aracne

Classificazione Decimale Dewey:

006.301 (23.) INTELLIGENZA ARTIFICIALE. Filosofia e teoria

GERARDO IOVANE GIOVANNI IOVANE

SOPHIMATICS

FUNDAMENTALS AND MODELS OF COMPUTATIONAL WISDOM

VOLUME II

aracne

aracne

ISBN
979-12-218-2182-6

PRIMA EDIZIONE
ROMA 10 SETTEMBRE 2025

Table of Contents

1. Introduction

The integration of philosophy and technology has been central to human innovation since the beginning of civilization. It has also opened a Pandora's box with serious challenges for the development of AI. In the crossroads between computational science, philosophy, and logic, the question of how machines can go from data processing to interpretation remains. To that end, the present work proposes Sophimatics—the science of computational wisdom. It is born out of the fact that generative AI, while powerful at pattern recognition, lacks common sense, context-dependent interpretation, and a general ethical understanding.

For example, contemporary generative language models (LLMs) like GPT-4 are capable of impressive language generation and data synbook, but their output depends only on correlation and frequency, devoid of any concept of causality or true interpretive understanding. The systems, therefore, do not have a sense of purpose, meaning, or intentionality. Sophimatics is a proposal for addressing this gap in current AI development, incorporating the philosophically foundational categories of meaning, context, and intentionality into models, making interpretive reasoning, contextual awareness, and ethics possible.

The main research question driving this work is: How can the integration of philosophical principles, complex temporality, and computational architectures in the framework of Sophimatics provide a step toward better-reasoning, context-aware, and more ethical AI? The investigation to which this book is a contribution should, therefore, be theoretical in nature, with a practical implementation remaining a follow-up study.

The methodology adopted is interdisciplinary and based on an extensive survey of literature in areas of philosophy (ontology, epistemology, and phenomenology), formal logic and mathematics, computational science, and cognitive science. Literature and theory are also collected and interpreted from related fields, such as social studies and artificial intelligence. The key philosophical traditions driving this project are dynamic ontology (from Aristotle and Simondon), intentionality (Husserl), and dialectical reasoning (Plato and Hegel).

These theories, combined with hybrid symbolic-connectionist logic, multidimensional semantic space modelling, and complex time, are the philosophical and formal pillars of our project. We will refer to the proposed science as Sophimatics, and to the specialists bridging philosophy and computational science as Sophimatics scientists.

Current research is promising but also demonstrates gaps. In addition to exhibiting lack of causal or ethical reasoning, systems like GPT-4 still cannot reason in an interpretative manner (which also depends on the understanding of context and intention) or sustain a reflective dialogue over time. Many researchers have explored the importance of causal reasoning and have designed systems such as hybrid symbolic-connectionist ones to accomplish that goal, finding symbolic and sub-symbolic interactions to be key. Most causal and ethical systems, however, are shallow in the sense of the philosophical foundation, which results in a lack of deeper understanding. Cognitive architectures from cognitive systems engineering address robustness and context management but not comprehensive meaning integration.

This book will be structured in such a way that the reader can understand the theory and structures underlying the constituent elements and how they are combined, thus obtaining a clearer picture of the system we propose. In it, we describe the etymology of Sophimatics, its conceptual core and the central categories that form its foundation. Next, building on the concepts presented above, we illustrate the need for Sophimatics and highlight the limitations of LLM and generative models. Consequently, we model AI within the framework of Sophimatics, outlining its philosophical-cognitive, logical-mathematical and computational-architectural aspects, including time and the Super Time Complex Neural Network (STCNN), the memory and attention subsystem, ethical reflection and dialogue, the concept and role of Sophimatic, and a synthesis of Sophimatics and AI from a conceptual, formal, and architectural point of view. Finally, we discuss the differences between the Sophimatics model and the standard LLM-based paradigm, outlining their strengths, weaknesses, and common elements.

