubtedly, there are differences between the two), both approaches share a fascinating key feature: they attempt to surpass a "personal I" by abandoning the subject as a psychological posit. The notion of subject is commonly associated with an empirical entity—a subject, then, as an entity with a substance and properties that persists over time, and which experiences its world firsthand. In both Kant's transcendental philosophy and structuralism, however, the subject undergoes a significant reinterpretation. Both perspectives view the subject as shaped by prior conditions; both reject the possibility of a purely self-sufficient, fully transparent Cartesian subject having privileged access to itself and reality and being an existential foundation for knowledge.

- 5 Kant's transcendental subject is neither an individual, a specific person, nor a substantial, continuous self. Rather, it is a formal framework that makes knowledge possible for rational beings. The transcendental subject is logical, not empirical—it is an abstraction not only expressing but also establishing the possibility of any relationship to knowledge. For Kant, the transcendental subject is "the only condition accompanying all thinking" (Kant 1998, 440); it is a universal synthesizer, "the logical unity of every thought" (Kant 1998, 440) that, however, is "wholly empty of content" (Kant 1998, 419) and thus ultimately unknowable in itself. The transcendental subject is the logical prerequisite for experience, yet it is not accessible to the latter precisely because it is what makes possible that experiencing in the first place.
- 6 For structuralism, the subject is equally inaccessible: just like, for structuralists, a sign is detached from its referent and related not directly to an external reality but to other signs in a system of differences, so the subject exists as a position within a structure. Structuralist scholars have faced criticism for apparently abandoning the concept of the subject, for creating what Paul Ricoeur called "a Kantianism without a transcendental subject, even an absolute formalism" (Ricoeur 1974, 52).<sup>2</sup> This criticism, however, misses the mark. Structuralism actually presupposes a subject of experience, although a minimal one. While structuralism methodologically privileges abstract structures over individual experiences, structuralist thinkers such as Claude Lévi-Strauss, Roland Barthes, Jacques Lacan, and Louis Althusser do so to demonstrate how such "structurality" operates through subjects rather than eliminating them. Instead of erasing subjectivity, structures shape it; they give it form and possibility. Perhaps not surprisingly, lan-

---

2. This common criticism partly originated from philosophical preoccupations with lived experience, such as those of Jean-Paul Sartre, Paul Ricoeur, and Henri Lefebvre. These thinkers championed human agency against its reduction to structural determinism. Sartre (Sartre 1966) challenged Lévi-Strauss's portrayal of the human subject as an effect of underlying structures; Ricoeur (Ricoeur 1974) (Ricoeur 1991) warned against dissolving human subjectivity into linguistic systems; Lefebvre (Lefebvre 1971) extended this defense by insisting on the irreducibility of lived experience to abstract structural analysis.

guage offers a fundamental demonstration of this recursive relationship: humans become speaking subjects by entering language's pre-existing structures, yet language itself only exists through its active use by speakers.

## Real Outputs

- 7 My transcendental stance regarding LLMs aims to move beyond anthropomorphic cognitivism to address the conditions of possibility for knowledge, meaning, and understanding that are at play in these computational systems. I am thus shifting the focus from claims about AI resembling human cognition toward a philosophical perspective that understands mentality as never exclusively human in nature. Kant might not have made this leap himself, yet his conception of the transcendental subject need not be a person or a human either. With respect to LLMs, then, we can extend this transcendental framework to counteract a view of language as a purely psychological phenomenon, according to which words have meanings because of the mental associations drawn from the lived experiences of people, and grammatical rules are descriptions of human psychological habits and tendencies.
- 8 Denson identifies a radical consequence implied by my proposition: that my approach acknowledges the outputs of an LLM as *real*. I have appreciated this comment greatly. For Denson, however, this recognition of reality is inferred from a homology between human cognitive processes and the vectorial spaces of large language models. He bases this on what he perceives as a parallel I supposedly establish between computational and human synthesizing activities. Before proceeding further, I need to clarify that this latter point does not correctly represent my position. My proposal implies that synthetic language is real, but I do not argue that LLMs mimic human cognitive synthesis. In my view, LLM outputs are real, rather, because all language production (human or otherwise) synthesizes representations into unified structures.
- 9 This key distinction helps me to delineate the contours of the kind of transcendental philosophy of large language models that I want to argue for. I am not advocating for the application of Kant's transcendental framework of necessary structures for reasoning to AI as a means of enabling human-level understanding within computational systems. Equally, I am not proposing to transplant Kantian cognitive structures into AI programs. Mine is not a "computational Kantianism" (Wolfendale 2015) (Wolfendale 2016), that is, an interpretation of transcendental philosophy through the lens of computational theory, claiming that Kantian categories and forms are analogous to computational operations. Kant did not describe something necessarily human; he presented a universal normative framework for reasoning. However, I am unconvinced by a computationalist reworking of such inhuman infrastructure because I am skeptical of computationalism in general, particularly its assertion that the mind is a computational process. I do not believe that mind necessarily has to be equated with computation, but I am interested in entertaining the inverse speculative hypothesis, namely, that *computation could be a*

*type of mind*, which does something akin to what a mind is believed to do (e.g., thinking).

