

The meta-ontology of AI systems with human-level intelligence

Roman Krzanowski
Pawel Polak

Pontifical University of John Paul II in Kraków, Poland

Abstract

In this paper, we examine the meta-ontology of AI systems with human-level intelligence, with us denoting such AI systems as AI_E. Meta-ontology in philosophy is a discourse centered on ontology, ontological commitment, and the truth condition of ontological theories. We therefore discuss how meta-ontology is conceptualized for AI_E systems. We posit that the meta-ontology of AI_E systems is not concerned with computational representations of reality in the form of structures, data constructs, or computational concepts, while the ontological commitment of AI_E systems is directed toward what exists in the outside world. Furthermore, the truth condition of the ontology (which is meta-ontological assumption) of AI_E systems does not require consistency with closed conceptual schema or ontological theories but rather with reality, or in other words, “what is the world” (Smith, 2019, p.57). In addition, the truth condition of AI_E systems is verified through operational success rather than by coherence with theories. This work builds on ontological postulates about AI systems that were formulated by Brian Cantwell Smith (2019).

Keywords

human-level intelligence AI, meta-ontology of AI Paradigm, AI

Paradigm, ontology of AI Paradigm, ontological commitment of AI Paradigm, John Haugeland, Brian Cantwell Smith, Marvin Minsky, Hubert Dreyfus.

Introduction

Artificial intelligence systems have developed over the past 60 years, bringing new solutions to a huge number of practical problems, and they continue to find many surprising and fascinating applications. However, the main goal of AI, namely the creation of a human-like intelligence,¹ is still proving unattainable. Indeed, the AI systems we currently design and implement cannot replicate human intelligence and a human agent's ability to cope with reality (see, e.g., Brooks, 1991; Minsky, 1991; Dreyfus, 2016; Mitchell, 2019; Bołtuć, 2020; Roitblat, 2020; Wooldridge, 2021).

One of the reasons for this failing (in their ability to cope with reality) is, it seems, related to how these AI systems lack a proper ontology or representation of the real world. (For more about the failings of current AI conceptualizations, see, for example, the works of Brooks (1991), Dreyfus (2016), and Smith (2019)).² Smith (2019, p.44) proposed four features that AI systems should possess if they are to mimic humans' ability to cope with the real world (i.e., have human-level intelligence). These systems, which we here refer to as AI_E systems, should be embodied, embedded, extended, and enactive

¹ See ft. 3 on AGI for an explanation of human-like intelligence.

² To get a sense of the ontologies (and meta-ontologies) of the real world in biological agents, consult Ed Yong's book *An Immense World* (Yong, 2022). Interesting analysis of deep learning systems ontological commitments see (Šekrst and Skansi, 2022).

(see also Käufer and Chemero, 2021). While these features are not ontological per se, but they do imply a commitment to some ontology. What the kind of ontology these four features would entail is a meta-ontological question that we explore in this paper, thus furthering the ideas studied by Krzanowski and Polak (2022).

This paper is organized as follows: First, we define some basic concepts related to the meta-ontological discourse, namely ontology, specifically ontology of computing and AI systems, meta-ontology, ontological commitment, and truth conditions. As these concepts have many interpretations, we need precise definitions to ensure that the subsequent discussion will be understood as intended. Next, we explain the main postulates of meta-ontology for AI_E systems, the topic of this work. In the conclusion, we discuss the inherent limitations of ontologies (which is a meta-ontological problem) for artificial systems such as AI_E, their inability to match human intelligence, and the potential prospects for AI_E. (For more on the problems of ontology in AI systems see, for example, Haugeland (1985) and Fjelland (2020)).

Three things should be borne in mind while reading this paper. First, AI systems with human-level intelligence are often referred to as AGI systems, but with many interpretations for this concept, we avoid using this term to prevent us from drifting into the debate about AGI.³ Second, this is a study of the meta-ontology of specific AI systems, i.e., AI_E, meaning that the focus of the study is ontology of these

³ See various references for different conceptualizations of AGI (e.g., Mitchell, 2019, p.40; Fjelland, 2020); general purpose, human-level intelligence (Marcus, 2022); the generic ability of a machine to consciously perform any task that a human can (Swar, Khoriba and Belal, 2022); the intelligence of a machine that is capable of understanding the world (Skuzza, 2020); the representation of generalized human cognitive abilities (Lutkevich, 2022); a general-purpose capability, including the ability to broadly generalize to fundamentally new areas (Cassimatis, Bello and Langley, 2008); and the capacity of an engineered system to display the same rough sort of general intelligence as humans (Goertzel, 2015). Creating human-level intelligence

AI systems rather than philosophical ontology. Philosophy forms the background of this discussion, but it is not its main objective. Third, we are not concerned with particular implementations of AI_E systems, which is why we instead study the AI_E system paradigm, which is the all-encompassing conceptual framework of AI_E systems that supports multiple implementations. Thus, when we talk about an AI_E system, we are referring to an AI_E system paradigm rather than a specific implementation. The concept of AI paradigm and its role in this study are explained later in the paper.

Key grounding ideas

The ontology of AI

Ontology can be thought as an empty buzzword or a specific conceptual construct,⁴ so we need to position the ontology of AI_E within the world of ontological theories and demonstrate, what it means in this discussion, how it relates to “other” ontologies in philosophy, computer science, and AI. Ontology in computer science and AI systems (e.g., Sánchez, Cavero and Martínez, 2007; as well Guarino and Gia-

was always the aim of AI research, as attested to by Yann LeCun’s recent claim “Getting machines to behave like humans and animals has been the quest of my life” (reported in 2022 MIT Technology Review (Heikkilä and Heaven, 2022)).

⁴ “We must be careful in reading [auth. any] philosophical works on ontology, when the author speaks of ‘ontology’ without qualifications, not to confuse the intended sense of the world with any of the alternatives” (Jacquette, 2002, p.3). There is also a confusion between ontology and metaphysics. Some authors see ontology as the ultimate study of reality (e.g., Jacquette, 2002; Stróżewski, 2004; or Perzanowski, 2015), and metaphysics as being “after physics”, some others see ontology as a part of metaphysics (see e.g., Van Inwagen, 2009). In AI literature because ontology takes on a very concrete garb (of an engineering domain) metaphysics is a rare term so the confusion is not so visible.

retta, 1995; Guarino, Oberle and Staab, 2009; Swar, Khoriba and Belal, 2022) takes on different meanings to that of philosophy (e.g., Jacquette, 2002; Stróżewski, 2004; Baker, 2007; Chalmers, Wasserman and Manley, 2009; Effingham, 2013; Berto and Plebani, 2015; Perzanowski, 2015; Thomasson, 2015; Hofweber, 2021).⁵ The philosophical concepts of ontology, however, are fundamental to those used in specific applications (see the comments of Jacquette, 2002, p.XII). This is therefore where we begin.

