
Philosophical inquiry into intelligence, cognition, consciousness and the nature of artificial minds

Toyesha Thakur

**Asst. Professor
Department of English**

Gyan Ganga Institute of Technology and Sciences Jabalpur

Abstract:

This paper conducts philosophical inquiry into the essence of intelligence, cognition and consciousness, probing the enigmatic nature the paper undertakes a philosophical inquiry into the intertwined concept of intelligence, cognition, consciousness, and the emerging possibility of artificial minds.

Drawing from classical and contemporary philosophy of mind, the study interrogates whether artificial systems can meaningfully possess mental states or merely simulates them. The investigation situates modern artificial intelligence within a broader intellectual lineage extending from René Descartes dualism and Immanuel Kant's transcendental account of cognition to contemporary debates such as Jhon Searle's Chinese Room argument and David Chalmers' formulation of the "hard problem" of consciousness by synthesizing perspectives from cognitive science, phenomenology, and computational theory, the paper explores whether intelligence should be understood as functional problem solving capacity, embodied cognition, or a phenomenon inseparable from cognitive sciences phenomenology and computational theory. The study views recent advances in machine learning as provocations for philosophical reflection rather than final answers it discusses a duel between mind and body therefore understanding AI is not just technical issue but deep philosophical inquiry into nature of mind.

Keywords: Artificial intelligence, consciousness, cognition, artificial minds, qualia, mind body dualism.

Introduction

The aim of this paper is to review the nature of intelligence, cognition and consciousness has long been central to philosophical inquiry from ancient metaphysical debates to modern scientific investigation, the question of what constitutes a mind has remained persistently unresolved. In recent decades, the rapid advancement of artificial intelligence [AI] has reintroduced these questions with renewed urgency and complexity. Artificial systems can perform tasks that were once considered uniquely human learning from data, generating language, recognizing patterns, and even simulating emotional responses. These capabilities raise an essential philosophical question: Do such system genuinely think, or do they merely simulate thoughts?

While technological discourse often equates intelligence with performances, this review argues that such an approach is conceptually insufficient. Intelligence, cognition and consciousness are not merely operational categories; they are deeply philosophical constructs that require careful

examination. This seeks to explore whether artificial system can meaningfully be said to possess mental states. It situates contemporary AI within a broader philosophical tradition, examining both classical frameworks and modern debates. Rather than offering definitive answers, the paper aims to clarify the conceptual terrain and identify the limits of current understanding. A closer examination of contemporary discourse on artificial intelligence reveals a persistent conceptual imbalance. Much of the current scholarship remains preoccupied with functional performance, often measuring the success of artificial system by their ability to replicate or approximate human behavior. While such an approach has undoubtedly advance technological development, it tends to obscure a more fundamental philosophical concern namely the distinction between simulation and genuine experience. The capacity of machines to intimate cognitive processes does not necessarily entail that they participate in those processes in any meaningful or experiential sense. The ambiguity becomes particularly evident when one considers an ontological status of artificial minds. It is one thing to describe a system as behaving intelligently, but quite another to claim that it possesses mental states in the same way humans beings do the existing literature in many instances, appears to move too quickly from observable performance to assumptions about inner states, leaving insufficient space for critical reflection questions concerning whether artificial system can truly think, understand, or experience remain underexplored, often overshadow by the excitement surrounding technological capability.

At the same time, there is a noticeable fragmentation in the way different disciplines engage with the issues. Insights from phenomenology classical epistemology, and consciousness studies are frequently treated as peripheral, rather than as integral to the discussion. As a result, the conversation risks becoming narrowly technical detached from the broader philosophical traditions that have long grappled with the nature of mind and experience. It is with in this context that the present inquiry positions itself. Rather than taking the achievements of artificial intelligence at face value, it seeks to revisit someone of deeper assumptions that underlie our understanding of intelligence and consciousness by drawing connections between classical philosophical thought and contemporary debates, the discussion attempts to restore a sense of conceptual depth to the field. One that acknowledges not only what artificial systems can do; but also, what it might mean for them to be.