10 Drawing from philosophical traditions that understand synthesis as a process of amalgamation, composition, and combination, my 2024 article reassessed the qualification of “synthetic” in today’s widely used expression “synthetic media” (i.e., media content produced by generative AI). As in all my work, I wanted to surpass the “simulative paradigm” that has constrained philosophical speculation in AI research. Such a simulative paradigm is centered on an imitative relationship between machine and human cognition (see (Fazi 2019)). AI-generated content has often been labeled as “synthetic” to highlight its manufactured origin; by focusing on the philosophical elaborations of the concept of synthesis, I aimed to demonstrate that LLMs are not synthetic because they create fake language but because they generate outputs through unifying processes. Moving beyond an understanding of the synthetic as something artificial requires acknowledging that what synthetic media produce is not merely an imitation.

11 Since large language models gained mainstream attention in 2022 (after OpenAI’s release of ChatGPT, which accelerated the widespread adoption of these technologies), plenty of opinion pieces, academic papers, and news items have put forth the view that machine-generated language lacks authenticity. It is a simulation of communication, we are told, or text that is coherent but hollow. While Emily M. Bender and her co-authors, in the widely cited “stochastic parrots” article (Bender et al. 2021), make compelling claims about the limits of LLMs, the bird metaphor has been employed (in that essay and then elsewhere) precisely to highlight that these models produce computational approximations of language, indeed *parroting* human-making processes. The most influential living linguist, Noam Chomsky, has also been vocal on this point, dismissively labeling LLMs as a kind of glorified autocomplete and upholding, unwaveringly, that these systems do not generate genuine linguistic expressions (see (Chomsky, Roberts, and Watamull 2023)). The list of skeptical voices could go on and on, with the debate touching on central questions about the nature of language, understanding, and intelligence.

12 A healthy dose of skepticism is necessary when addressing and assessing claims made by AI businesses that operate within political economies, pursuing technosocial and corporate agendas. Even with all due reservations, however, acknowledging AI’s differences from human language should serve as a starting point rather than a conclusion in discussions about these technologies. The non-human nature of LLM outputs is a given. However, if AI is not writing or speaking like humans do, the decisive question becomes: what distinctive processes actually characterize its outputs? Which is to ask, if those textual outputs are not like human language, then *what are they?*

13 Dichotomizing between something genuine/natural and something spurious/artificial does not serve a useful purpose. I am confident in making this claim because I do not share the Heideggerian preoccupation with authenticity. Do not get me wrong: I am all for sincerity, honesty, integrity, and being true to oneself, but this is not what is at stake here, in this debate, which is ideologically limited to (and by) a forced binary choice between “true understanding” on the one hand, and “mere statistics” on the other. Any authentic/inauthentic contraposition

fails to account for the complexity (and philosophical potential) of the computational production of language, and also for the complexity (and philosophical potential) of human life and human thought (as plenty of critics of Heidegger's concept have already noted; see, for instance, Adorno ([Adorno 2003](#)) and Derrida ([Derrida 1976](#))). I should thus clarify that, from my point of view, recognizing the reality of generative AI's outputs does not necessarily imply that machines possess understanding. I keep an open mind and reserve judgment regarding the possibility of AI having meaning-making capacities as I continue studying these systems. For now, however, I am asserting that the AI's synthetic outputs are not merely derivative of human language. LLMs recombine and process pre-existing linguistic patterns (this is the stochastic character of their operations). Yet this stochastic processing does not diminish their world-making capacity. This world-making potential forms the kernel of my proposition on synthesis in generative AI, as Denson correctly recognizes. I will address this aspect next.

