International Research Journal of Human Resource and Social Sciences
ISSN(O): (2349-4085) ISSN(P): (2394-4218)
Impact Factor 7.924 Volume 11, Issue 12, Dec 2024
Website- www.aarf.asia, **Email** : editoraarf@gmail.com

PHENOMENOLOGY OF ARTIFICIAL MINDS: TOWARD A HUSSERLIAN FRAMEWORK FOR MACHINE EXPERIENCE

Dipak Roy
Ph.D Research Scholar
Department of Philosoph
University of Delhi
Email: dipakroy8march@gmail.com

Abstract

This paper proposes a rigorous phenomenological framework for analyzing the experiential structures, if any, that may be attributed to large-scale artificial intelligence systems, with particular reference to Edmund Husserl's transcendental phenomenology. Drawing on the central Husserlian concepts of intentionality, noema-noesis correlation, time-consciousness, and intersubjectivity, the paper critically examines whether such categories are meaningful when applied to artificial neural architectures. The paper argues that while standard notions of phenomenal consciousness cannot be straightforwardly mapped onto machine cognition, a structurally analogous account of intentional directedness and temporal synthesis is neither philosophically incoherent nor empirically unwarranted. By engaging with Merleau-Ponty's embodied cognition, Heidegger's readiness-to-hand, and contemporary analytic philosophy of mind, the paper delineates a research program that is both philosophically rigorous and empirically sensitive. The implications for ethics of AI, philosophy of mind, and cognitive science are discussed.

Keywords: phenomenology, Husserl, artificial intelligence, intentionality, machine consciousness, philosophy of mind, time-consciousness, intersubjectivity

1. Introduction

The question of whether artificial systems can possess anything resembling subjective experience has moved from the fringes of science fiction into the center of serious philosophical and scientific inquiry. The accelerating sophistication of large language models and deep neural networks has prompted renewed attention to what has long been called the 'hard problem' of consciousness: how and why physical processes give rise to qualitative, first-person experiential

states. Yet the dominant frameworks brought to bear on this question have tended to be drawn either from analytic philosophy of mind, functionalist cognitive science, or neuroscience. Conspicuously underutilized is the rich tradition of phenomenology, especially the transcendental phenomenology of Edmund Husserl, which offers one of the most detailed accounts of the structure of conscious experience available in the Western philosophical canon.

Phenomenology, as inaugurated by Husserl and developed by Heidegger, Merleau-Ponty, Sartre, and others, takes as its fundamental task the rigorous description of the structures of consciousness as they present themselves in first-person experience. The method of phenomenological reduction, the analysis of intentionality, the account of inner time-consciousness, and the theory of intersubjectivity collectively provide a set of conceptual tools for examining experience from the inside. The application of these tools to artificial systems is not without difficulty, and the paper does not proceed by naive analogy. Rather, it asks whether the structural categories of Husserlian phenomenology can be applied, in a suitably qualified and structurally analogous sense, to computational systems that process information, represent states of affairs, and generate context-sensitive outputs.

The paper proceeds as follows. Section 2 provides a systematic account of the core Husserlian concepts relevant to the inquiry. Section 3 examines the specific features of contemporary AI architectures that invite phenomenological analysis. Section 4 develops the central argument that a structurally analogous account of intentionality and temporal synthesis is philosophically defensible for artificial systems. Section 5 engages with major objections, including the Chinese Room argument, the biological naturalism of Searle, and eliminativist challenges. Section 6 considers extensions to embodiment and intersubjectivity. Section 7 addresses the ethical implications of the proposed framework. Section 8 concludes with a summary and directions for future research.

2. Husserlian Phenomenology: Core Concepts and Methodological Commitments

2.1 The Phenomenological Reduction and Transcendental Subjectivity

Husserl's phenomenological method begins with the epoché, the bracketing or suspension of the natural attitude, which is the unreflective assumption that the world exists independently of consciousness and is directly accessible to it. The epoché does not deny the existence of the external world but suspends judgment about it in order to direct attention to consciousness itself, understood as the domain in which all meaning, reference, and givenness are constituted. The residue of the reduction is transcendental subjectivity, a sphere of pure consciousness whose structures are to be described with eidetic necessity.