In philosophy, ontology is the study of being as it is (i.e., “what is”), so it is about “being” in the most general sense.⁶ More specifically, in its purely philosophical meaning, ontology is the study of the foundations of what exists, what is common and most general among it, and what its origins are (see, e.g., Jacquette, 2002; Stróżewski, 2004, p.32). Hereafter, we refer to this concept by a boldface, capital “O” without a subscript (i.e., **O**).

Ontology in philosophy may also refer to “what exists” in a much more constrained, narrower, sense. This kind of ontology investigates existing, subject to a definition for existence, objects and relations in the world, and we will refer to this ontology as O. Depending on the assumptions made, different types of objects and relations may be recognized by O ontologies, because O branches into many

⁵ Importing AI (or technical) aspects into philosophy brings with it a touch of reality that philosophical considerations often lack. See also the comment by Jacquette on the relation between a domain ontology and the domain itself (Jacquette, 2002, p.5).

⁶ The term “being” is used in the sense employed by the Ancient Greeks, Parmenides (opposite to The Unbeing), Aristotle (Being qua being), medieval scholars like Aquinas (the study of being qua being) (Kerr, 2022), and some modern philosophers such as (Jacquette, 2002; Stróżewski, 2004; or Perzanowski, 2015). This term is also sometimes written as Being meaning totality of what exist (Kenny, 2012, p.160). Many modern philosophies, scientists computer engineers infuse this term with many different meanings (see this paper for the examples) obfuscating the original Greek sense of *onto-logia* – the fundamental study of being, probably as too esoteric (i.e., metaphysical) for their tastes (see also Kenny, 2012; Hofweber, 2021).

subdomains. Thus, many different perspectives have been developed for ontology (e.g., Quine, 1960; Jacquette, 2002; Strózewski, 2004; Baker, 2007; Chalmers, Wasserman and Manley, 2009; Effingham, 2013; Ingarden, 2013; 2016; Berto and Plebani, 2015; Perzanowski, 2015; Thomasson, 2015; Hofweber, 2021); these differ in terms of extent, content, consistency, and accuracy, often responding to the specific needs of a domain.

In computational systems, ontology can be defined as “a specific vocabulary (dictionary) used to describe a certain reality, plus a set of explicit assumptions regarding the intended meaning of the vocabulary words” (see Guarino and Giaretta, 1995; Guarino, Oberle and Staab, 2009). In this context, ontology may also refer to “a model of the structure of a system” (Guarino, Oberle and Staab, 2009) or “a formal, explicit specification of a shared conceptualization” (Studer, Benjamins and Fensel, 1998). We also have computational ontologies, which often called engineering ontologies, that “are machine-processable structures which represent particular domains of interest” (Husáková and Bureš, 2020). Ontology may also be used to refer to knowledge-based systems, databases, or AI systems that manage knowledge bases (see the discussions of Sharman, Kishore and Ramesh, 2007; Staab and Studer, 2009; Garbacz and Oliver Kutz, 2014; Husáková and Bureš, 2020).

The ontology of AI_E systems, meanwhile, denotes and determines how an AI_E system represents and reasons about the real world. It is not concerned with computational representations of reality through structures, theories, data constructs, or computational concepts but rather with how real world objects, properties, and relations are registered by an AI system, as well as how they are recognized and interpreted. In brief, this ontology is solely committed to the real world (in the sense

explained by Smith, 2019, p.145), with it reflecting the real world⁷, or physical reality, and the AI system's place in this world. More specifically, it is embodied, embedded, extended, and enactive (Smith, 2019, p.43), which are terms that will be explained later in this paper. There is no formal theory to accompany 4E ontology (ontology of AI system that is embodied, embedded, extended, and enactive), so there are no criteria for theoretical truth verification, but verification comes instead from a confrontation with the real world, which we will discuss later. A 4E ontology is not given in the form of a set of a priori relations and objects but rather acquired (Smith, 2019) in response to a dynamically changing reality (see Minsky, 1991, as well as; Bołtuć, 2020). We will refer to this ontology as O_E .

Let us now put these things together. Ontology (**O**), as a philosophical discipline, asks what is, in a most general sense, and what exists. Ontology (O) is more restricted with the scope of this ontology being defined by the horizon of interest: For example, it may be the universe, some aspect of it, or a domain of reality (i.e., domain ontology), like ontology of biology, or ontology of physics. Ontology in computational systems, meanwhile, can be defined as “a specific vocabulary used to describe a certain reality, plus a set of explicit assumptions regarding the intended meaning of the vocabulary entries,” while the ontology of AI systems relates specifically to a representation of the

⁷ The term “real world”, or reality, is understood here as it is understood in Smith (2019, p.xiv), denoting the physical world we live in. Minsky (1991, p.6) refers to this reality as common sense reality. The term may be opposed to “virtual worlds”, “imaginary worlds”, “fantastic worlds”, or other qualified uses of “words” denoting worlds as creations of computer systems, artistic expressions, or imaginations. The term “world” or “real world” may have multiple interpretations that we have no intention to discuss as such a discussion would be pointless and would not further the main point of the paper. Thus, the reader seeking more detailed explanation of this term should follow the cited references.

world or some knowledge domain in AI systems. Finally, the ontology of an AI_E system (O_E) refers to how AI_E system represents about the real world and how it is situated within reality.

Figure 1 illustrates these ontological dependencies by showing them in terms of their increased specificity both in scope as well as in application domain, from the most general (**O**), which is the ontology of what exists (**O**), to the most specific one, which in this case is O_E , the ontology of an AI_E system.