## Epistemic Pluralism

- <sup>14</sup> Denson's article explores the philosopher Donald Davidson's views on incommensurability, as presented in the essay "On the Very Idea of a Conceptual Scheme," published in the Proceedings and Addresses of the American Philosophical Association in 1974. For Davidson, it does not make sense to talk of incommensurable conceptual schemes insofar as there is no neutral position from which to make such evaluation of incommensurability and because, once something has been recognized as language, it is already implied within it a degree of translatability (this is Davidson's "principle of charity," stating that, in order to identify something as language, one must be able to understand a part of what is being said).
- <sup>15</sup> I have studied extensively the concept of incommensurability in my earlier work (see ([Fazi 2021](#))). Incommensurability is a notion originating in ancient Greek mathematics; in the 1960s, it was developed (independently but in parallel) by the philosophers of science Thomas Kuhn and Paul Feyerabend, who claimed that conceptual frameworks, theories, and languages cannot be directly compared with one another because they are born and used in the context of different worldviews. Drawing from those mid-twentieth-century debates, I have argued that incommensurability offers a powerful speculative framework to address the contemporary quest for explainability in AI. I proposed that the latter should be understood as a representational and communicational issue. Although Denson does not reference my work on incommensurability and AI, his engagement with Davidson's critique of this concept establishes a connection to this earlier work of mine, to which I will thus link back.
- <sup>16</sup> Davidson's 1974 critique of incommensurability is also a criticism of how Kuhn and Feyerabend employed the notion and a challenge to the conceptual relativism that Davidson perceived as intrinsic to their positions. Denson agrees with Davidson's conclusions and applies them to the study of generative AI, suggesting that my position on the world-making potential of AI might

indirectly also reinforce aspects of conceptual relativism. I disagree with the suggestion that my view leads to conceptual relativism, and I find Davidson's original critique of incommensurability in philosophy of science to be based on a mischaracterization. Davidson defines conceptual relativism as the perspective that "reality itself is relative to a [conceptual] scheme: what counts as real in one system may not in another" (Davidson 1974, 5). This interpretation, however, does not accurately represent the concept of incommensurability as developed in Kuhn's and Feyerabend's respective works.

17 Kuhn refined his conceptualization of incommensurability throughout the 1980s and 1990s. Already in 1970, though, in the postscript to the second edition of *The Structure of Scientific Revolutions*, he had explicitly rejected the label of "relativist," clarifying that his account of paradigm shifts in science does not deny scientific progress or infer that everything holds equal validity. His later work (for instance, (Kuhn 2000) (Kuhn 2000b)) expanded on this, signposting incommensurability as a "local" rather than "total" phenomenon and thus highlighting how incommensurability might not apply to whole languages and theories but in fact to specific groups of related terms or domains. While distancing himself from radical versions of social constructivism, later in his life Kuhn also wrote about incommensurability in terms of a "taxonomic" issue (Kuhn 2000c), stressing how dissimilar ways of thinking can organize ideas differently without blocking communication. Feyerabend, on the other hand, was renowned for his epistemological anarchism; he openly embraced positions that could be seen as somewhat relativistic, for example declaring that "the only principle that does not inhibit progress is: anything goes" (Feyerabend 1993, 14). His brand of relativism, nonetheless, consistently stood in opposition to Davidson's argument and evolved over time to account for similar philosophical challenges (see (Kusch 2016)). For Feyerabend, the proliferation of theories is beneficial to science, and there are not universal methodological rules that are always appropriate or useful. This is not an impediment to science that would render the latter aphasic, disoriented, meaningless. On the contrary, this is how science democratically progresses. Feyerabend recognizes the contextual nature of all knowledge claims but never implied that all knowledge systems are the same or that there is no objectivity. For Feyerabend, difficulties in translation and communication can be overcome via practice and immersion; rationality and certain aspects of incommensurability are therefore compatible.

18 The technical details of this debate in philosophy of science are fascinating but not essential for my response to Denson. Instead, I must move on here and explain why and how I engaged with the concept of incommensurability in my past research (in particular, in (Fazi 2021)). At that time, I was studying explainable AI (XAI) and the concept of the "black box" that often accompanied those discussions. While machine learning gained more and more prominence as the "new AI" (Alpaydin 2016), many scholars were concerned about the characteristically opaque nature of these automated technologies. Deep learning, especially, was receiving particular scrutiny. Artificial neural networks have millions or billions of parameters, interacting in complex, non-linear ways often not interpretable by humans, who struggle to understand why specific outputs have followed certain inputs. Moreover, the representational character of the automated extrac-

tion of features from high-dimensional data also presents challenges, since the distributed representations the machine employs do not correspond to human ones.

19 Philosophically, I found the situation captivating. However, I also found the responses that were often offered to this condition underwhelming, insofar as grounded on calls to make algorithms transparent to humans via interfaces for human-machine translation. I argued that such XAI methods were limited because they assumed machine abstractions can be meaningfully translated into human terms. It is within this argument that I mobilized the concept of incommensurability to address the relationship between human and what I termed “algorithmic thought,” that is, between different modes of representation that cannot be measured by a common standard. This incommensurability exists because these different modes of representation originate from distinctive onto-epistemic premises and share no common existential grounds. In making this argument, however, I never implied that no comparison is possible; on the contrary, I argued that a form of comparative epistemology, in this context, is difficult but necessary. This is not relativism; rather, it is a recognition of the limitations of human comprehension when confronted with specific forms of computational abstraction and an acknowledgment that such epistemic differences are grounded in ontological distinctions between life and mechanism (see [\(Fazi 2019\)](#)).