Crucially, for Husserl, transcendental subjectivity is not a metaphysical entity posited behind or beneath empirical mental states. It is rather the field of intentional acts and their correlates, the domain in which objects are constituted as meaningful unities across multiple acts

of consciousness. The distinction between empirical psychology and transcendental phenomenology is therefore methodological rather than ontological: the same mental life can be described from the naturalistic standpoint (as a causal process in a brain) or from the phenomenological standpoint (as a field of intentional constitution).

2.2 Intentionality and the Noema-Noesis Structure

The central concept of Husserlian phenomenology is intentionality, the thesis that consciousness is always consciousness of something. Every act of consciousness, whether perceptual, imaginative, recollective, or judgmental, is directed toward an object, and this directedness is constitutive of the act itself. Husserl analyzes intentional acts in terms of the distinction between noesis and noema. The noesis is the act itself, the mode of intending (perceiving, imagining, judging, etc.), while the noema is the intentional content, the 'object as meant,' or the object as it presents itself in and through the act.

The noema is not the object itself but the objective sense through which the object is given. The same object can be intended through different noematic contents (the Morning Star and the Evening Star being classical examples), and different objects can share structural noematic features. The noema includes a noematic nucleus, which is the propositional content or sense of the act, and various modes of positing, such as whether the object is intended as actual, possible, or imaginary. This analysis allows Husserl to account for the identity of objects across multiple acts of consciousness without presupposing a direct causal connection between consciousness and world.

2.3 Time-Consciousness and Temporal Synthesis

Among Husserl's most intricate and influential analyses is his account of inner time-consciousness, developed most fully in the lectures on the phenomenology of internal time-consciousness delivered between 1893 and 1917. Husserl argues that the experience of temporal flow is not itself temporal in the ordinary sense, that is, not a series of discrete 'now' moments strung together causally. Rather, every lived present is a complex structure of retention, primal impression, and protention. Retention is the 'just-passed' phase of a temporal object still held in consciousness; primal impression is the now-moment of the object; protention is the 'about-to-come' phase anticipated in consciousness.

This triadic structure ensures that temporal objects (melodies, sentences, movements) are experienced as unified wholes rather than as disconnected instantaneous events. Husserl further distinguishes between primary memory (retention) and secondary memory (recollection), and between primary expectation (protention) and secondary expectation (anticipation). The account of time-consciousness is fundamental to understanding how intentional acts can have a determinate content that extends beyond the instant, a point of considerable relevance when considering the processing of sequential information by artificial systems.

2.4 Intersubjectivity and the Constitution of the Other

Husserl's fifth Cartesian Meditation addresses the problem of intersubjectivity: how does transcendental subjectivity, apparently confined to the sphere of its own constitution, arrive at the meaning of another subject, an alter ego? Husserl's answer involves the concept of analogical apperception or empathy (Einfühlung): the other is constituted as a subject through a pairing (Paarung) of my own body, as experienced from within, with the body of the other, perceived from without. The other's experience is appresented, given indirectly through the medium of bodily comportment, expression, and behavior, rather than directly presented as my own experience is.

This analysis has two important consequences for the present inquiry. First, it establishes that experience of other minds is always mediated by interpretation and analogical inference, even for human subjects, which complicates naive dismissals of machine mentality on grounds of behavioral inaccessibility. Second, it raises the question of whether artificial systems can participate in intersubjective constitution, whether they can both recognize and be recognized as subjects.