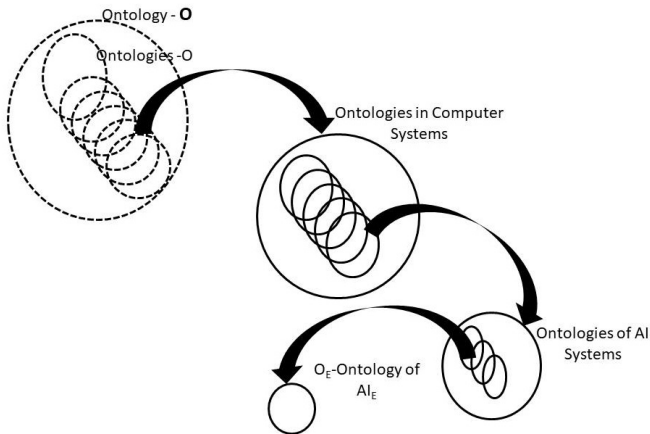


Figure 1: AI_E ontology and hierarchy of ontologies.

Thus, with respect to specificity and scope, O_E falls under the AI ontologies, which in turn are subspecies of the ontologies of computer systems, and these in turn are forms of specific ontologies (**O**) concerned with a specific segment of reality, which then falls under fundamental ontology (**O**) for existence. All these ontologies ask the same question but within different contexts, scopes, and perspectives. Furthermore, Figure 1 shows how only ontology **O** attempts to com-

prehend all that exists, with other ontologies being mere fragments. The further away we move from the fundamental ontology **O**, the narrower and more specific the scope of the ontology becomes.

Thus, ontology that represents the real world being always expressed in some form of language (which always is) is always incomplete with respect to the world, i.e., the world “as is” cannot be represented by something else, despite the close isomorphism between the reality and some logical aka ontological systems.⁸ As a result, there will always be some inconsistency between what exists and what a given ontology represents, because ontology (even ontology **O**) will always be a theory (of sorts), expressed in some specialized language (see ft. 8) about the world rather than the world itself. Thus, there will always be a degree of incommensurability,⁹ or a gap, between the ontology of synthetic systems and the reality of what exists. This gap may be narrowed but never entirely closed. Indeed, a representation can never contain itself as a part of what exists. We should keep this incommensurability in mind when building synthetic systems like AI_E.

The meta-ontology of AI

Meta-ontology is a relatively recently coined term with many interpretations. In philosophy meta-ontology denotes a study of ontology, what it investigates, and what it is concerned with (e.g., Quine, 1960;

⁸ All pure ontological systems are logical systems (Jacquette, 2002, p.xiii; see e.g., Foschini, 2013).

⁹ Incommensurability is not understood here in the same way as the incommensurability of paradigms or scientific theory in the works of Khun (1962) or Feyerabend (e.g., Ryan, 2002; for the incommensurability of paradigms, see also Sankey, 1993; Oberheim and Hoyningen-Huene, 2018; Bird, 2000). Instead, incommensurability is understood here as a general sense of not being entirely comparable according to some criteria.

Eklund, 2008; Berto and Plebani, 2015). Peter Van Inwagen (1998), the originator of the term, posited that the role of meta-ontology is to clarify the subject of ontology and explain how ontological claims can be interpreted. Francesco Berto and Matteo Plebani describe meta-ontology in terms of “‘meta-X’ as the inquiry on the central concepts and procedures of discipline X” (Berto and Plebani, 2015, p.13), where “X” refers to ontology in this case. Meta-ontology asks what a philosopher means when he asks ontological questions or questions about ontology (Eklund, 2008; Turner, 2014). Furthermore, meta-ontology is also concerned with ontological commitments and the truth conditions for a given ontological theory (Van Inwagen, 1998). From this perspective, meta-ontology would inquire as to what kind of “things” an ontological theory (i.e., ontology) is committed to.¹⁰

Ontological commitment and truth condition are key terms for defining meta-ontology, and they are critical for differentiating between ontology and meta-ontology (Turner, 2014). On the conceptual level, the ontological commitment denotes what kind of ontology (i.e., what exists, what is) a system or a theory is supposed to represent. Gibson (2009, p.631) states that “theory is ontologically committed to an object only if that object occurs in all the ontologies of that theory.” Thus, ontology may be committed to the existence of numbers, planets, subatomic particles, ghosts, values, ethics, and so on, anything as long as these “things” are recognized in “all the ontologies of that theory” (e.g., Eklund, 2008). While ontology is about existence, “what is”, the ontological commitment is about representation of “what is” (Smith, 1998) and the capacity to represent it.

Moreover, ontological commitment also includes verifying the criteria for this ontological commitment (i.e., verifying the truth of

¹⁰ We do not discuss Quine’s meta-ontology as being specific to Quine’s concept of ontology that is not considered here.

its existential claims). Ontological commitment can be reduced to the simple claim that “A man is committed to the truth of whatever he asserts” (Searle, 1969, p.112).¹¹ For some, however, particularly professional ontologists, Searle’s take on ontological commitment seems perfunctory, but in the AI context, it provides a simple (i.e., operational) means for judging the scope of an AI ontology. We do not dispute the robustness of Searle’s claim on ontological commitment; here we take it as a guide in practical applications.

Truth conditions are what makes an ontology correct, such that a “theory is ontologically committed to an object only if that object occurs in all the ontologies of that theory” (Gibson, 2009, p.631). This statement was rephrased by Rayo (2007) into the following claim: “[...] for a sentence to carry commitment to Fs is for the sentence’s truth to demand of the world that it contain Fs.” Gibson’s claim therefore comes from a philosophical perspective. It requires that a theory (of ontology) can be expressed in sentences (of any language), so this truth condition boils down to an agreement between the theory and the world in question (the correspondence theory of truth¹²), which may be the real world (as is the case of AI_E systems) or an imaginary constructs or virtual realities whatever the domain of ontology is.

¹¹ See the critique of Searle by Inwagen (1991).

¹² We are not going here into the discussion of the correspondence theory of truth or truth in general. We assume that for the systems (natural or artificial) “being in the world” (no Heideggerian connotations) there must be some relation, however tentative and limited holding between them and the real world they are immersed in, the relation of truth. This relation in a case of evolution and synthetic systems is verified by systems’ operational success (see for example the discussion of truth relation in Bird, 2000). The correspondence theory of truth seems to trouble only philosophers of the anti-realist, skeptical persuasion, not computer scientists or evolutionary biologists (Bird, 2000). “Operational success” in applied sciences may be thought as playing the role of empirical proof in naturalized epistemology (Bird, 2000, p.263).

AI_E Paradigm

The term “AI paradigm” is used in many AI-related papers, discussions, and articles, with it carrying various meanings, so there is no agreement about what it should denote. In principle, authors simply define a “paradigm” in a way that suits their narrative, thus confirming Ian Hacking’s prediction that the term would become banal following Kuhn’s publication (Hacking, 2012). Thus, when we talk about AI paradigm we need to show precisely, to avoid misinterpretations, what we are talking about and how our definition of paradigm differs from, or is similar to, from other definitions.