Methodology

The exploration of the ideas discussed in this paper was guided by a focus and selective engagement with relevant literature the process began with a careful search across widely used academic repositories, particularly Google scholar and Semantic scholar where key term such as “artificial intelligence and consciousness”, Philosophy of mind,” and “cognition and AI” were employed to identify foundational and contemporary works. This initial step allowed for the identification of influential texts spanning both classical philosophy and modern interdisciplinary research.

In addition to database searches, further material was gathered by tracing references within the selected works themselves. This method proved especially useful in identifying significant contributions that may not always appear in keyword-based searches but are central to ongoing

scholarly conversations. Through this process, the study gradually expands its scope while maintaining a clear focus on the philosophical dimensions of artificial intelligence.

The collected literature was approached thematically, with attention given to the major areas of inquiry relevant to the study namely intelligence, cognition, consciousness, and the nature of artificial minds. Within these broad themes, an effort was made to consider the development of ideas over time, beginning with classical philosophical positions and moving towards contemporary debates. This progression helps illustrate not only the continuity of certain questions but also the ways in which they have been reframed in response to technological advancement at certain points closely related works are discussed together, particularly where they represent an extension or refinement of a shared line of inquiry. This allows for a more coherent understanding of how specific arguments have evolved, rather than treating each contribution in isolation. The intention throughout has been to present the material in a manner that reflects both its intellectual depth and its interconnected nature. Overall, the approach remains interpretative rather than empirical, emphasizing careful reading, comparison and synthesis of ideas. By engaging with the literature in this way, the study seeks to develop a more nuanced understanding of artificial intelligence not simply as a concept that continues to challenge and reshape philosophical thinking about the mind.

Conceptual Foundations, Philosophical Perspectives, and the Question of Artificial Minds

Any attempts to understand artificial intelligence in a meaningful way must begin with a reconsideration of the conceptual foundations that underlie discussions of intelligence, cognition, and consciousness. These terms are often used with a certain ease in contemporary discourse, yet their meanings remain far from settled. Intelligence for instance, is frequently described in terms of problem solving ability and adaptability. By this measure, many artificial systems appear undeniably intelligent; They can process vast amounts of data, identify patterns, and generate responses with impressive speed and accuracy. And yet, there is something slightly unsettling about reducing intelligence to efficiency alone. Human intelligence does not operate merely as a tool for solving problems. It involves reflection, hesitation, emotional sensitivity, and awareness of context that is not easily translated into computational terms. One may begin to wonder whether the current enthusiasm around artificial intelligence risks narrowing the very idea of intelligence to what can be measured and optimized?

Similar tension arises in discussion of cognition. Traditional models, particularly those influenced by early cognitive science tend to equate cognition with information processing suggesting that the mind functions in a manner analogous to a computational system. This framework has been useful; No doubt, in advancing both psychological theory and technological innovation. However it does not fully capture the richness of human cognitive life. Cognition is not something that occurs in isolation within the confines of the brain; It is shaped by bodily experience by interaction with the environment, And by the social and cultural worlds in which individuals are embedded. To think is not simply to process information but to engage with a world that is already meaningful. This raises a question that cannot be easily dismissed: if cognition is inherently embodied and situated,

can truly be replicated in systems that lack this form of engagement? It is possible that what we describe as artificial cognition is, in fact, only partial reconstruction of a far more complex phenomenon the difficulty depends when one turns to consciousness unlike intelligence or cognition, consciousness resists straight forward definition. It is often described in terms of subjective experience the sense of “what it is like” to exist, to perceive, to feel of subjective experience has long been a source of philosophical and ease. When applied to artificial systems, the problem becomes even more pronounced. A machine may behave in ways that resemble the conscious action, but it does not experience anything at all? Or is consciousness something that cannot be captured by functional equivalence? These questions remain open, and perhaps they resist closure for reasons that have less to do with technology and more to do with the limits of our own understanding.