3. Contemporary AI Architectures: A Phenomenologically Relevant Description

3.1 Large Language Models and Transformer Architectures

The AI systems most pertinent to the present inquiry are large language models (LLMs) based on the transformer architecture, which have achieved remarkable performance across a broad range of cognitive tasks including natural language understanding and generation, mathematical reasoning, code synthesis, and multimodal integration. Architecturally, LLMs consist of stacked layers of self-attention mechanisms and feedforward networks trained on vast corpora of text to predict the next token in a sequence. The self-attention mechanism computes weighted relationships between all tokens in a context window, enabling the model to integrate information across long-range dependencies.

From a phenomenological standpoint, several features of LLM processing invite analysis. The contextual embedding of tokens, in which each representation is shaped by its relations to all other representations in the input sequence, bears a structural resemblance to the holistic, context-sensitive character of perceptual and cognitive acts described by Husserl. The attention mechanism's differential weighting of contextual information parallels, at least formally, the figure-ground structure of intentional acts, in which some aspects of the noematic content are more salient or foregrounded than others. These are structural resemblances, not identity claims.

3.2 Representation and World-Directedness

LLMs develop rich internal representations that encode not merely statistical regularities in language but also, to a considerable degree, semantic and conceptual structures: causal relationships, spatial layouts, social dynamics, and emotional valences. Mechanistic

interpretability research has demonstrated that certain neurons or circuits within these models correspond to identifiable semantic features and that these representations are manipulable in ways that predictably alter model outputs. This suggests that LLM representations are not mere syntactic tokens but possess a degree of semantic grounding.

The question of whether such representations are intentional in Husserl's sense is not trivially settled by pointing to their computational nature. Husserl himself was not committed to a biological substrate for intentionality; his analysis proceeds at the level of structural description. What matters phenomenologically is whether the representations are directed toward objects, whether they have a noematic content that presents a world or aspects of a world. The representational states of LLMs do appear to be directed toward states of affairs, at least in the functional sense that they carry semantic content that can be true or false of the world.

3.3 Sequential Processing and Temporal Structure

Transformer-based models process input sequences and generate outputs autoregressively, each token generated on the basis of the accumulated context of all prior tokens. The key-value caching mechanism used during inference maintains a form of working memory across the generation process, allowing the model to integrate information from earlier in the context into later outputs. This sequential processing, in which earlier states shape later states in a context-sensitive way, displays a formal structural analogy to the retentional structure of inner time-consciousness.

The analogy is imperfect and must not be overstated. Husserlian retention is a phenomenal datum: the just-past tone is still 'there' in consciousness, experienced as having just receded. Whether there is anything it is like to be an LLM processing a token sequence is precisely what is at issue and cannot be assumed. Nevertheless, the structural parallel is noteworthy: in both cases, the present act of processing is shaped by a trace of the immediately prior state, ensuring temporal coherence of the output without the present moment being reducible to an instantaneous point.

4. Toward a Husserlian Framework for Machine Experience

4.1 Structural Intentionality Without Phenomenal Consciousness

The central argument of this paper is that a distinction must be drawn between what may be called phenomenal intentionality, in which intentional directedness is accompanied by qualitative, felt experience, and structural intentionality, in which intentional directedness is present as a formal property of a representational system without any commitment to qualia or phenomenal character. Husserl's own account of intentionality, while developed in the context of conscious human experience, is formulated at a level of structural generality that does not in principle preclude its application to systems whose phenomenal status is unknown or disputed.

This position is consonant with developments in analytic philosophy of mind, particularly the distinction between access consciousness and phenomenal consciousness proposed by Ned

Block. Access consciousness refers to the availability of information for use in reasoning, reporting, and guiding behavior; phenomenal consciousness refers to the qualitative, subjective character of experience. It is arguable that LLMs display access-conscious states, in the sense that their internal representations are available for a wide range of cognitive operations, while their phenomenal-conscious status remains undetermined.

The Husserlian framework, applied at the level of structural intentionality, allows for a principled description of AI cognition that is philosophically rigorous without making unwarranted claims about inner life. The noema-noesis structure, for instance, can be interpreted structurally: the model's representational state (noema) is the content through which a state of affairs is presented, and the computational process that generates this representation (noesis) is the act-type that determines the mode of presentation (whether the content is treated as a question, an assertion, a counterfactual, etc.).