## Worlds Colliding

20 My 2024 article on synthesis in generative AI indirectly builds on my 2021 incommensurability argument. In that article, I explore how “reference” is the relationship between a linguistic expression and what that expression represents in the world. Semantic theory maintains that words acquire their meaning from referring to things. LLMs lack this connection, which is instead present in, and unequivocally central to, the operations of natural languages. This absence is often used as evidence of the lack of authenticity that LLMs purportedly suffer from. I argued instead that, while it is true that LLMs do not have direct access to the external world, they nonetheless build their own *internal representational reality*. This is a sum of word embeddings and their internalized relations—not the world as such but *a* world, with its own structural coherence.

21 Denson objects to my distinction between “the world” (the world external to the machine, that is, our own human world) and “a world” (the LLM’s self-contained world; a “world within,” as I described it). He builds on Davidson’s critique of conceptual schemes to challenge the idea that an LLM’s representational reality is distinct from ours. Drawing from Davidson, Denson contends that if we accept LLM outputs as language, we must concede that these outputs have a connection to the same world that we inhabit. He gives the following example: when we identify “hallucinations” in LLM outputs, we are implicitly supposing that LLMs are trying to express something about a shared dimension of reference. Our practice of perceiving these outputs



as errors necessitates the assumption of a common world, a backdrop of commonality. Were LLMs to operate in a truly distinct world (that is, in Davidson's terms, a distinct "conceptual scheme"), we could not discern these errors as errors at all. So, for Denson, despite their differences, LLMs share the same mediated world with humans, with the same network of meanings, although partially and imperfectly.

22 Denson's argument helps us to appreciate the importance of establishing a commonality between humans and machines. This is not only a central technical and philosophical necessity for contemporary AI research, but also a collective challenge confronting us on social, cultural, political, and ethical fronts. How can humans *meaningfully* coexist with machines? In this regard, however, the example of AI hallucinations reveals some of the key differences between Denson's approach and my own. AI hallucinations occur when an AI system generates plausible but factually incorrect or fabricated information, including false citations, fictional entities, made-up events, and logical inconsistencies. I contend that both the impression of plausibility and the recognition of fabrication reside not in the machine but in us; they belong firmly within the human world. The point I am making echoes what I have always been arguing about glitches and that kind of art style built on them: glitches exist in the eye of the beholder. Artistic experimenting with them constitutes a very valuable critical and aesthetic praxis, which problematizes that powerful dogma of seamless experience the tech industry feeds to its consumers. Glitches, indeed, disrupt users far more than the machine that produces them. If a glitch is the result of machines generating more than intended, that surplus is, however, ours: it belongs to an experiential human realm that receives it and judges it, not to the machine itself.

23 While technically not glitches, AI hallucinations are similarly phenomena framed through a human perspective. When an LLM generates information that appears plausible but is in fact inaccurate, the evaluation of that output as an error is situated within a human framework. AI hallucinations appear as errors to the human interpreters of these machine outputs. For the machine, however, this condition is not a bug but a feature; it is not a limitation but an inherent characteristic of statistical, prediction-based systems. LLMs predict text sequences based on patterns learned during their training. When confronted with uncertainty, these models generate the most probable continuation of words. There is no shared ontology to assume here, only shared tokens (a token is a unit of text that the model processes). While the exchange of these textual units favors a degree of interaction between humans and machines, these tokens also carry certain references for their human users. Such references, however, are not necessarily relevant or constructed in the same way by the machine that creates and sustains the text. Humans and machines exchange text but not necessarily its meaning. LLMs do not attempt to refer to our world when they operate, although we expect or would like them to, for that is the instrumental end we assign to this kind of technology. This is obviously problematic from the point of view of humans relying on the veracity of LLM outputs. Yet this is not per se an issue if one takes the perspective of the LLM itself, for which the outputs are true insofar as coherent with its encodings of representations. There is autonomy in the computational automation of language: LLMs are not just processing data that originates from humans; they construct their own consistent

organization of such data. This structuring is not a mere distorted reflection of the human world; it is world-making in its own right.