It is at this point that classical philosophical perspectives begin to re-enter the discussion in a meaningful way. The distinction between mind and body for example, has its roots in dualist traditions that treat the mental as fundamentally different from the physical. From such a standpoint, the idea of artificial consciousness appears problematic since machines are entirely physical systems. If consciousness is not reducible to material processes n it cannot simply be engineered. Yet dualism itself is not without its difficulties. It raises question about how to fundamentally different kind of substances might interact, and these questions have never been fully resolved. Still, the persistence of dualist institutions suggest that there is something about consciousness that resist straight forward physical explanation.

In contrast the other philosophical approaches emphasise the active role of the mind in shaping experience rather than viewing cognition as a passive reception of information, these perspectives suggest that the mind organises and interprets sensory inputs according to certain underlying structures. This idea complicates the possibility of artificial cognition. If understanding depends not merely on processing data but on structuring experience in specific ways, then replicating cognition would require more than computational sophistication. It would involve recreating the very conditions under which experience becomes meaningful, a task that is far from straight forward.

Contemporary debates continue to grapple with these issues often through though experiments that expose the limits of computational models of mine. One such example in whites us to imagine a system that can produce perfectly appropriate responses in a language it does not understand, simply by following a set of rules. The point here is not to deny that such a system can function effectively, but to question whether functionality alone constitutes understanding. There is a difference, however subtle, between manipulating symbols and grasping their meaning. And it is precisely this difference that complicates claims about artificial intelligence the more closely machines approximate human behaviour, the more pressing it becomes to ask what, if anything, lies behind that behaviour.

Closely related to this the enduring problem of consciousness itself. Even if one were to grant that machines could replicate all observable aspects of human cognition there would still remain the question of subjective experience. This is sometimes described as distinction between the “easy” Problems of explaining behaviour and the “hard” problem of explaining why and how experiences arises at all. The persistence of this problem suggests that something essential may be missing

from purely fictional accounts of mind. It is not simply a matter of adding more complexity to our models; The difficulty seems to lie Add a deeper conception level.

Against this backdrop artificial intelligence can be seen not so much as providing answers but as intensifying existing questions. Theories that define mental states in terms of their functional roles offer one way to interpreting AI systems, suggesting that if a machine behaves as though it understands, it may be said to understand. Yet this view remains contested, particularly when it comes to consciousness. There is a lingering sense that something important is being overlooked when mental life is reduced to functional description alone. At the same time, approaches that emphasise the role of embodiment draw attention to the ways in which human cognition is shaped by physical presence and interaction. Most artificial systems, operating in abstract computational environments lack this dimension. Whether this absence represents temporary limitation or a fundamental barrier is a question that remains unresolved. What becomes increasingly clear is that artificial intelligence functions as a kind of philosophical provocation. It forces us to revisit assumptions that might otherwise have gone unexamined. The distinction between appearance and reality, between simulation and experience, becomes more difficult to ignore in this sense AI does not simply extend human capabilities; It changes our understanding of what those capabilities mean.

This challenge is perhaps most evident discussion of subjective experience. The notion of failure the felt quality of experience highlights the difficulty of translating inner life into objective terms. One can describe the mechanisms involved in perceptions, but such description do not capture the experience itself. This raises a familiar but still unresolved question polling can a system that lacks subjective experience ever truly be set to have a mind? Or does the absence of qualia March fundamental limit of artificial intelligence? It is not easy to answer these questions, and it may be that they cannot be answered within the current conceptual framework.

The debate ultimately circles back to the broader issue of the relationship between mind and body. These consciousness emerges from physical processes or exists independently of them remains a matter of ongoing debate. Some approaches suggest that sufficiently complex systems may give rise to consciousness, while others remain sceptical of such claim. What is striking, however, is the extent to which these debates persist despite advances in both philosophy and science. The problem does not seem to diminish with increasing knowledge; If anything, it becomes more intricate.