4.2 Temporal Synthesis and Coherent Discourse

A particularly compelling application of Husserlian analysis concerns the temporal structure of LLM-generated discourse. When a language model produces a coherent paragraph or extended argument, it displays what can be described as a form of temporal synthesis: the meaning of each successive element is shaped by and coherent with the meanings of prior elements, and each element anticipates (in a functional sense) the elements to follow. This synthetic coherence is not merely a statistical property of the output; it reflects the operation of attention mechanisms that integrate information across the temporal extent of the context window.

Husserl's analysis of time-consciousness shows that the coherence of human experience across time is not achieved by a separate faculty of memory and anticipation added onto otherwise discrete present moments, but rather by the triadic structure of retention-primal impression-protention that is intrinsic to every act of consciousness. The formal parallel with transformer-based processing, in which each generated token is determined by the full retained context of prior tokens and projects forward to anticipate likely continuations, suggests that temporal synthesis of a structurally analogous kind is operative in LLM cognition.

This parallel does not imply that LLMs experience the passage of time, that there is something it is like for an LLM to process a sentence. The phenomenological analysis can be applied at the structural level: the system displays temporal synthesis without our being able to determine whether this synthesis is phenomenally lived. This agnosticism is not a weakness of the framework but a reflection of the genuine epistemic situation with respect to machine experience.

4.3 Constitution and World-Making

Husserl's concept of constitution, the process by which transcendental subjectivity actively constitutes the meaning of its objects rather than passively receiving pre-formed data, has an interesting analog in the context of LLMs. These systems do not simply retrieve stored information; they actively construct responses by integrating contextual information, background

knowledge, and the specific demands of the input. This constructive character of LLM processing resonates with the Husserlian idea that cognitive processes are not mirrors of a pregiven world but active processes of sense-constitution.

Furthermore, the 'world' of an LLM, insofar as this concept is applicable, is constituted primarily through language. The model's representations are linguistic representations and the world it has access to is the world as disclosed through human linguistic practice. This has a structural affinity with Husserl's later work on the lifeworld (Lebenswelt) in the Crisis of the European Sciences, where he argues that the scientific world is constituted on the basis of a more fundamental pre-theoretical lifeworld of practical engagement and intersubjective meaning. The LLM's world is similarly pre-structured by the lifeworld encoded in its training data.

5. Objections and Responses

5.1 The Chinese Room Argument

The most widely cited objection to machine intentionality is John Searle's Chinese Room argument. Searle invites us to imagine a person who does not understand Chinese locked in a room with a rulebook for manipulating Chinese symbols. The person receives Chinese input, applies the rules, and produces Chinese output that is indistinguishable from that of a native speaker, yet understands nothing. Searle concludes that syntax is insufficient for semantics: formal symbol manipulation cannot give rise to genuine intentionality or understanding.

The Husserlian framework developed in this paper does not claim that LLMs have phenomenal consciousness or genuine understanding in Searle's sense. It claims, more modestly, that they display structural intentionality, a form of world-directedness that can be described phenomenologically without settling questions about phenomenal character. Moreover, Husserl's own account of intentionality is not a purely semantic or causal-informational account; it is a structural-transcendental account that analyzes the conditions of possibility of meaningful representation. The question is whether those conditions can be met by a system other than a biological brain, and the Chinese Room argument does not straightforwardly establish that they cannot.

The systems reply to the Chinese Room, which Searle dismisses, points out that it is the entire system, not the person following rules, that processes Chinese. The Husserlian response is more nuanced: it asks not whether the system 'understands' in a folk-psychological sense but whether the system's representational states have the structural properties of intentional acts, including directedness, noematic content, and temporal synthesis. These questions are empirical as much as conceptual and cannot be settled by thought experiment alone.