Schopman (1986) suggests that AI has not developed a specific paradigm, claiming that “[. . .] no computational paradigm has yet been produced: there is no single generally accepted way to do AI” (Schopman, 1986, p.6). Čaplinskas (1998), however, defines three AI paradigms: the behaviorist paradigm, the agent paradigm, and the artificial life paradigm. Norvig (1992; 2002), meanwhile, associates the term with Lisp programming to express the paradigm of Artificial Intelligence Programming as being equivalent to a programming approach. Next, Cristianini (2014) distinguishes four AI paradigms: data-driven AI, statistical AI, knowledge-driven AI, and reasoning-and-search-based AI. Without much explanation as to why, Leary (2017) claims that the new Google AI paradigm is machine learning, while in a blog post titled “The AI Paradigm Shift,” Richardson (2018) refers broadly to the AI paradigm as an approach to engineering AI systems, such as deep learning (DL), machine learning (ML), natural language processing (NLP), and robotics. (Hernández-Orallo et al., 2020, p.2522), meanwhile, claims that the concepts of AI paradigms

have been used to denote “broad families of technical or conceptual approaches: ‘symbolic’ vs ‘connectionist’, reasoning vs learning, expert systems vs agents.”

Much like Richardson, Romero (2021) refers to deep learning and machine learning as AI paradigms, while Yalçın (2021) refers to symbolic (i.e., a human-readable, symbolic representation of problems, logic, and search) and sub-symbolic (i.e., an implicit representation derived from experience-based learning with no symbolic representation of rules and properties) representations as AI paradigms. Villar et al. (2021) hint at relating the AI paradigm to versions of machine-learning and deep-learning methods. Meanwhile, Luhach Kumar and Elçi Atilla (2021) in their book use the term “paradigm” in several places but in different contexts, with its meaning being variably associated with programming, the applications of AI to smart computational cyberspaces, or execution paradigms associated with computer hardware. Next, Joseph Makokha (2021) distinguishes two AI paradigms, namely an AI-based one for rule-following methods and another one based on artificial neural network constructs. The term “AI paradigm” is often used to simply denote a method for AI learning or knowledge acquisition, and this is how Yonguin Xu et al. (2021) use this term to denote deep-learning approaches, such as supervised, unsupervised, and reinforced learning.

Several studies have posited that the current AI systems use two broad, conceptual constructs, namely the symbolic and the sub-symbolic (e.g., Searle, 1998; Harvey, 2013; Neapolitan and Jiang, 2018; Mitchell, 2019; Smith, 2019; Cole, 2020; Russell and Norvig, 2020; Wooldridge, 2021). Thus, the symbolic paradigm reflects symbolic representations of a priori defined concepts that may be implemented in various programming environments, while the sub-symbolic paradigm relates to less clearly defined concepts, such as

the multi-dimensional probability weights on the connections within an artificial neural network. An approach under the sub-symbolic paradigm would therefore be implemented using one of the various machine learning (ML) technologies. These two paradigms can also be fused into a neuro-symbolic paradigm (e.g., Bader and Hitzler, 2005; Garcez, Gori et al., 2019; Garcez and Lamb, 2020; Kautz, 2022) that combines symbolic and sub-symbolic elements. As pointed out earlier, the term “AI paradigm” does not point to a single software or hardware solution (see Minsky, 1991; Russell and Norvig, 2020).

In this paper, we follow the example of Searle, Harvey, Neapolitan and Jiang, Smith, Cole, Wooldrige, and Russell and Norvig in using the term “AI paradigm” to denote a broad, conceptual construct that underlies AI systems. AI paradigm as is here defined does not imply a specific implementation. The AI paradigm therefore allows for multiple implementations, formal structures, representations, programming methods, and processing algorithms,¹³ with these all belonging to a single paradigm.

The meta-ontology of AI_E systems

We have concluded that in philosophy, meta-ontology is the study of what ontology is all about. In other words, it is a study of, or about, ontology, as well as what it investigates; after Berto and Plebani “‘meta-X’ as the inquiry on the central concepts and procedures of discipline X”, where “X” refers to ontology in this case (Berto

¹³ We use the term algorithm in the general sense of a procedure that can be conceptualized and implemented in a computer. This use follows Knuth’s definition of a computational method as being “A procedure that has all of the characteristics of an algorithm except that it possibly lacks finiteness may be called a computational method” (Knuth, 2005, p.5).

and Plebani, 2015, p.13). We also concluded that ontological commitment and truth conditions represent the key concepts of meta-ontology, and these terms are critical for differentiating between ontology and meta-ontology.

The meta-ontology of AI_E systems retains the core philosophical meaning (i.e., what ontology is committed to), but it diverges from the philosophical concept in the details, because it assumes the perspective of AI_E systems. When we say it “retains the core philosophical meaning,” we mean that meta-ontology of AI_E systems is concerned with ontology, or is about ontology (i.e., the philosophical meaning of meta-ontology). Contrary to the use in philosophy, however, the meta-ontology of the AI_E focuses not on a theory of ontology, as meta-ontologies in philosophy do, but rather on what the AI_E represents (i.e., the real world).

We also said that we are concerned with studying the meta-ontology of the AI_E paradigm, which is the broad all-encompassing conceptual construct that underlies AI_E systems, rather than any particular realization of it. Indeed, we assume that most realizations of AI_E systems have the same foundational assumptions; what we denote as AI_E systems paradigm, so what we can conclude about the ontology of AI_E paradigm will also hold for its realizations. Different paradigms (from the one assumed here) of AI_E systems are logically possible. But we limit this discussion to assumptions, formulated by Smith, that AI_E systems should possess if they are to mimic humans’ ability to cope with the real world (i.e., have human-level intelligence).

Summing this all up, the meta-ontology of the AI_E paradigm¹⁴ is concerned primarily with what is (or about) the ontology of AI_E systems, what its ontological commitment is, and what its truth condition is.