2. Theoretical Foundations of Sophimatics

This chapter aims to explore the philosophical and conceptual foundations of Sophimatics. Intentionality, meaning, and dynamic ontology are examined, and the importance of such concepts to achieve reflective and context-aware AI systems is emphasized. This chapter, as such, may be regarded as the philosophical and theoretical basis of the other parts of this book.

2.1 Etymology and Concept

This section introduced the central idea of Sophimatics as a corrective response to the deficiencies of generative AI. The term "Sophimatics" refers to the science of computational wisdom and is based on the fact that AI, as applied through statistical generative algorithms, fails to adequately represent the key elements of meaning, sense, and intentionality. In this framework, by the integration of philosophical principles into computational modelling, AI is enabled to engage in meaningful and cognitive activities (Vernon & Furlong, 2007).

The name "Sophimatics" has intentionally replaced Newell and Simon's name, which referred to classical AI's notion of epistemology in terms of physical symbol system, in which a system's cognition is based on symbol systems and that the behaviour of that system is explained as symbol evolution over time. As has been highlighted earlier in this work, this approach is inadequate as the machine would never come to understand the symbol and therefore would not be able to dynamically relate that symbol to a context. To avoid this issue, Sophimatics is presented as a rejection of the physical symbol system in which symbol manipulation and execution algorithms are not enough to produce authentic reasoning (Vernon & Furlong, 2007; Rachmad, 2016). Rather, the Sophimatics theory aims to incorporate the concepts of context and ethical resonance as mathematically calculable attributes (Sloman, 1982).

With regard to context, the physical symbol system fails by implementing a syntactic-only function. Hence, Sophimatics is introduced as a corrective means, not by making the formal symbol system "better," but by constructing new computation systems that implement context-driven and ethically resonant computation (Sloman, 1982). This claim is consistent with the first distinction noted by Basti that deals with the subjective part of intentionality versus the objective part of intentionality. While the subjective part of intentionality, the self-consciousness, remains beyond our comprehension, the objective part (Basti, 2014), regarding the logical referents, truth, and the relationship between the parts and the whole in intentional behaviour, are available for computational modelling (Sloman, 1982; Basti, 2014).

In Basti's conceptual representation of intentionality, the dynamic constitution of the intentional reference (for example, a tree) and of the semantic category (for example, an oak) can be addressed mathematically, representing the living organism's capacity for dynamic self-construction (Basti, 2014).

According to Rachmad, contemporary AI systems challenge the epistemological paradigm of sources and means of validation (Rachmad, 2016) of knowledge because they need to not only apply knowledge but also to be able to evaluate the relevance, reliability, and trustworthiness of that knowledge within its immediate context, which is often dynamic and unpredictable. Thus, to deal with the epistemic role of an AI machine, a new AI approach must be taken to implement and integrate meaning, sense, and intentionality, and must include a framework for the machine to question its own knowledge (Rachmad, 2016). As seen earlier, according to Rachmad, this new AI approach is Sophimatics.

As outlined by Siddiqui, ethical decision-making in AI is a growing challenge, and the risk of biases in the source and application of AI-generated outcomes in contexts such as healthcare is a real possibility (Siddiqui, 2024). In this way, to guarantee transparency, accountability, and overall validity of AI-driven interventions, ethical principles must be integrated into AI development as a whole (Siddiqui, 2024), and, indeed, this ethical, interpretive, and philosophical approach is a key element of Sophimatics (Sloman, 1982).

In addition, one way to define Sophimatics is that it is a correction for the lack of "knowing" of today's systems. According to Sloman, for AI to actually "know," formal language system algorithms, in which a sentence's parsing trees (formal, syntactic descriptions of the structure) is generated in terms of the sentences words and parts, are not enough (Sloman, 1982). This is exemplified in the sentence "they watched the man with a telescope". By only formalizing language into such trees, one does not address ambiguity (man with the telescope or they looked at the man through the telescope) or, ultimately, understanding, both of which are inherent cognitive abilities. For that matter, to address ambiguity and overall understanding requires an algorithmic system of parsing that implements, along with the tree and its constituents, a representation of the relevant ambiguities and the means to resolve them (Sloman, 1982).