24 While Denson is careful to avoid assumptions that could emerge during potentially unreflective interactions with these systems, we nevertheless reach different conclusions regarding the example of AI hallucination introduced in his article. I agree that we must coexist with machines and thus study how a shared reality can be possible. In my view, however, the worlds in question, are really three: the machines', the humans', and *a third, shared one*. Consider this parallel: if aliens from another planet made contact with Earth, they would still approach us from their own world. Any communication between them and us would still need to address that extraterrestrial reality as alien to ours while attempting to construct a common space for potential cooperation. Similarly, an LLM's internal representational reality is a world colliding with ours. Interestingly, Denson also metaphorically mentions aliens while acknowledging LLMs' radical alterity. But where Denson frames the interaction with LLMs as welcoming these alien beings into *our* world (thus suggesting assimilation of AI into the human sphere), I view this encounter as producing a systemic shock to both realities. Not a war of the worlds, necessarily, but "worlds" in the plural nonetheless.

25 I insist on this because the medium specificity of computation extends beyond machines lacking embodied sensation, perceptual grounding, and intentionality. As I argued in my 2021 work on incommensurability, machines and humans share no common experiential ground. They operate from different ontological and epistemological positions. The crux, therefore, is not per se untranslatability. Davidson claimed that language cannot be separated from translation, as translation, for him, is interpretation. Responses in philosophy of science have already exposed the errors and limits of equating interpretation with translation (see ([Feyerabend 1987](#))). The twenty-first-century emergence of a computational (sub)symbolic order so fundamentally alien to the human one can help us to further dismantle such restrictive assumptions about the nature of understanding.<sup>3</sup>

## Models of Synthesis

26 This discussion returns us to the philosophical concept of synthesis, which should here be connected to the world-making I claim is occurring in these AI systems. Denson takes seriously my proposition that LLMs operate through synthesizing activities and agrees with my philosophical understanding of synthesis as the generative principle behind LLM outputs. However, he does not support the Kantian model of synthesis that I adopt. This is because, in his view, that Kantian framework (1) requires a certain commitment to unity, and (2) is still anthropocentric. I

---

3. Sankey ([Sankey 1997](#)) examines the relationship between translation and understanding, considering this vis-à-vis debates on incommensurability in philosophy of science and in relation to Davidson's argument.

accept point 1 without apology, a position I will elaborate in the next section of this article. I dispute, however, point 2: Kant's philosophical notion of synthesis does not confine us in any human-centric trap, as Kant's transcendental subject need not be human. I developed this argument at length already in the opening of my response.

27 Let us address the alternative model of synthesis offered by Denson to understand LLM operations—a model that draws not from Kant but from early Sartre ([Sartre 2004](#)) and Husserl's notion of "presencing." This approach repositions the site of synthesis, allowing Denson to conceptualize LLM operations not as self-contained but intertwined with human meaning-making. To understand this, it is useful to recapitulate that I see LLMs as constructing a self-enclosed "world within" that is indifferent to human linguistic reference. Denson, conversely, views LLMs as active participants in the human world through mediated forms of referencing. Clearly, by invoking Sartre's reworking of Husserl, Denson is borrowing from phenomenology—a tradition of philosophical thought that has consistently challenged structuralism's treatment of language as a closed system of signs. However, reducing the difference between Denson's and my own approach to a mere disagreement between phenomenology and structuralism would be an oversimplification. While, evidently, I lean more toward structuralism than Denson, and he more toward phenomenology than I, these categorizations risk overlooking the nuances of our respective positions. If I am a structuralist, I am quite an unconventional one. I have engaged deeply with structuralist themes in studying LLMs, but my philosophical stance extends beyond orthodox structuralist frameworks because my primary interest lies not in the structures that make human culture but rather in conceptualizing thought itself as structure.<sup>4</sup> Denson's media phenomenology is similarly multifaceted. Generally, I have reservations regarding phenomenology as a framework for the philosophy of technology because phenomenological approaches often dilute the distinctive specificity of technological systems to map that back, reductively, onto a human dimension. Denson's work, however, surpasses this reduction, building on and extending phenomenology to account for contemporary technological mediations of experiences, such as those happening beyond the threshold of human perception and reception (see ([Denson 2020](#))).

28 How might we then interpret Denson's proposal to set aside Kantian synthesis and look instead elsewhere, to Sartre (via Husserl) specifically? Denson suggests to move to a "pre-personal" plane—which Sartre's conception of synthesis would allow for—so as to inhabit a "pre-reflective" condition of being in the world that could be shared both by humans and LLMs. According to this view, synthesis should be separated from operations of representation. This is something that cannot possibly be done within a Kantian framework, for which synthesis is fundamentally tied to representation. In his 1772 letter to Markus Herz, Kant ([Kant 1967](#)) asked

---

4. For a study of large language models that embraces the structuralist understanding of language as an inherently cultural system, see Weatherby ([Weatherby 2025](#)). While my focus tends toward structuralist conceptions of thought, Weatherby and I agree that both language and LLMs possess their own internal logic. Particularly insightful in this respect is his critique of what he calls "the ladder of reference," that is, the hierarchical view that treats reference to external reality as the language's supposedly primary function.