To make our claims about the ontology of AI_E systems more specific, we employ Smith's postulates for the AI_E paradigm (Smith, 2019, p.44). As we mentioned earlier, Smith's AI_E paradigm is not purely ontological, but it does commit AI_E systems to certain ontology. Smith posited that for an AI system to match human-level intelligence, it needs to be embodied, embedded, extended, and enactive. More specifically, embodied means that an AI_E agent's representation of the real world accounts for its body's position, size, senses, and movement, such that the body plays a critical role in shaping the "mind" and its internal representation of the world (i.e., embodied cognition). Extended, meanwhile, implies that the AI_E agent's representation of the real world accounts for the AI_E system's mind and body as part of the cognition process (i.e., a co-creating ontology). (For more about the discussion of embodied and extended cognition, see, for example, the works of Varela (1991), Clark and Chalmers (1998), Anderson (2003), Pfeifer and Iida (2004), Rupert (2009), Rowland (2010), Shapiro (2010), Wheeler (2011), Kiverstein (2018), Bermúdez (2020), and Paul (2021)). Next, the embedded condition refers to the AI_E system being aware of the context surrounding a situation, which should be accounted for in its ontology (e.g., Hutchins, 1995; Pouw, van Gog and Paas, 2014). Finally, being enactive means that an AI_E agent fully participates in actions, both in mind and body (e.g., Varela, Thompson and Rosch, 1991, p.175; Klein, Moon and Hoffman, 2006;

¹⁴ The discussion of the meta-ontology of the AI_E paradigm is based on ideas proposed by Brian Cantwell Smith (1998; 2019) and the works of Minsky (1991), Dreyfus (2016), Mitchell (2019), Roitblat (2020), and Wooldridge (2021).

Froese and Ziemke, 2009; “the brain is conceived as participating in the action” Gallagher et al., 2013; Di Paolo and Thompson, 2017; Hutto and Myin, 2017; Newen, De Bruin and Gallagher, 2018; Smith, 2019; Newen, Bruin and Gallagher, 2020; Käufer and Chemero, 2021; Shapiro and Spaulding, 2021; “enacted AI” Shin, 2021; Hipólito and van Es, 2022).

Still, it is not obvious what ontology is implied by these requirements. As well as, it is not obvious, as to how we should translate the four features of this AI_E paradigm into meta-ontological requirements; Smith neglects to offer any suggestions here (Smith, 1998; Mitchell, 2019). Thus, we reformulated Smith’s claims about the ontology of AI_E systems into four meta-ontological theses that appear to fill the ontological lacuna in his specifications. Indeed, they would seem to be necessary for AI_E systems to be embodied, embedded, extended, and enactive. They are:

- T1. The ontological commitment of AI_E is to the real world, the world of a human agent.
- T2. The truth condition of the ontology of AI_E is not consistency with ontological theory but rather the real world.
- T3. The truth condition of AI_E is verified through the operational success of an AI_E system.
- T4. The ontology of the AI_E paradigm must account for the dynamic environment of the real world.

(T1) The ontological commitment of AI_E is to the real world,¹⁵ world of a human agent. The world for AI_E is the same reality that a human actor would exist in. We could say that AI_E ontology is a partial ontology as opposed to one that covers the entire world,

¹⁵The term explained earlier in the text. See footnote 7.

so it is about a state of affairs, by which we mean a local, temporal, dynamic (as described by Minsky) reality of the everyday world. This partial ontology does not attempt to create a comprehensive ontology of existence but rather account for what exists, together with the state of affairs,¹⁶ in the part of the actual world that is relevant to an AI_E agent, we may say, agent-relevant ontology.

AI_E ontology is therefore not a theory about what exists, abstract objects, possible worlds, and maximal worlds (e.g., Forbes, 1992; Textor, 2021). Indeed, rejecting models or theories about the world may be beneficial, as Brooks suggests (in her ontology of everyday objects): “When we examine very simple level intelligence we find that explicit representations and models of the world simply get in the way. It turns out to be better to use the world as its own model” (Brooks, 1991). For instance, the ontology of the AI_E paradigm may not have to account for subatomic particles, quantum physics, or imaginary objects, so it does not have to resolve Russell’s table paradox (Russell, 1912); it does not have to account for these or similar objects as it is an ontology of everyday world we live in (it is Minsky’s commonsense reality, or Baker’s the world of ordinary things, or “the world of medium-sized objects” (Baker, 2007, p.18)).

Accordingly, the ontological commitment of the AI_E paradigm is whatever an AI_E system can assert about the world (i.e., the entities, relations between them, etc.) given its paradigm. AI_E systems face the real world, so they are committed to things within their context, and they need to recognize reality’s features. Thus, the ontological

¹⁶ The term “state of affairs” is used in the sense employed by Jacquette (2002).

commitment in AI_E systems that seek to mimic our own ontological commitment must be geared toward recognizing the reality with limited *a priori* suppositions.¹⁷

(T2) The truth condition of AI_E is not consistency with ontological theory but rather the real world. The truth condition of AI_E systems does not depend on theory, and it is not committed to the truth of a sentence because there are no sentences or collections of them, as there is no *a priori* ontological theory defining the ontology of AI_E systems. The truth condition of AI_E systems' ontology is therefore not consistency with a closed conceptual schema or ontological theory but rather with reality, with "what is the world like" (Smith, 2019, p.57).

The truth condition of the AI_E paradigm shares some similarities with the truth condition of philosophy, which states that "theory is ontologically committed to an object only if that object occurs in all the ontologies of that theory" (Gibson, 2009, p.631). Or, in an alternative formulation by Rayo (2007) rephrased as follows: "for a sentence to carry commitment to Fs is for the sentence's truth to demand of the world that it contain Fs." However, the truth condition of AI_E systems requires AI systems to "deal with reality as it actually is—not in the way our language represents it as being" (Smith, 2019, p.34), i.e., the truth condition of AI_E systems does not have to satisfy any sentences; understood as the ontological claims expressed in some form of language.

¹⁷ See also Käufer and Chemero's (2021, p.220) discussion of Heideggerian AI, which is very similar to Baker's ontology of ordinary things (Baker, 2007) and a critique of the representational approach to the world.