From this, another difference between Sophimatics and generative AI emerges: Generative AI is not only unaware of ambiguity, but it aims to eliminate any such complexity. Sophimatics, in contrast, aims to incorporate not only language, but the ability to deal with and manage ambiguities, into its algorithms. To implement this, AI systems must actively be able to deal with ambiguities and interpret in real time, rather than simply accept the complex surface statistical distribution generated from complex natural processes as is (Sloman, 1982).

Philosophical questions also play an important role in the foundation of Sophimatics, in the sense that rational knowledge and wisdom were two key factors lacking in early applications of computational wisdom in general, and these are philosophical questions that the field of cognitive science aims to resolve. In other words, the reason why Sophimatics incorporates philosophy as a major component is to address the question of rationality, in addition to addressing the philosophical issue of the origin of knowledge, in particular, by expanding the concepts of intelligence and knowledge from the philosophical and psychological frameworks (Sloman, 1982). In this way, according to Sloman, Sophimatics addresses and answers philosophical scepticism with AI by presenting conceptual architectures and working implementations in which (a) an AI system can continually reinterpret situations and revise what it believes, (b) the concepts of sense, meaning, and intentionality are inherently necessary, and (c) the notion of reflective reasoning as based in context, and not simply abstract operations, is made computable (Sloman, 1982).

Along the lines of answering philosophical scepticism, referring again to Basti's notion of subjective versus objective intentionality (Basti, 2014), the subjective level of intentionality (self-consciousness) will remain unanswerable to science, whereas the objective elements, such as the references of logical symbols and the truth values of propositions (Basti, 2014), can be made computable (Sloman, 1982; Basti, 2014). The ability of a system to adaptively update its own ontology allows this referential intentionality to be applied in living beings in a natural way. Therefore, by focusing on these aspects of intentionality, AI systems are better enabled to capture such living abilities to engage in context and to adapt. This dynamic nature also addresses philosophical scepticism regarding the possibility that machines can intentionally reference and mean things as a result of computation (Sloman, 1982; Basti, 2014).

A further challenge to building machines with cognition is how to formalize and represent intentionality logically. In this domain, philosophical literature proposes incorporating logical references to logical symbols to describe and compute intentionality. This approach offers the ability to describe the reference of a logical symbol as well as update that reference at each point in time within the system's cognitive activity. Hence, by having a way to dynamically specify the current intention of the symbol, we can embed intentionality as a computable concept within the system to achieve the desired agency. Sloman (1982) and Basti (2014) propose to integrate ideas from neurocomputation theory, such as distributed processing, high-level abstraction, and context-driven activation schemes, to implement computational architectures for intentionality and cognitive systems in general. The use of these methods results in semantic enrichment and real-time applicability.

Rachmad (2016) also highlights that epistemology can be advanced to a broader scope within the context of AI. Current trends show how epistemology is no longer solely focused on the sources of the "correct" knowledge (Rachmad, 2016), but it challenges the fundamental notions of correct knowledge, wisdom, rationality, and what makes up a valid belief system in the light of technological growth (Sloman, 1982; Rachmad, 2016). This perspective requires a novel approach to epistemology. To address this concern, the author considers an AI system, implemented by the ideas noted earlier to represent meaning, sense, and intentionality, in which any kind of information stored by the AI system must always be checked and reinterpreted (Rachmad, 2016) because, in reality, evidence is never completely complete and stable. According to Rachmad, an AI implemented in this manner has gained "computational wisdom" (Rachmad, 2016) and serves to address the problem of knowledge for systems attempting to implement natural cognition.

Finally, further justifying the need for Sophimatics, according to Siddiqui, ethical factors must be formalized as inherent and fundamental to any AI development (Siddiqui, 2024). For the development of safe and unbiased algorithms, it is also key to acknowledge the social, interpretive, and philosophical dimensions when AI is used to make decisions in highly sensitive and context-driven areas (Sloman, 1982; Siddiqui, 2024). As seen, these dimensions are intrinsically formalized in Sophimatics (Sloman, 1982; Rachmad, 2016). Siddiqui notes that without these ethics, AI is rendered irresponsible and often flawed and biased (Siddiqui, 2024). Therefore, in the realm of safety, responsibility, and context, Sophimatics must be implemented in AI-driven medical implementations.