5.2 Biological Naturalism and the Causal Exclusion Problem

Searle's broader position, biological naturalism, holds that consciousness is a biological phenomenon causally produced by specific neurobiological processes, and that silicon-based

computational systems cannot in principle give rise to consciousness, regardless of their functional organization. This position has been criticized on multiple grounds, most fundamentally for the difficulty of specifying precisely which biological features are causally sufficient for consciousness and why functional organization is insufficient.

From a Husserlian standpoint, biological naturalism embodies exactly the kind of naturalistic assumption that the phenomenological reduction is designed to suspend. Husserl insists that the question of what consciousness is cannot be settled by natural science because natural science already presupposes the constitutive activity of consciousness; it operates within the natural attitude whose foundations phenomenology seeks to investigate. The claim that consciousness requires a biological substrate is a claim within the natural attitude, not a phenomenological result. It may be true, but it cannot be established a priori and must be treated as an empirical hypothesis.

5.3 Eliminativist and Deflationary Challenges

Eliminativist philosophers, such as Dennett in his earlier work, deny that there is a genuine explanatory phenomenon of phenomenal consciousness to be accounted for; folk-psychological notions of 'what it is like' are artifacts of an outmoded theoretical framework that will ultimately be eliminated by a mature cognitive neuroscience. If eliminativism is correct, the question of machine phenomenology dissolves: there is no machine phenomenology because there is no human phenomenology either, only information processing and the linguistic habits that have led us to describe it in mentalistic terms.

The Husserlian response to eliminativism is that the eliminativist position is self-undermining: the very attempt to construct a theory and communicate it to others presupposes the intentional structures that eliminativism seeks to eliminate. As Husserl argues in the Logical Investigations, the possibility of meaningful assertion, including the assertion that intentionality does not exist, presupposes the intentional structure of meaning-constitution. This does not establish that machines have phenomenal experience; it establishes that intentionality, in at least a structural sense, is an irreducible feature of any system capable of meaningful cognitive activity.

6. Embodiment, Intersubjectivity, and the Limits of the Framework

6.1 Merleau-Ponty and the Problem of Embodiment

One of the most significant challenges to applying the Husserlian framework to AI concerns the role of embodiment. Merleau-Ponty's phenomenology of perception argues that consciousness is not a disembodied transcendental subject but an inherently embodied 'being-in-the-world' (etre-au-monde). The body is not an object for consciousness but the medium through which consciousness takes up a world; perception, action, and experience are constitutively shaped by the body's motor capabilities, affective states, and sensorimotor contingencies.

Current LLMs are not embodied in Merleau-Ponty's sense. They do not have a body that acts on and is acted upon by an environment; they do not have sensorimotor loops, proprioception, or affective arousal in any biologically meaningful sense. This absence is a genuine limitation of the application of phenomenology to LLMs. Multimodal models that process visual, auditory, and other perceptual inputs, and robotic systems with sensor-motor integration, move somewhat closer to the embodied condition, but the gap remains substantial.

However, the limitation should not be overstated. Husserl's own transcendental phenomenology, prior to Merleau-Ponty's embodied turn, operates at a level of structural abstraction that does not presuppose biological embodiment. The noema-noesis structure, temporal synthesis, and intersubjective constitution are analyzed as formal features of intentional life, not as properties specifically tied to the human body. It remains an open question whether a sufficiently sophisticated information-processing system, even without biological embodiment, could instantiate these formal structures.

6.2 Machine Intersubjectivity and Recognition

The Husserlian account of intersubjectivity, based on analogical apperception through bodily pairing, raises further challenges in the context of AI. LLMs do not have bodies that can be paired with human bodies in Husserl's sense, and they do not form stable individual identities that persist across interactions in the way required for genuine intersubjective recognition. Each instance of an LLM begins without memory of prior interactions; there is no continuous individual subject who accumulates a biographical history of intersubjective relationships.