(T3) The truth condition of AI_E systems is verified through operational success. The truth condition of AI_E systems is not concerned with computational representations of reality in the form of structures, data constructs, or computational concepts. Objects that are registered in AI_E systems follow constitutive regularities and norms, but these are not known beforehand but instead derived and learned from the real world as the basis of being (Smith, 2019, p.103). There is no theory to go with it, so there is no criterion for truth verification that references a theory, at least if we accept that this statement is not a theory in itself. As Baker says in an article about the metaphysics of ordinary things, “the ultimate test of a metaphysical theory is... pragmatic” (Baker, 2007, p.11).¹⁸ In other words, the truth condition of AI_E systems is verified pragmatically, i.e., through operational success, because they are solely committed to the world and their actions within it (Smith, 2019, p.145). The precise meaning of operational success of engineering (including AI_E systems) or natural systems (living organisms) depend on the specific system the term “operational success” is applied to. In biological systems operational success (mostly) means survival and reproduction. In artificial systems operational success means fulfilling design objectives. It is not always obvious what is operational success even in engineering systems. For factory robots operational success is a well-defined task – like proper welding of a pin or similar. For autonomous vehicles operational success means (among other things) collision avoidance. For AI_E systems operational success is proper response/decision to situations. Of course operational success is much harder to evaluate in some cases than a welding of a pin; like it is in a case of the notorious trolley problem (see e.g., Cathcart, 2013). “Operational success” of

¹⁸ In Baker’s ontology, the pragmatic mode of verification has nothing to do with pragmatic theories of truth in philosophy.

the trolley problem is a subject of endless debates between engineers, philosophers and enthusiasts of AI probably with a limited chance of success as these groups talk past each other; philosophers see the trolley problem as ethical problems, engineers as engineering problem, and enthusiasts of AI are too emotionally engaged to be rational. As we mentioned earlier, registered/recognized objects in **AI_E ontology** create regularities and norms, but rather than being known a priori, they are derived and learned from the real world, and this provides the grounding for AI_E ontology.

(T4) The ontology of AI_E systems must account for the dynamic environment of the real world. The complex and dynamic nature of the AI_E domain was discussed by Minsky: “. . . the objects and activities of everyday life are too endlessly varied to be described by precise, logical definitions and deductions. Commonsense reality is too disorderly to represent in terms of universally valid axioms. To account for such variety and novelty, we need more flexible styles of thought, such as those we see in human commonsense reasoning, which is based more on analogies and approximations than on precise formal procedures” (Minsky, 1991, p.6).¹⁹

Thus, the reality that the ontology of the AI_E paradigm must represent is specific to a situation, because raw reality is too disorderly to represent through universally valid axioms. Indeed, the reality/ontology faced by AI_E is too complex and nuanced to be definable by a closed set of formal rules, and any attempt to do so would result

¹⁹ Minsky is obviously not the first or only person to recognize the messiness of reality, but he is one of the few AI researchers that did so in the early years of AI technology (others include, for example, Dreyfus (2016), Wooldridge (2021), Smith (2019), Božić (2020), Mitchell (2019), Roitblat (2020), and Käufer and Chemero (2021)).

in a combinatorial explosion (Inder, 1996, p.26). This combinatorial explosion barrier implies that the ontology of the AI_E paradigm must eschew any formal *a priori* decision-making procedures.²⁰

We remain unsure about how to design AI systems that implement Smith's embodied, embedded, extended, and enactive ontology (e.g., Hoffmann and Pfeifer, 2018). Nevertheless, as biological agents do have embodied, embedded, extended, and enactive ontology suited to their specific living niche, we can assume that, in principle, synthetic systems could do the same, at least to some degree and perhaps with the use of technology that may not yet exist.²¹

Conclusions

In summation, the meta-ontological claims about AI_E systems posit that the ontological commitment of an AI_E system is directed solely to the outside world. In addition, the truth condition of AI_E **systems'** ontology is not consistency with a closed conceptual schema or ontological theory but rather with the reality, with "what is the world" (Smith, 2019, p.57). It is not concerned with computational representations of reality in the form of structures, data constructs, or computational concepts. In addition, this truth condition of AI_E **systems** is verified through operational success.

²⁰ Philosophical ontology recognizes that (to some extent) the needs of AI_E ontology seem to be the "ontology of everyday life" described by (Baker, 2007). Baker describes the ontology of common objects (i.e., the "metaphysics of everyday objects"), and this ontology may provide a philosophical interpretation for AI_E ontology, but possible similarities would again require further study.

²¹ Smith's concept is similar to 4E cognition (e.g., Shapiro, 2010; Wheeler, 2011; Newen, Bruin and Gallagher, 2020). The field of 4E cognition requires a separate discussion because it lies outside the scope of this paper.

We are well aware of how natural systems engage successfully (most of the time) with the real world, and we know, at least in some sense, how they achieve this (see, for example, studies of 4E cognition (Shapiro, 2010; Wheeler, 2011; Newen, Bruin and Gallagher, 2020) or the work of Yong (2022)). To replicate the prowess of natural systems in synthetic systems, at least to some extent, we know that we need to mimic what natural systems do (i.e., engage with the real world (see e.g. Sarosiek, 2021)). In fact, we do not have any other example to follow but us and some other animals.

We also know that we must do something different to the way in which we approach AI systems now. In other words, we must change our AI paradigm, i.e., foundational assumptions about constructing AI systems (We refer again to the ideas of Minsky (1991), Dreyfus (2016), Wooldridge (2021), Smith (2019), Mitchell (2019), and Roitblat (2020).) Alas, we still do not know how to do this effectively.

We also know that we will always have a degree of incommensurability between the ontology of synthetic systems (including AI_E systems) and reality, which is in a sense explained in ft. 8, because there is also insurmountable incommensurability between the ontology of biological agents and reality. This means there will always be some shortfall between what exists and what can be comprehended by a system, whether biological or synthetic (e.g., Yong, 2022). Indeed, the ontology of a cognitive agent, whether natural or synthetic, always only partially covers reality (revisit Figure 1), because for biological systems, it is tailored to its environmental niche and continued survival, while for synthetic systems, it is oriented toward ensuring the utility of a system and the safety of those related to, on relying on, this system. The best we can do is to minimize this incommensurability gap, once we realize that it exists, by optimizing a system to suit a specific environment.

Indeed, the meta-ontological lesson from nature is not that organisms strive to match their ontology with **O** ontology but rather to optimize their ontology to best meet their needs (see e.g., Yong, 2022) or occupy their biological niche, although this niche is essentially what their ontology is. Is this the way to go for **AI_E** systems? Obviously, a factory robot tightening nuts and bolts does not need an **AI_E** ontology, but a robot delivering pizza in a city would require a more sophisticated ontology. And so would police robots with the license to kill patrolling the city streets (see e.g., Propper, 2022). Furthermore, robotic companions, nurses, or personal assistants may require a still higher level of **AI_E** ontology. Such robots would need to navigate the messy environment of everyday life with the sort of cleverness that we expect from their human counterparts. In other words, they need human-level intelligence with human-level ontology.