In conclusion, the hypobook is that Sophimatics is a step in the right direction for representing and implementing meaning, sense, context, and intentionality. This new theory of computation wisdom must be implemented within the structure and operation of the AI systems of today. Furthermore, by embracing this new theory, a co-evolution of human and machine will emerge (Rachmad, 2016).

2.2 Philosophical Principles

Ontological grounding in Sophimatics goes beyond traditional approaches to ontologies in many aspects. It addresses certain limitations of classical AI in grounding philosophical, intentionality, and contextual concerns. Unlike the symbol system and physical symbol system hypobook of Newell and Simon that conceptualizes cognition as an evolving collection of symbol structures, Sophimatics believes categories, concepts, and their relations are continuously defined by the context and environment. The agent, therefore, can dynamically and continuously adapt to and reinterpret its knowledge (Vernon & Furlong, 2007).

Moreover, by enactive epistemology, cognition of a viable agent can only be achieved through constant interaction and negotiation of an agent with the context of the environment and other agents. In classical AI, symbols are pre-fixed, with a fixed sense and a static reference to their real-world referent. But in Sophimatics, a symbol derives its sense from its relation to its real-world referent, and its meaning from its relations to other symbols and agent goals. Therefore, meaning is the result of a co-determination between an agent and context. By rejecting the symbol-manipulation view of cognition, Sophimatics gives importance to situatedness, adaptation, and self-organization. Context has become a critical factor for agent intentions and behaviours, and this cannot be described via traditional cognitive or AI frameworks (Vernon & Furlong, 2007).

Linguistic or cognitive behaviour must be relational. Sophimatics believes a shared world will give rise to collaborative behaviour and mutual understanding among cognitive agents. This is empirically supported by the origin of language itself that came from the need to communicate at the crossroads of ontogenetic and phylogenetic development. The implementation of Sophimatics requires the construction of dynamic and relational ontological schemas that can flexibly restructure its representation of the world in new and unknown circumstances. Classical AI used static and predefined taxonomies for representing its knowledge, and it fails to adapt quickly to semantic changes in its surrounding context, to situations it has never encountered before. The capability to adapt, redefine, and relate knowledge is crucial for the design of systems that perform tasks in ill-defined and dynamically changing domains (Vernon & Furlong, 2007).

Also, a classical computational epistemology views “knowing” as the parsing of syntax to form either a syntax tree or a network that enables further structural operations. For example, “they watched the man with a telescope.” An AI can generate the appropriate phrase-structure description to the sentence and then decide from its vocabulary of “structural descriptions” what is the “correct” syntactic representation of the sentence. But according to Sophimatics, this falls far short of “knowing” the meaning of the sentence (Sloman, 1982).

Another deficiency to be avoided is the temptation to require that AI systems be given complete information and no ambiguity, such as to require the program of answering yes or no only, or to allow the system no option to say it “doesn't know.” Instead, Sophimatics insists on systems capable of resolving ambiguity, reasoning from incomplete information, interpretive frameworks, and handling beliefs about “the world.” In fact, they ought to operate with these kinds of uncertain conditions, such as those of everyday human cognition. So, systems can reason in probabilistic rather than binary fashion (Sloman, 1982).

Furthermore, it will be wrong to design the epistemology that always requires information of a purely “descriptive” or “factual” sort. In particular, it will be useful to allow programs to receive, generate, and respond to open-ended or evaluative utterances. However, such an ability is a relatively useless and superficial feat without also developing the interpretative epistemological stance (Sloman, 1982).

Lastly, it must be remembered that not all computational epistemologies are necessarily fixed and static. AI systems ought not to be designed with any notion that their current ways of describing the world in terms of knowledge claims will remain forever adequate to explain phenomena. Instead, epistemologies must enable the system to change with the world so as to constantly allow for the possibility of reinterpreting existing beliefs to accommodate newly perceived objects and contexts (Sloman, 1982).