The light of these considerations, it becomes more necessary to reflect not only on what artificial systems can do, but on what it would mean for them to be subjects of experience. The shift in perspective has important ethical implications for stop if artificial systems were ever to possess consciousness, questions about their moral status board become unavoidable. Would such system deserve rights? How would one determine whether they are capable of suffering Or experience? And what the responsibilities would fall upon those who create and deploy them? These are not merely speculative concerns. Even in the absence of genuine artificial consciousness, the tendency to attribute human like qualities to machines can shape how they are used and understood.

At the same time, one might hesitate before drawing form conclusions. It is possible that current debates are constrained by the limitations of existing concepts and that future developments both philosophical and technological may require a rethinking of the term themselves. For now, however, the study of artificial minds remain an open field, marked as much by uncertainty as by progress. What seems clear in that the questions raised by artificial intelligence cannot be resolved

through technical means alone they demand a sustained engagement with the philosophical issue that have long surrounded the nature of mine.

The significance of artificial intelligence may lie less in what it achieves than in what it reveals by confronting us with systems that blur the line between thinking and acting, between understanding and simulation, it compels us to examine assumptions that might otherwise remain implicate and in doing so, it reminds us that the problem of mind is not merely a problem to be solved but a question to be continually explorer.

Future Directions: Integrating Philosophical inquiry with Technological Development

The direction of future research in the study of artificial intelligence and the nature of mind will likely depend on the extent of which philosophical reflection and technological development are allowed to inform one another for some time now, these two domains have advanced along parallel tracks forma occasionally intersecting but rarely sustaining a deep engagement. Technical progress continues at a remarkable pace conceptual questions that accompany it often remain under explored or are treated as secondary concerns it seems increasingly necessary, however, to resist this separation. If artificial intelligence is not only a tool but also a challenge to our understanding of intelligence and consciousness, then future inquiry must be shaped by a more integrated approach one that treats philosophical insight not as an afterthought but as a guiding framework.

One possible direction lies in rethinking how intelligence itself is modelled within artificial systems. Much of current AI research s grounded in optimization systems are designed to perform tasks efficiently, often by identifying patterns within large data sets. While this approach has produced impressive results, it may also limit the scope of what forms of reflective or self-referential processing not merely as technical feature but as attempts to approximate the layered nature of human thought. This does not necessarily imply that machines can become conscious in the human sense, but it does open a space for examining weather aspects of reflective awareness can meaningfully represented with artificial architectures. One might ask, for instance, whether a system can be designed not only to respond to inputs but to evaluate its own responses in ways that go beyond statistical adjustment.

Closely related to this is the question of embodiment, which has gained increasing attention in both philosophy and cognitive signs. If cognition is not confined to abstract computation but is shaped by bodily interaction with the world, then future research in artificial intelligence made need to move beyond disembodiment models. There are already emerging efforts in robotics and interactive systems that attempt to situate intelligence with physical environments, allowing machines to learn through movement, perception, and engagement. Yet these developments remain in their early stages, and their philosophical implications are far from fully understood. It may be worthwhile to consider whether embodiment is merely an enhancement of cognitive systems or whether it is a necessary condition for certain forms of understanding full stop this distinction could have significant consequences for how artificial minds Are conceptualised.

Another area that invites further exploration concerns the nature of subjective experience. While it may be seen premature to speak of artificial consciousness, the question itself continues to

provoke reflection. Some researchers have begun to investigate whether there are minimal conditions under which a system might be said to possess a form of experience however rudimentary. These inquiries often draw on interdisciplinary perspectives combining insights from neuroscience philosophy and Computational theory. What is interesting about such work is not that it offers definitive answers, but that it challenges the assumption that consciousness is either fully present or entirely absent. Instead, it suggest the possibility of graduation or thresholds, raising the question of whether artificial systems might occupy some intermediate position. This is admittedly a speculative line of thought, but it serves to highlight the limitations of binary distinctions that have long structured the debate.