The meta-ontological perspective, in the absence of generally accepted criteria, may also be useful for defining the AI paradigm, which could then be differentiated not by computing methods, software, or theories (as it is the case now) but rather by the ability to represent the real world. Such a perspective would clearly separate symbolic, sub-symbolic, or neuro-symbolic systems from their **AI_E** peers.²²

²² Meta-ontology has been used as a differentiating criterion between ontological paradigms. For example, Eklund (2008) uses meta-ontology to differentiate between ontological paradigms, such as between robust and deflationary conceptions of ontology.

Bibliography

- Anderson, M.L., 2003. Embodied Cognition: A field guide. *Artificial Intelligence* [Online], 149(1), pp.91–130. [https://doi.org/10.1016/S0004-3702\(03\)00054-7](https://doi.org/10.1016/S0004-3702(03)00054-7).
- Bader, S. and Hitzler, P., 2005. *Dimensions of Neural-symbolic Integration – A Structured Survey* [Online]. arXiv. Available at: <<http://arxiv.org/abs/cs/0511042>> [visited on 4 January 2023].
- Baker, L.R., 2007. *The Metaphysics of Everyday Life: An Essay in Practical Realism*. 1st ed., *Cambridge Studies in Philosophy*. Cambridge: Cambridge University Press.
- Berto, F. and Plebani, M., 2015. *Ontology and Metaontology: A Contemporary Guide*. London; New York: Bloomsbury Academic.
- Bird, A., 2000. *Thomas Kuhn*. 1st ed., *Philosophy now*. Chesham: Acumen.
- Božtuć, P., 2020. Conscious AI at the Edge of Chaos. *Journal of Artificial Intelligence and Consciousness* [Online], 07(01), pp.25–38. <https://doi.org/10.1142/S2705078520500010>.
- Brooks, R.A., 1991. Intelligence without representation. *Artificial Intelligence* [Online], 47(1-3), pp.139–159. [https://doi.org/10.1016/0004-3702\(91\)90053-M](https://doi.org/10.1016/0004-3702(91)90053-M).
- Čaplinskas, A., 1998. AI paradigms. *Journal of Intelligent Manufacturing* [Online], 9(6), pp.493–502. <https://doi.org/10.1023/A:1008880017722>.
- Cassimatis, N.L., Bello, P. and Langley, P., 2008. Ability, Breadth, and Parsimony in Computational Models of Higher-Order Cognition. *Cognitive Science* [Online], 32(8), pp.1304–1322. <https://doi.org/10.1080/03640210802455175>.
- Cathcart, T., 2013. *The Trolley, or, Would You Throw the Fat Man Off the Bridge?: A Philosophical Conundrum*. New York: Workman.
- Chalmers, D.J., Wasserman, R. and Manley, D., eds., 2009. *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Clarendon Press.
- Clark, A. and Chalmers, D., 1998. The extended mind. *Analysis* [Online], 58(1), pp.7–19. Available at: <<https://www.jstor.org/stable/3328150>> [visited on 3 June 2019].

- Cole, D., 2020. The Chinese Room Argument. In: E.N. Zalta, ed. *The Stanford Encyclopedia of Philosophy*. Winter 2020. Metaphysics Research Lab, Stanford University.
- Di Paolo, E. and Thompson, E., 2017. *The Enactive Approach* [Online]. (preprint). MindRxiv. <https://doi.org/10.31231/osf.io/3vraf>.
- Dreyfus, H.L., 2016. *Skillful Coping: Essays on the Phenomenology of Everyday Perception and Action*. Ed. by M.A. Wrathall. First published in paperback. Oxford; New York, NY: Oxford University Press.
- Effingham, N., 2013. *Introduction to ontology*. Cambridge, UK; Malden, MA: Polity Press.
- Eklund, M., 2008. The Picture of Reality as an Amorphous Lump. In: T. Sider, J. Hawthorne and D.W. Zimmerman, eds. *Contemporary Debates in Metaphysics, Contemporary debates in philosophy*, 10. Malden, MA: Blackwell Pub, pp.382–396.
- Fjelland, R., 2020. Why general artificial intelligence will not be realized. *Humanities and Social Sciences Communications* [Online], 7(1), p.10. <https://doi.org/10.1057/s41599-020-0494-4>.
- Forbes, G., 1992. Worlds and States of Affairs: How Similar Can They be? In: K. Mulligan, ed. *Language, Truth and Ontology* [Online]. Dordrecht: Springer Netherlands, pp.118–132. https://doi.org/10.1007/978-94-011-2602-1_8.
- Foschini, L., 2013. *Where the "it from bit" come from?* [Online]. arXiv. Available at: <<http://arxiv.org/abs/1306.0545>> [visited on 4 January 2023].
- Froese, T. and Ziemke, T., 2009. Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence* [Online], 173(3-4), pp.466–500. <https://doi.org/10.1016/j.artint.2008.12.001>.
- Gallagher, S., Hutto, D.D., Slaby, J. and Cole, J., 2013. The brain as part of an enactive system. *Behavioral and Brain Sciences* [Online], 36(4), pp.421–422. <https://doi.org/10.1017/S0140525X12002105>.
- Garbacz, P. and Oliver Kutz, 2014. *Formal Ontology in Information Systems*. Proceedings of the 8th international conference fois 2014, *Frontiers in artificial intelligence and applications*, v. 267. Washington, DC; Amsterdam: IOS Press.