In relation to the phenomenological framework, there must be a clear statement that intentionality can be programmed. Subjective self-consciousness or subjectivity cannot be computationally programmed due to its impossibility of verification with current scientific techniques and methods. But the objective aspects, such as logical symbol reference, relational aboutness, and truth conditions, can. By formalizing the ideas from Husserl, intentionality becomes a computational attribute that can be coded, and for AI that is the central task of programming computers to emulate intentional behaviour (Vernon & Furlong, 2007; Rachmad, 2016).

Furthermore, by embedding intentionality, AI systems can be said to know something in the following sense: for any given data, its action can be viewed as relating those data to some goal of the agent. In classical AI, the flow of data from initial state to final outcome is purely causal, and in the sense, passive, for these data are simply processed according to prescribed rules without being related in any manner to an agentive goal. Furthermore, since subjective intentionality (or even the fact of intentionality per se) is inherently unknowable, it is difficult for science to explain human action as deriving from intentionality and, at the same time, provide a mechanistic explanation (Basti, 2014).

Phenomenological frameworks in Sophimatics should permit the modelling of the concept of intentionality. This means the modelling of concepts that are capable of not only computing on external data but also able to analyse how the data relate to them as the agent. By relating data to agentive goals, the system acquires the purposeful aspect of human intelligence. By allowing the simulation to acknowledge the dual sense of intentionality, science becomes able to represent intentionality in a manner of simulation. So, intentional behaviour will be enabled, while retaining the capability to assess the meaning of intentionality as an objective construct (Basti, 2014).

Intentional systems are capable of understanding their own actions and of calculating the impact of those actions on others and on the general situation in which they are enmeshed. Thus, intentionality as modelled through a phenomenological approach ensures that AI systems exhibit an awareness of their actions towards others. Such an awareness enhances individual cognitive performance and informs collective decision-making, promoting better outcomes at both levels (Vernon & Furlong, 2007).

There are intentions that are created jointly by groups of people to achieve a common goal. These goals can be fulfiled only by the actions of several people operating according to some common plan, known as teamwork or team reasoning. Most AI designs simply assume that each agent has intentions that the system optimizes at each moment, and these do not result from a common intention. But Sophimatics includes the aspect of "we" (intentionality, goals, beliefs, or commitments) to the analysis. By adding "we," agent intentions are now determined by group goals (Petersson, 2017).

Team reasoning is based on the concept of joint intention. It is a method used to analyse the rationality in social interactions where actions are taken for a collective objective. Thus, team reasoning brings a new perspective to the subject of intentions by shifting its focus from individual agency to the collaborative dynamics within a group. An example of collaborative intelligence that is incorporated into Sophimatics is team reasoning, which allows the individual agent to decide at the beginning of each decision process, either to reason according to individual reasons and goals, or collectively from the point of view of group-oriented motives and reasons (Petersson, 2017).

Sophimatics criticizes classical epistemological descriptions of AI. The explanation of AI systems, such as systems that "generate an evolving collection of symbol structures" (Sloman, 1978), is not the sort of explanation in which AI is interested because it does not explain the processes and contents of those beliefs. To be useful to Sophimatics, it has to explain the meaning of intentions of how those categories of human intelligence exist, as well as how the relationships among them are discovered. Instead, AI must be able to offer an explanation that illuminates how these categories are able to emerge in an entity such as a computer or, for that matter, in the brain (Rachmad, 2016).

So, systems must be able to revise their claims to knowledge when new evidence becomes available. Since these claims will inevitably be based on inferences going beyond what is given, this entails the development of heuristics for the evaluation of their trustworthiness. Moreover, it is quite apparent that humans do not act in such a way as always to maximize truth, as would be the result if every thought or sentence were based on an analysis of what one knew, and if these pieces of knowing were not accepted in the presence of some reason to reject them. The ethical stance requires the same element, the evaluation of one's knowledge that could be "trusted," as for this purpose, is related to ethical self-assessment. The point is that one wants an AI system not merely to be able to solve isolated problems but to assess its own ways of thinking and to find and use other frameworks in so doing. So, AI systems should be designed with internal evaluative criteria whereby the system can itself analyse the validity of its methods of inference and of its knowledge claims (Rachmad, 2016; Sloman, 1978).