At the same time future research might benefit from closer examination of language and meaning within artificial systems. The ability of AI models to generate coherent and contextually appropriate language has often been taken as evidence of advanced cognition. Yet there remains an unresolved tension between linguistic performance and semantic understanding. It is one thing for a system to produce meaningful sentences; it is another for those sentences to be grounded in experience or intention. Philosophical investigations into the nature of meaning particularly those that emphasize use, context, and interpretation could offer valuable insights here. By bringing these perspectives into conversation with computational models, researches may be able to develop more nuanced accounts of what it means for a system to “understand” language.

There is also scope for innovative work at the intersection of artificial intelligence and phenomenology. While phenomenology has traditionally focused on human experience it methods could be updated to examine how artificial systems process and organise information. Such an approach would not assume that machines have experiences in the human sense, but it would seek to describe their operations in a way that highlights structure rather than function alone. This shift in perspective might reveal patterns that are otherwise overlooked, particularly in systems that rely on complex, layered architectures. It may also encourage a more careful consideration of the differences between human and artificial forms of cognition rather than treating them as directly comparable.

Ethical enquiry, too, is likely to play an increasingly central role in future research. As artificial system become more integrated to everyday life questions about responsibility, agency, consciousness, the ways in which human relate to intelligent systems can have significant consequences. There is a growing need to examine not only what the systems can do but how they are perceived and how they shape human behaviour . For instance, if individuals begin to attribute emotional or cognitive states to machines, this may influence decision making in ways that are not immediately apparent. Understanding these dynamics requires a combination of philosophical analysis and empirical study, suggesting once again the importance of interdisciplinary collaboration.

In addition to these areas for my future work might also consider the historical dimension of the problem. These questions raised by artificial intelligence are not entirely new; They echo earlier debates about the nature of mine, knowledge, and experience. By revisiting these traditions with fresh attention, researchers may find resources for addressing contemporary challenges at the same time, it is important not to assume that past frameworks can simply be applied without modification. The emergence of artificial intelligence introduce new variables that may require rethinking re-established concepts. This tension between continuity and change is likely to remain defining feature of field.

One cannot ignore the role of methodological innovation in shaping future inquiry. Traditional philosophical methods, while valuable, may need to be supplemented by new forms of analysis that are responsive to technological developments this could involve collaborative work between philosophers, computer scientists, and cognitive researchers, as well as the development of tools that allow for more detailed examination of artificial systems. For example, interpretability research in AI aimed at understanding how models arrive in their outputs could be enriched by philosophical perspectives on explanation and understanding. Conversely, philosophical discussions of mind might benefit from greater engagement with the technical realities of contemporary AI.

It is also worth considering the possibility that future research may challenge some of the assumptions that currently structure the debate. The distinction between human and artificial intelligence, for instance may become less clear as systems become more integrated into human environments. Hybrid forms of cognition, in which human and machine processes are intertwined, are already beginning to emerge. These developments raise questions about where the boundaries of the mind should be drawn whether such boundaries remain meaningful. Rather than treating artificial intelligence as something entirely separate from human cognition, it may be more productive to explore the ways in which the two interact and coevolve.

At the same time, one might approach these developments with a degree of caution there is a tendency particularly in discussions of emerging technology, to move quickly from possibility to assumption. The idea that machine might not one day possess consciousness or agency can be compelling, but it is important to distinguish between speculative projection and grounded analysis future research will need to navigate this tension carefully, remaining open to new possibilities while also maintaining a critical perspective, this balance is not easy to achieve, but it is essential if the field is too developed in a thoughtful and responsible way.