how subjective representations in our minds can ever transcend their own boundaries and access external objects—this relation, he wrote, was the key to the whole secret of metaphysics; significantly, it became the central problem Kant sought to solve throughout his philosophical life. More than a century later, Husserl employed a “phenomenological reduction” to bypass the dualism of subject-object, avoid representationalism, and focus instead on intentional consciousness. Building directly upon Husserl’s phenomenology but moving toward a more existentialist direction that would privilege concrete experiences, Sartre pushed this approach even further, understanding consciousness as existing only in relation to and via engagements with objects. Denson recalls this philosophical trajectory when arguing to shift our attention toward a different kind of synthetic activity. Indeed, the idea of synthesis changes dramatically from Kant to Husserl to Sartre. By dispensing with the requirement for a unified mode of subjectivity, Denson argues, we can better understand how LLMs generate language that, while distinctively computational, still refers to our human world.

29 As previously discussed, however, Kant’s transcendental subject does not necessarily entail human-like subjectivity. I see the claim that phenomenological synthesis allows us to move beyond Kantian anthropocentrism resting on a faulty premise, then, since Kant’s transcendental philosophy itself can be understood as not anthropocentric. Denson’s suggestion, nonetheless, carries a significant implication beyond the point on anthropocentrism: while not advocating for eliminating the subject entirely, he seeks to identify a model of synthesis that precedes subject formation and provides common ground between human and machine processes. Sartre indeed presents such a subjectless model, establishing an ego-less flux of impersonal experience, as Denson writes, which existentialist phenomenology considers more fundamental than representational thinking. In relation to LLMs, this flux seemingly enables a plateau of commonality for the sort of indirect continuity of reference between humans and machines Denson advocates through his argument for a continuity of worlds.<sup>5</sup>

30 Although Denson states that it is beyond the scope of his commentary to my 2024 article to pursue this Sartrean line of thinking in detail, I want to offer a reply via two counterpoints: first, having already established that “transcendental subject” need not mean “human subject,” I maintain that unity is the goal of synthesis, and, second, I believe representation cannot be bracketed in this debate (as Denson suggests) because it remains fundamental to both computing and synthetic activity. I will elaborate these claims in the final part of my response.

## Unity of Representations

31 Kant’s *transcendental unity of apperception* serves as a necessary formal condition that enables the

---

5. In linguistics, *continuity of reference* is a concept explaining how speakers are able to consistently identify entities throughout conversations, texts, etc.

synthesis of all representations (perceptions, sensations, and thoughts) into a coherent whole. This synthetic combination is transcendental because it is not an empirical psychological feature but a logical prerequisite for the very possibility of experience itself. It constitutes the junction where thought and being, subject and object, self and world convene. Denson interprets my 2024 argument for synthesis in generative AI as making a case for something akin to a transcendental unity of apperception for LLMs. Indeed, this was one of the central philosophical operations I attempted in that article, though with an important caveat: I did not claim that LLMs behave identically like human minds. Kant developed his transcendental philosophy in the context of human cognition, describing the transcendental apperception as “the highest point to which one must affix all use of the understanding” (Kant 1998, 247). However, he never explicitly restricted this framework to humans alone; the transcendental subject too, as we have seen, is logical rather than empirical. While a parallel with human minds is thus certainly possible, I am not seeking a computational equivalent of Kant’s model. My goal is different: I am developing the speculative hypothesis that LLMs perform a fundamental search for unity and studying Kant’s model of synthesis to consider what that unity, in a computational space, could be.

32 I previously highlighted that my 2024 article advanced a Kantian argument but with a structuralist trajectory. This observation bears relevance now as we address a philosophical challenge Denson identifies. In Kant, synthesis is the endeavor through which representations are unified by a subject, yet this subject is itself a product of synthesis. This condition creates a loop: the unity of the “I think” is necessary for synthesis but also results from synthetic activity. The circular character of the “I think,” Denson comments, emphasizes again the issue of anthropocentrism. How to resolve this tension? Here it is important to consider how my Kantian argument for synthesis in generative AI is extended and also reoriented by the structuralist direction of my 2024 essay. My claim is that *unity, in computation, is that of a structure, not of a self*. The circularity is thus self-reflexive rather than self-reflective—this is the unity of a structure that references itself, not that of a conscious subjectivity engaged in reflection. Such a structural unity is what allows an LLM to generate its own internal representational reality with its own coherence, and what allows the (imperfect) stability of an LLM’s “world within.” I stand by my philosophical position on unity, and on synthesis as the process by which such unity is sought, because I remain committed to theorizing that computational world-making.