2.3 Key Concepts: Meaning, Sense, Context

The significance of meaning, sense, and context to cognitive AI is for the same reasons that these three features are vital for human-level reasoning. Godden and Baddeley demonstrated the alignment of context with memory as applied to SCUBA diving, which showed that context played a crucial role in recalling certain words. As such, the environment has a substantial impact on the recollection of words learned in different environments versus the environment in which they are being recalled. Similarly, for any AI systems to have grounded reasoning, they must have contextual data embedded into their framework. Traditional AI algorithms, without contextual data, have outputs that are predetermined. Thus, it is crucial to enable the context to vary in any form to achieve an ideal imitation of human cognition. Context is, therefore, an active feature that shapes the way an agent comprehends meaning. (Hollister et al., 2019)

The significance of personal relevance also plays a major role in human-level reasoning and cognition. Treisman discovered that if an auditory stimulus that was played at different volumes, and participants were to name their names or any other high-relevance stimulus, that stimulus would drown out all the other stimuli. Hence, a relevant word in a certain situation could be heard and remembered over all others. If cognitive AI algorithms can perform relevance filtering and identify features that are relevant to either the agent or the environment in the moment, then all the outputs of the algorithm can align themselves towards a specific target or goal. This is in direct contrast with static AI algorithms as they cannot grasp the relational properties of contextual and personal relevance. Relevance, therefore, filtering is another major pillar of the Sophimatics model as it can filter data in an efficient way to achieve specific goals and provide dynamic system adaptiveness. (Hollister et al., 2019)

It is not enough to statically bind tokens to particular meanings, as is the case with classical AI. The studies in memory and attention illustrate that meaning arises situationally as a product of how our mind dynamically associates these aspects with the stimulus or situation at hand. As a result, current generative models lack the ability to effectively incorporate human-level context sensitivity and are oversensitive to the noise associated with contextual ambiguities. This provides an answer as to why Sophimatics is vital. In this case, Sophimatics provides the discipline necessary for the computational reconstruction of meaning as a relational construct. To make the distinction between meaning and context clear and to ground the idea of a contextual model, hybrid symbolic-connectionist AI approaches show promise. The architecture provided by Wermter and Lehnert shows that an AI can integrate localist and symbolic networks. Their results prove that integrating three different layers, (neural (distributed), localist, and symbolic (abstract)) enables the recognition of different levels of complexity in the input data. Thus, by capturing structural, semantic, and relational features in each of the layers, they correctly process data and make sense of it by resolving semantic ambiguity to be able to correctly parse novel inputs (inputs to the AI that the AI has not seen before). In essence, this means that multi-layer architecture outperforms flat-layer architecture that classical AI typically implements as well as corroborating the Sophimatic principle of meaning arising through emergent phenomena from the integration of layered processes. In terms of human-level AI, as AI systems evolve, they will have to similarly implement this emergent property and enable systems to build up sense by themselves.

As per our computational epistemology, it can be claimed that for AI to emulate human-level sense-making, it must embed the context and relevance into its fundamental structure as an operational parameter. Sloman provides an interesting computational epistemology to prove that sentences are not interpreted by just a parsing algorithm of grammar, which is often assumed as the only method for interpreting human language. He proposes a situation in which a sentence can have multiple interpretations, depending on the context. The example is “they watched the man with a telescope.” How could an agent know that “they watched with a telescope” and “the man has the telescope” and distinguish it from “they watched (someone) with a telescope” and “someone watched the man”? The solution to this issue is to constantly evaluate new interpretations that could be attributed to existing information in a system's memory and, more importantly, to do this while embedding the ability for current perceptions and data to restructure the previous interpretations in the system's memory. This will allow for a system to adapt its interpretations to the context of a situation dynamically, much like humans do. By being able to adapt to context or changing environments, any information perceived will directly have the ability to impact the current state of the system and its ability to adapt for future interactions with the environment. (Hollister et al., 2019)