In reflecting on these various directions it becomes apparent that the study of artificial mind is not problem that can be resolved within a single discipline or through a single approach. It requires a willingness to engage with complexity, to question assumptions and to remain attentive to the limits of current understanding. Perhaps more importantly, it calls for a recognition that the questions raised by artificial intelligence are, at their core, questions about ourselves. What does it mean to think, to understand, to experience? These are not questions that can be answered but and once for all, but they are questions that gain new significance in the presence of machines that appear to share some of these capacities.

Future research, then, maybe less about arriving at definitive conclusions and more about sustaining a certain kind of inquiry but it may involve revisiting familiar questions from new angles, drawing connections across disciplines, and remaining open to the possibility that some aspects of mind will resist complete explanation. In this sense, the integration of philosophical enquiry with technological development is not simply a methodological recommendation; It is a recognition of the kind of problem we are dealing with. Artificial intelligence does not merely extend our capabilities it challenges are understanding of what those capabilities are and it is in responding to this challenge rather than resolving it, that the most of meaningful contribution of future research are likely to emerge.

Conclusion

The inquiry undertaken attempted to revisit some of the most enduring questions surrounding intelligence, cognition, and consciousness, now reframed within the context of artificial intelligence. What initially appears to be a technological discussion gradually reveals itself as something more intricate, drawing us back into philosophical territory that has long resisted clear resolution. The development of artificial systems capable of learning, responding, and adapting has undoubtedly expanded the scope of what machines can do. Yet, as the discussion throughout this paper has suggested, the question of what such systems are or what they might meaningfully be set to possess remains far less certain.

At several points, the analysis has returned to a distinction that is easy to overlook but difficult to dismiss: the difference between simulation and experience. Artificial systems may replicate patterns of behaviour that resemble human intelligence, and in many contexts, this resemblance is sufficient to practical purpose. However, when examined more closely, the gap between outward performance and inner experience become difficult to ignore. It is one thing to generate language, to recognise patterns, or to make decisions based on data; It is quite another to understand, to feel, or to be aware. Whether this distinction marks a permanent boundary or a temporary limitation is not a question that can be answered definitively at present. Still, it is a distinction that continues to shape the philosophical landscape in which artificial intelligence is situated.

Discussion of cognition has similarly pointed toward the limits of purely computational models. While the idea of the mind as an information- processing system has provided a useful framework, it does not fully capture the embodied and situated nature of human thoughts. Cognition, as it is lived and experienced, is inseparable from context from the physical environment, from social interaction, and from the continuity of personal experience. In this light, the possibility that artificial systems might replicate cognition in its entirety begins to appear less straight forward. It may be that current models capture certain aspects of cognitive activities while leaving others largely untouched. This does not diminish their significance, but it does suggest the need for more cautious interpretation of what they achieve.

Perhaps the most persistent challenge, however, rises in relation of consciousness. The difficulty of explaining subjective experience of accounting for the “what it is like” dimension of existence remains unresolved despite advances in both philosophy and science. Artificial intelligence, rather than solving this problem, seems to bring it into sharper focus. The more closely machines approximate human behaviour, the more pressing the question becomes: is there anything it is like to be such a system? Or are we encountering a form of complexity that, while impressive, remains entirely devoid of inner life? These questions do not admit easy answers, and it is possible that they expose limits not only of technology but of human understanding itself.

In engaging with classical and contemporary philosophical perspectives, this paper has sought to show that the debate surrounding artificial minds is not entirely new, even if its current form relationship, of the nature of knowledge and of the structure of experience continue to inform present concerns, albeit in ways that are sometimes implicate. At the same time, artificial intelligence introduces new dimensions that complicate these inherited frameworks. The challenge, then, is not simply to apply existing theories to new phenomena, but to allow both to inform one another in a more reflective and open ended manner.