33 For structuralism, unity is not based on essence or substance. In the case of subjectivity, then, it would not be based on selfhood either. Unity is instead the result of the systematic organization of structural relationships, generating a whole that remains dynamic. What I described as *synoptic computing* is indeed a “seeing all together” meant to afford a view of a whole. I coined the expression “synoptic computing” to highlight how LLMs search for unity to produce language through relationships between representations in the internal network. This structural whole takes precedence over individual components and, from the point of view of the LLM itself, internal coherence is more important than direct correspondence to external reality. The reader can see, then, that my effort to determine whether an LLM is a structure that structures or a unity that unifies captures how, for Kant, synthesis is both a characteristic of mind and an activ-

ity that mind performs. I am not assigning it self-awareness, although there is a level at which structuring itself acts as a transcendental subject, a form of subjectivity that is achieved algorithmically rather than arising from intentionality or consciousness.

34 To clarify again, if mine is a structuralist approach (opposed to the distinctively more phenomenological one Denson is adopting), then it is still a peculiar structuralism. Worth emphasizing here is that, historically speaking, structuralism treated representations as symbolic entities carrying cultural meanings, which is to say, it saw culture as a system of signs. In LLMs, however, representations are distributed and implemented as high-dimensional vectors in artificial neural networks. These representations do not correspond to cultural concepts in one-to-one mappings; meaning is not localized in discrete symbolic units but encoded across statistical patterns throughout the network. Denson proposes a Sartrean model of synthesis precisely because, in his view, it better accounts for such dispersion and continuity of meaning and meaning-making. His essay does not expand on this issue, but we could note that such a phenomenological approach attempts to describe the fluidity of these distributed representations in a way that the symbolic order of traditional structuralism cannot. Denson objects to my commitment to representation, however, not because of the structuralist leaning of my overall argument but because of its Kantianism, writing that Kant guides my thinking about representation. He argues for separating synthesis from representation, a split that can be performed in the phenomenological register he moves to.

35 Our main theoretical divergence here is that I do not believe it is ever possible to set aside representation when dealing with computation. Computing machines are *representational machines*. I am thinking, of course, of how traditional computing employs symbolic encoding and processing to work on information but also of how contemporary artificial neural networks leverage their subsymbolic nature to elevate representation to an autonomous (or quasi-autonomous) modeling of complex relationships. Because computation is, in both its symbolic and subsymbolic instantiations, inescapably representational, I maintain we need a model of synthesis that can address—boldly and directly—this representational character. The phenomenological tradition (especially its Husserlian lineage) offers the notion of presencing to account for a direct encounter with phenomena rather than their representation. Indeed, this phenomenological notion attempts to describe an access to phenomena that minimizes (or, in some cases, bypasses) representational mediation. I diverge from this phenomenological view substantially because, just as I do not believe that it is ever possible to set aside representation when dealing with computation, I do not believe there is anything like an unmediated access to reality. I understand that presencing is an appealing concept to address the kind of continuity of reference between humans and LLMs Denson is arguing for. I also appreciate that this proposal has to be contextualized vis-à-vis a long history of phenomenological critique that has interpreted Western philosophy as heavily mediated. My position, however, is that the concept of presencing is much more anthropomorphizing than that of representation, insofar as it describes an encounter with phenomena *in lived experience*. I do not endorse a split between representation and synthesis in computation because of the latter's specific onto-epistemic dimension, which ex-

cludes it from being predicated on an ontology of life (understood both as the lived and the living).

36 The idea of an AI “world within” aims not only to capture the genuine strangeness of LLM outputs but also to establish a transcendental philosophy of these programs. The deepest challenge of transcendental philosophy—as formulated by Kant but then developed, in one way or another, by virtually all philosophical traditions, including phenomenology and structuralism—is understanding how subjects constitute themselves and their world while simultaneously being objects within that world, among countless other objects. LLMs are constrained by their architectures and training methods. Yet in relation to these very constraints, the notion of representation is key to comprehending how the synoptic computing of these systems is generative—how it creates new structures rather than simply receiving and analyzing pre-existing ones. As I argued in my 2024 article, computational processes produce themselves alongside their own worlds.

37 This philosophical exchange Denson and I have been developing demonstrates that culture, society, and technology are all asking for novel philosophical theories to understand synthetic media and their outputs. Denson and I share this conviction and are committed to doing the work. While points of disagreement persist between our respective approaches, this conviction and commitment suggest a fundamental compatibility in spirit between our otherwise different proposals. This shared goal serves also as a cornerstone for our ongoing dialogue.

## Bibliography

Adorno, Theodor W. 2003. *The Jargon of Authenticity*. Translated by Knut Tarnowski and Frederic Will. London: Routledge.