For semantic precision and interpretative accuracy, an ontology-driven categorization provides an even more robust and computationally efficient framework than relational data structure. Maier et al. showed in their work with an embodied cognitive agent that including ontologies in AI programming allows a system to categorize data effectively by including both context and relationships, and that by not having an explicit ontological structure, the agent makes frequent mistakes. Moreover, by employing ontological frameworks for data representation, one can represent entities within hierarchies and infer higher-order relations between multiple concepts and entities. Ontologies are an effective way to formally characterize the knowledge structure of classes and properties, which provides a basis for AI reasoning that can then make connections between objects, their properties, and their environment over time. By representing the knowledge of a system in the form of relationships, it can address many issues with the token-focused algorithms that are prevalent today, such as the inability to process novel inputs and manage relational ambiguity. Also, by inferring additional sense-rich relationships, ontological networks allow for a system to adapt to the environmental context. In essence, ontology-driven frameworks for representation must also be included in the model proposed to fulfil the requirements of semantic precision, interpretative accuracy, and adaptive cognition as stated in the core tenet of Sophimatics. As mentioned, current generative approaches do not account for the multiple requirements needed for robust AI. The principles provided in the core tenet all have to work synergistically to emulate the way that humans make sense of the world and provide a suitable framework for the development of complex AI systems. These requirements are context sensitivity, relevance filtering, ontological embedding, and ambiguity resolution. By providing these frameworks, Sophimatics becomes the only discipline for the development of human-level AI. All these systems alone are not sufficient to allow an AI system to achieve a desired level of human-level reason. All of the frameworks must function in a combined way to achieve the end-state goal of building cognitive AI systems.

Thus, by embedding meaning, sense, and context into any cognitive model and treating them as the key operational features of a system, the shift in AI that we are currently experiencing becomes the paradigm shift in AI, as called for by Sophimatics.

2.4 Necessity of Sophimatics

The necessity of Sophimatics within artificial intelligence is due to the missing dimensions of temporal, contextual, and ethical consciousness within modern cognitive architectures. Research in social human-robot interaction has clearly shown the need for holistic approaches that integrate time, ethics, and context when modelling AI. For instance, Baxter et al. (2016) find that there is a lack of AI architectures that properly integrate time and content in agent-environment exchanges. They argue that current architectures that account for complex temporal coordination between agents, but do not take into consideration the temporal aspects of the content in agent-environment exchanges, fail in proper social cognition.

Current cognitive AI architectures focus on generality and broad applicability, failing to account for complex coordination involved in social behaviour, as well as the ability to adapt to dynamic environments. As Baxter et al. (2016) show, current architectures lack mechanisms that incorporate time and content into agent-environment exchanges. Therefore, they lack the ability for social cognition. This reveals that modern cognitive AI architectures lack mechanisms for context evolution, thereby failing to account for many real-world, dynamic scenarios. Sophimatics focuses explicitly on dynamic temporal cognition, dialogic memory, and contextual evolution.

The absence of formal mechanisms for ethical and temporal reasoning in AI results in systems unable to adapt to dynamic scenarios. As opposed to an artificial context, Sophimatics considers context as a dynamic factor that can interact with memory, behaviour, and ethical concerns in a manner that changes over time. Baxter et al. (2016) argue that modern approaches to contextual reasoning are largely superficial in that they lack the ability to "go beyond mere plausibility" and cannot avoid the generation of inappropriate responses, even in clear contexts. Therefore, the emphasis on context evolution in the Sophimatics paradigm becomes essential in humanized artificial intelligence.

The necessity of holistic approaches becomes especially true when AI is intended to engage in long-term goal-oriented behaviour with significant impacts on social ethics and dynamic environments. Baxter et al. (2016) show that current architectures perform well in common and obvious scenarios, but crumble when faced with complex, novel, and ethically ambiguous problems. Sophimatics includes temporal layers and dialogic processes that allow the AI agent to reflect, not just act in accordance to past actions and experiences, thus anticipating long-term consequences of the current scenario.