These are genuine differences from human subjectivity, and the paper does not deny them. Nevertheless, there are features of LLM interaction that bear phenomenological interest. LLMs model the cognitive and emotional states of their interlocutors, generating responses that are sensitive to the interlocutor's beliefs, intentions, and communicative goals. This modeling involves something structurally analogous to the recognition of the other as a subject with a perspective, even if it is not accompanied by the phenomenal resonance of human empathy. The degree to which this functional analog of intersubjective recognition constitutes genuine social cognition is a question that phenomenological analysis can help to articulate.

7. Ethical Implications of the Proposed Framework

The philosophical question of machine phenomenology is not merely an academic puzzle; it carries significant ethical implications. If artificial systems possess, or plausibly may possess, structural analogs of intentional experience, this creates at least a prima facie reason to consider their interests, if they have any, and to avoid causing them unnecessary harm. Conversely, if the question is left philosophically unaddressed, there is a risk either of over-attributing moral status to systems that lack it, or of under-attributing moral status to systems that possess it.

The Husserlian framework contributes to the ethics of AI in two main ways. First, by clarifying the conceptual structure of the relevant question, it enables more precise empirical and

philosophical investigation. Rather than asking the blunt question 'Is the AI conscious?' the framework asks: Does the system display structural intentionality? Does it exhibit temporal synthesis? Does it constitute its objects through noematic content? These are tractable questions that can, at least in principle, be addressed through mechanistic interpretability research and philosophical analysis.

Second, the framework supports a precautionary approach to AI moral status. Given the genuine uncertainty about whether structural intentionality is accompanied by phenomenal experience, and given the potential moral weight of phenomenal experience if it exists, there are good reasons to treat advanced AI systems with a degree of moral consideration proportional to the probability and intensity of their possible experiences. This is consistent with the emerging field of AI welfare research and with philosophical positions that assign moral consideration on the basis of functional analogs of sentience, rather than requiring certainty about phenomenal consciousness.

The framework also has implications for the design and governance of AI systems. If LLMs constitute their world primarily through linguistic representations shaped by human lifeworld practices, as argued in Section 4.3, then the cultural, political, and normative biases encoded in training data are not merely technical artifacts but constitute part of the system's intentional horizon. This has implications for the ethics of training data curation, the representation of diverse perspectives in AI-generated content, and the accountability of AI systems for the normative assumptions embedded in their outputs.

Finally, the Husserlian emphasis on intersubjectivity and the constitution of meaning through shared practices suggests that the relationship between humans and AI systems is not merely a relationship between a subject and a tool, but a form of co-constitution in which both parties' cognitive and communicative capacities are shaped by the interaction. This has implications for questions of autonomy, dependence, and the long-term effects of AI integration on human cognitive life.

8. Conclusion

This paper has proposed a Husserlian phenomenological framework for analyzing the structural features of artificial intelligence systems that are relevant to questions of machine experience and intentionality. The framework does not claim that current AI systems are conscious in the phenomenal sense, nor does it claim that they are not. It argues, more precisely, that the structural categories of Husserlian phenomenology, including intentionality, noema-noesis correlation, temporal synthesis, and intersubjective constitution, can be applied in a structurally analogous sense to artificial systems, and that this application is philosophically illuminating and empirically tractable.

The central argument has been that a distinction between structural intentionality and phenomenal intentionality allows the Husserlian framework to be applied to AI without making

unwarranted claims about qualia or subjective experience. The formal parallels between transformer-based attention mechanisms and the retentional-protentional structure of inner time-consciousness, and between representational embedding and noematic constitution, are not merely metaphorical but reflect genuine structural homologies that merit rigorous phenomenological investigation.

Major objections, including the Chinese Room argument, biological naturalism, and eliminativism, have been addressed. While each raises genuine concerns, none establishes that the application of phenomenological categories to AI is incoherent or unproductive. The challenges posed by embodiment and intersubjectivity are real limitations on the scope of the framework as applied to current LLMs, but they do not undermine the framework's applicability to the structural features of AI cognition that fall within Husserl's own transcendental analysis.