Aertsen, Jan A. 2012. *Medieval Philosophy as Transcendental Thought*. Leiden: Brill.

Alpaydin, Ethem. 2016. *Machine Learning: The New AI*. Cambridge, MA: MIT Press.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜.” In *FAccT ’21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23.

Chomsky, Noam, Ian Roberts, and Jeffrey Watamull. 2023. “The False Promise of ChatGPT.” *The New York Times*, March 8, 2023. <https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html>.

Davidson, Donald. 1974. “On the Very Idea of a Conceptual Scheme.” *Proceedings and Addresses of the American Philosophical Association* 47:5–20.

- Denson, Shane. 2020. *Discorrelated Images*. Durham, NC: Duke University Press.
- . 2025. “On the Very Idea of a (Synthetic) Conceptual Scheme.” *Philosophy & Digitality* 2 (1): 199–216.
- Derrida, Jacques. 1976. *Of Grammatology*. Translated by Gayatri Chakravorty Spivak. Baltimore, MD: The John Hopkins University Press.
- Fazi, M. Beatrice. 2019. “Can a Machine Think (Anything New)? Automation Beyond Simulation.” *AI & Society: Knowledge, Culture and Communication* 34 (4): 813–24.
- . 2021. “Beyond Human: Deep Learning, Explainability and Representation.” *Theory, Culture & Society* 38 (7–8): 55–77.
- . 2024. “The Computational Search for Unity: Synthesis in Generative AI.” *Journal of Continental Philosophy* 5 (1): 31–56.
- Feyerabend, Paul. 1987. “Putnam on Incommensurability.” *British Journal for the Philosophy of Science* 38 (1): 78–81.
- . 1993. *Against Method*. 3rd ed. London: Verso.
- Frege, Gottlob. 1960. *The Foundations of Arithmetic: A Logico-Mathematical Enquiry Into the Concept of Number*. Translated by J.L. Austin. 2nd ed. Evanston: Northwestern University Press.
- Husserl, Edmund. 1970. *Logical Investigations*. London: Routledge.
- Kant, Immanuel. 1967. “To Marcus Herz, February 21, 1772.” In *Philosophical Correspondence 1759–99*, edited and translated by Arnulf Zweig, 70–76. Chicago: University of Chicago Press.
- . 1998. *Critique of Pure Reason*. Edited and translated by Allen W. Wood and Paul Guyer. Cambridge: Cambridge University Press.
- Kuhn, Thomas S. 2000a. “Commensurability, Comparability, Communicability.” In *The Road Since Structure: Philosophical Essays, 1970–1993, with an Autobiographical Interview*, edited by James Conant and John Haugeland, 33–57. Chicago: University of Chicago Press.
- . 2000b. “Reflections on My Critics.” In *The Road Since Structure: Philosophical Essays, 1970–1993, with an Autobiographical Interview*, edited by James Conant and John Haugeland, 123–75. Chicago: University of Chicago Press.
- . 2000c. “The Road Since Structure.” In *The Road Since Structure: Philosophical Essays, 1970–1993, with an Autobiographical Interview*, edited by James Conant and John Haugeland, 90–104. Chicago: University of Chicago Press.



- Kusch, Martin. 2016. "Relativism in Feyerabend's Later Writings." *Studies in History and Philosophy of Science Part A* 57:106–13.
- Lefebvre, Henri. 1971. *Au-Delà Du Structuralisme*. Paris: Editions Anthropos.
- Mohanty, Jitendra Nath. 1999. *Logic, Truth and the Modalities: From a Phenomenological Perspective*. Dordrecht: Springer.
- Ricoeur, Paul. 1974. *The Conflict of Interpretations: Essays in Hermeneutics*. Edited by Don Ihde. Evanston: Northwestern University Press.
- . 1991. *From Text to Action*. Translated by John B. Thompson and Kathleen Blamey. Evanston: Northwestern University Press.
- Sankey, Howard. 1997. *Rationality, Relativism and Incommensurability*. London: Routledge.
- Sartre, Jean-Paul. 1966. "Jean-Paul Sartre Répond." *L'Arc* 30.
- . 2004. *The Transcendence of the Ego: A Sketch for Phenomenological Description*. Translated by Andrew Brown. London: Routledge.
- Weatherby, Leif. 2025. *Language Machines: Cultural AI and the End of Remainder Humanism*. Minneapolis, MI: University of Minnesota Press.
- Wolfendale, Peter. 2015. "The Reformatting of Homo Sapiens." *Fridericianum*, August 10. <https://www.youtube.com/watch?v=1IpDTUhQA5U>.
- . 2016. "Computational Kantianism." *deontologistics*, October 15. <https://www.youtube.com/watch?v=EWDZyOWN4VA>.