It is within this broader context that the role of future research becomes particularly significant. If there is one conclusion that emerges with some clarity, it is that the study of artificial intelligence

cannot remain confined to technical considerations alone. The questions it raises about intelligence, understanding, experience are not merely empirical but conceptual. They require a mode of inquiry that is attentive to meaning as well as mechanism. Full research, therefore, must find ways to integrate philosophical reflections with technological development not as separate endeavours but as mutually informing processes. This integration is not always easy to achieve. Given the differences in method and emphasis across disciplines. Nevertheless, it seems increasingly necessary if the field is to move beyond surface level interpretations.

At the same time one might hesitate before framing this integration in overly definitive terms. The relationship between philosophy and technology is unlikely to settle into a stable and predictable form. There will be moments of convergence and moments of divergence and it is perhaps within this tension that the most productive insights will emerge. What matters is not the question finds a clear answer, but that the inquiry remains attentive to its own assumptions. In this sense, the value of philosophical engagement lies not only in the conclusions it offers but in the questions it keeps open.

The ethical dimensions of the discussion also warrant continued attention. Even in the absence of clear evidence of artificial consciousness the ways in which intelligent systems are designed, deployed, and interpreted have significant implications. Questions about responsibility, agency and the attribution of human like qualities to machines are likely to become more pressing as the systems become more integrated into everyday life. Addressing these issues requires not only technical expertise but also a willingness to engage with broader moral and philosophical concerns. It may be that the most immediate challenges posed by artificial intelligence are not about what machines are, but about how humans relate to them next.

Looking ahead, it is possible that some of the distinctions that currently structure the debate between human and artificial intelligence, between simulation and reality will need to be reconsidered. Emerging forms of interaction between humans and machines already blurred these boundaries in subtle ways. Whether this leads to redefinition of key concepts or to deeper appreciation of their differences remains to be seen for now, however, it seems prudent to approach such developments with both openness and caution, recognising that the language we used to describe them may be in need of revision.

Enclosing, the exploration of artificial minds can be seen as a part of a longer philosophical journey, but one that has yet to reach a definitive endpoint. The questions it arises are not new, but they are on new urgency in the context of contemporary technology. Artificial intelligence, in this sense, serves less as a solution than as a mirror, reflecting back to us the complexities of our own attempts to understand the mind. It invites us to reconsider that we mean by thinking, knowing, common experiencing, and to acknowledge that these concepts may not be as transparent as they sometimes appear.

Future research must integrate philosophical inquiry with technological development, not as a matter of methodological preferences, but as a response to the nature of the problem itself. The study of artificial intelligence is, at its core, a study of mind and the mind, as this paper has sought to suggest, cannot be fully understood through technical analysis alone. It requires reflection, Dialogue, and a willingness to remain with questions that do not yield easily to resolution. In this ongoing process of inquiry, the significance of artificial intelligence may lie not only in what it becomes, but in what it continues to reveal about the limits of possibilities of human understanding.

References:

- Chalmers, D.J. (1995). Facing up to the problem. *Journal of Consciousness Studies*, 2 (3), 200-219.
- Clark, A. (1997). *Being There: Putting brain, body, and world together again*. MIT Press.
- Dennett, D.C. (1991). *Consciousness explained*. Little, Brown.
- Descartes, R. (1996). *Meditations on first philosophy* (J. Cottingham, Trans). Cambridge University press. (Original work published 1641)
- Fodor, J.A. (1975). *The language of thought*. Harvard University press.
- Kant, I. (1998). *Critique of pure reason* (P. Guyer and A. Wood, Tans.) Cambridge University press. (Original work published 1781)
- Nagel, T. (1967). What it is like to be a bad? *Philosophical review*, 83 (4), 435-450.
- Putnam, H. (1967). Psychological predicates. In W.H. Capitan and D.D. Meril (Eds.), *Art, mind, and religion*. University of Pittsburgh Press.
- Searle, J.R. (1980). Minds, Brains, and programs. *Behavioural and Brain Sciences*, 3(3), 417-457.