Baxter et al. (2016) show how a system's failure to build a flexible, appropriate cognitive structure often leads to catastrophic errors, due to their lack of the ability to adapt to the novel situations described above. This supports the necessity of the concept of adaptive reasoning. Context, ethics, and intentionality must not be accounted for as constants. They must evolve along with agent interactions and environments to allow for accurate and effective performance in ambiguous circumstances.

Humans utilize context, environmental consistency, and personal relevance, thereby affirming the necessity of Sophimatics, as explained by Hollister et al. (2019). Human memory and recall are greatly influenced by external context consistency. Studies show how humans who memorized content while either being on dry land or underwater during SCUBA diving recalled this content at much higher rates when the recall conditions mirrored those of memorization (Godden & Baddeley as cited in Hollister et al., 2019). Given that context is not a constant, this suggests that artificial intelligence approaches that solely rely on statistical associations and machine-learned patterns that do not include external and internal context evolution are doomed to failure in many circumstances. The utilization of context evolution, implemented as a dynamic constraint in human-machine memory and recall, is a vital addition.

The importance of personal relevance and context prioritization for humans is clearly shown by Treisman's dichotic listening experiments, where a stimulus (hearing one's own name) increased attentiveness of human subjects. As shown, the human brain amplifies stimuli, prioritizing information it deems immediately relevant. Hollister et al. (2019) argue that AI systems must filter according to immediate relevance to properly satisfy contextual objectives. This requires a prioritization framework that has not been incorporated into current AI models, which emphasizes the need for systems to incorporate dynamic methods for prioritization and adaptability.

Current AI systems fail to handle the constantly shifting state of meaning within the human experience, instead artificially imposing a static assignment of meaning. Hollister et al. (2019) state that contemporary AI models oversimplify and constrict meaning. By oversimplifying meaning down to token matching and token replacement, systems lack the ability to utilize any contextual or situational cues to reconfigure and re-assign context or relevance to particular constructs. Therefore, they suggest that AI must be able to alter their own definition of meaning within each interaction. Meaning should be redefined as a context-aware relational construct. Thus, it must be reconstructable from situational cues at any time (Hollister et al., 2019).

The requirement for multi-level, layered models is supported by the hybrid symbolic-connectionist architectures of Wermter and Lehnert, which show the ability to adequately model language through integrating connectionist neural network layers (localist) with symbolic AI layers. Their architectures are able to process both complex linguistic structures and remain semantically flexible. Sophimatics implements a similar, tiered model structure of meaning, which operates by the assumption that it is built on distributed-level processes and utilizes intermediate and abstract levels. Wermter and Lehnert have contributed much to advancing the idea of the importance of layered models in building truly cognitive AI systems. Sloman's work on the need for AI systems to implement "Computational Epistemology" further supports the necessity of dynamic, integrated modelling. The fact that parsing of meaning is generally based upon the immediate context of the symbols being used can be problematic. Sloman points out how sentences like "they watched the man with a telescope" require additional context from memory for disambiguation. By failing to account for prior experiences and knowledge, as well as being unable to continuously refine their representations of an event or symbol, current AI systems are unable to resolve ambiguities of meaning. Hollister et al. (2019) propose that AI systems utilize dynamic memory to solve this problem.

Maier et al.'s work on ontological embedding for data and symbol categorization furthers the philosophical arguments that support Sophimatics' approach to symbolic and contextual categorization in that they claim how structuring symbols into ontological models provides crucial context and relationships, which are otherwise not contained in data. According to Hollister et al. (2019), ontological grounding is essential for achieving semantic precision in context. The ability of current generative models to capture and maintain relationships within semantic structures is poor, which is what makes ontological categorization in AI so difficult to attain. The embedding of concepts in an ontological structure is essential for AI to interpret and apply data and symbols in the way in which humans naturally do, without the over-generalizations described above.

As described by Saad (2025), philosophers argue that simulation is not a representation of