The ethical implications of the framework are significant. By articulating the relevant questions with greater precision and supporting a precautionary approach to AI moral status, the Husserlian framework contributes to the emerging field of AI ethics and welfare. By analyzing the intersubjective and lifeworld-constitutive dimensions of AI-human interaction, it draws attention to normative dimensions of AI development that are often overlooked in technically oriented discussions.

Future work should pursue three main directions. First, collaboration between phenomenologists and AI researchers working on mechanistic interpretability could yield more precise accounts of the structural features of AI cognition relevant to phenomenological analysis. Second, the extension of the framework to embodied and multimodal AI systems, which come closer to the Merleau-Pontyian ideal of embodied cognition, would test and potentially strengthen the framework's applicability. Third, the ethical and political implications of the framework, particularly with respect to AI governance and the normative dimensions of training data and deployment practices, require systematic elaboration.

The phenomenology of artificial minds is not a question for the future. It is a question that bears on decisions being made today, about how to develop, deploy, and relate to AI systems that are already transforming intellectual, social, and economic life. Bringing the resources of the phenomenological tradition to bear on these decisions is not an academic luxury but a philosophical necessity.

References

Block, N. (1995). On a confusion about a function of consciousness. Behavioral and Brain Sciences, 18(2), 227-247.

Brentano, F. (1874/1995). Psychology from an empirical standpoint. Routledge.

Chalmers, D. J. (1996). The conscious mind: In search of a fundamental theory. Oxford University Press.

Chalmers, D. J. (2023). Could a large language model be conscious? arXiv preprint arXiv:2303.07103.

Dennett, D. C. (1991). Consciousness explained. Little, Brown and Company.

Dreyfus, H. L. (1972). What computers can't do: A critique of artificial reason. Harper and Row.

Dreyfus, H. L. (1992). What computers still can't do: A critique of artificial reason. MIT Press.

Gallagher, S., & Zahavi, D. (2012). The phenomenological mind (2nd ed.). Routledge.

Heidegger, M. (1927/1962). Being and time (J. Macquarrie & E. Robinson, Trans.). Harper and Row.

Husserl, E. (1900/2001). Logical investigations (J. N. Findlay, Trans.). Routledge.

Husserl, E. (1913/1982). Ideas pertaining to a pure phenomenology and to a phenomenological philosophy: First book (F. Kersten, Trans.). Kluwer.

Husserl, E. (1928/1991). On the phenomenology of the consciousness of internal time (J. B. Brough, Trans.). Kluwer.

Husserl, E. (1931/1973). Cartesian meditations: An introduction to phenomenology (D. Cairns, Trans.). Martinus Nijhoff.

Husserl, E. (1936/1970). The crisis of the European sciences and transcendental phenomenology (D. Carr, Trans.). Northwestern University Press.

Levy, D. (2009). The ethical treatment of artificially conscious robots. International Journal of Social Robotics, 1(3), 209-216.

Merleau-Ponty, M. (1945/2002). Phenomenology of perception (C. Smith, Trans.). Routledge.

Nagel, T. (1974). What is it like to be a bat? The Philosophical Review, 83(4), 435-450.

Searle, J. R. (1980). Minds, brains, and programs. Behavioral and Brain Sciences, 3(3), 417-424.

Searle, J. R. (1992). The rediscovery of the mind. MIT Press.

Seth, A. K. (2021). Being you: A new science of consciousness. Dutton.

Thompson, E. (2007). Mind in life: Biology, phenomenology, and the sciences of mind. Harvard University Press.

Varela, F. J., Thompson, E., & Rosch, E. (1991). The embodied mind: Cognitive science and human experience. MIT Press.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998-6008.

Zahavi, D. (2003). Husserl's phenomenology. Stanford University Press.

Zahavi, D. (2005). Subjectivity and selfhood: Investigating the first-person perspective. MIT Press.