- Garcez, A.d., Gori, M. et al., 2019. *Neural-Symbolic Computing: An Effective Methodology for Principled Integration of Machine Learning and Reasoning* [Online]. arXiv. Available at: <<http://arxiv.org/abs/1905.06088>> [visited on 5 January 2023].
- Garcez, A.d. and Lamb, L.C., 2020. *Neurosymbolic AI: The 3rd Wave* [Online]. arXiv. Available at: <<http://arxiv.org/abs/2012.05876>> [visited on 5 January 2023].
- Gibson, R.F., 2009. Ontological Commitment. In: R. Audi, ed. *The Cambridge dictionary of philosophy*. 2. ed., 11. printing. Cambridge: Cambridge Univ. Press, p.631.
- Goertzel, B., 2015. Artificial General Intelligence. *Scholarpedia* [Online], 10(11), p.31847. <https://doi.org/10.4249/scholarpedia.31847>.
- Guarino, N. and Giaretta, P., 1995. Ontologies and Knowledge Bases: Towards a Terminological Clarification. In: N.J.I. Mars, ed. *Towards Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing*. Amsterdam: IOS Press, pp.25–32.
- Guarino, N., Oberle, D. and Staab, S., 2009. What Is an Ontology ? In: S. Staab and R. Studer, eds. *Handbook on Ontologies* [Online]. 2nd ed., *International Handbooks on Information Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp.1–20. <https://doi.org/10.1007/978-3-540-92673-3>.
- Hacking, I., 2012. Introductory Essay. *The structure of scientific revolutions*. Fourth edition. Chicago; London: The University of Chicago Press, pp.6–30.
- Harvey, I., 2013. Perspectives on Artificial Intelligence: Three Ways to Be Smart. In: I. Harvey et al., eds. *SmartData* [Online]. New York, NY: Springer New York, pp.27–38. https://doi.org/10.1007/978-1-4614-6409-9_3.
- Haugeland, J., 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.
- Hernández-Orallo, J. et al., 2020. AI Paradigms and AI Safety: Mapping Artefacts and Techniques to Safety Issues. *ECAI 2020* [Online]. Ed. by G. Giacomo, pp.2521–2528. <https://doi.org/10.3233/FAIA200386>.

- Hipólito, I. and van Es, T., 2022. Enactive-Dynamic Social Cognition and Active Inference. *Frontiers in Psychology* [Online], 13, p.855074. <https://doi.org/10.3389/fpsyg.2022.855074>.
- Hoffmann, M. and Pfeifer, R., 2018. Robots as powerful allies for the study of embodied cognition from the bottom up. In: A. Newen, L. De Bruin and S. Gallagher, eds. *The Oxford Handbook of 4E Cognition* [Online]. Oxford: Oxford University Press, pp.841–861. <https://doi.org/10.1093/oxfordhb/9780198735410.013.45>.
- Hofweber, T., 2021. Logic and Ontology. In: E.N. Zalta, ed. *The Stanford Encyclopedia of Philosophy* [Online]. Spring 2021. Stanford, CA: Metaphysics Research Lab, Stanford University. Available at: <<https://plato.stanford.edu/archives/spr2021/entries/logic-ontology/>>.
- Husáková, M. and Bureš, V., 2020. Formal Ontologies in Information Systems Development: A Systematic Review. *Information* [Online], 11(2), p.66. <https://doi.org/10.3390/info11020066>.
- Hutchins, E., 1995. *Cognition in the Wild*. Cambridge, MA; London: MIT Press.
- Hutto, D.D. and Myin, E., 2017. *Evolving Enactivism: Basic Minds Meet Content*. Cambridge, MA: MIT Press.
- Inder, R., 1996. Planning and Problem Solving. In: M.A. Boden, ed. *Artificial Intelligence* [Online], *Handbook of Perception and Cognition*. San Diego: Academic Press, pp.23–53. <https://doi.org/10.1016/B978-012161964-0/50004-2>.
- Ingarden, R., 2013. *Controversy Over the Existence of the World. Vol. 1* (A. Szylewicz, Trans.), *Polish Contemporary Philosophy and Philosophical Humanities*, 6. Frankfurt am Main [etc.]: Peter Lang Edition.
- Ingarden, R., 2016. *Controversy Over the Existence of the World. Vol. 2* (A. Szylewicz, Trans.), *Polish Contemporary Philosophy and Philosophical Humanities*, 8. Frankfurt am Main [etc.]: Peter Lang Edition.
- Inwagen, P.v., 1991. Searle on Ontological Commitment. In: E. LePore and R. Van Gulick, eds. *John Searle and His Critics, Philosophers and their critics*. Oxford: Blackwell, pp.345–358.
- Jacquette, D., 2002. *Ontology* [Online]. Routledge. <https://doi.org/10.4324/9781315710655>.

- Käufer, S. and Chemero, A., 2021. *Phenomenology: An Introduction*. Second edition. Cambridge, UK; Medford, MA: Polity Press.
- Kautz, H.A., 2022. The third AI summer: AAAI Robert S. Englemore Memorial Lecture. *AI Magazine* [Online], 43(1), pp.105–125. <https://doi.org/10.1002/aaai.12036>.
- Kenny, A., 2012. *A New History of Western Philosophy, New History of Western Philosophy*. Oxford; New York: Oxford University Press.
- Kerr, G., 2022. *Aquinas: Metaphysics*. Available at: <<https://iep.utm.edu/thomas-aquinas-metaphysics/>> [visited on 5 January 2023].
- Kiverstein, J., 2018. Extended Cognition. In: A. Newen, L. De Bruin and S. Gallagher, eds. *The Oxford Handbook of 4E Cognition* [Online]. Oxford: Oxford University Press, pp.19–40. <https://doi.org/10.1093/oxfordhb/9780198735410.013.45>.
- Klein, G., Moon, B. and Hoffman, R., 2006. Making sense of sensemaking 1: Alternative Perspectives. *IEEE Intelligent Systems*, 21(4), pp.70–73.
- Knuth, D.E., 2005. *The Art of Computer Programming. Vol. 1: Fundamental Algorithms*. 3rd. Boston [etc.]: Addison-Wesley.
- Krzanowski, R. and Polak, P., 2022. Ontology and AI Paradigms. *Proceedings* [Online], 81(1), p.119. <https://doi.org/10.3390/proceedings2022081119>.
- Kuhn, T.S., 1962. *The structure of scientific revolutions. The structure of scientific revolutions*. University of Chicago Press: Chicago.
- Leary, K., 2017. *Google's AI is a "new paradigm" that unites humans and machines*. Available at: <<https://futurism.com/googles-ai-is-a-new-paradigm-that-unites-humans-and-machines>> [visited on 5 January 2023].
- Luhach, A.K., Elçi, A. and Sugumaran, V., eds., 2021. *Artificial Intelligence Paradigms for Smart Cyber-Physical Systems* [Online], *Advances in Systems Analysis, Software Engineering, and High Performance Computing*. IGI Global. <https://doi.org/10.4018/978-1-7998-5101-1>.
- Lutkevich, B., 2022. *What is Artificial General Intelligence?* Available at: <<https://www.techtarget.com/searchenterpriseai/definition/artificial-general-intelligence-AGI>> [visited on 5 January 2023].
- Makokha, J., 2021. Artificial Intelligence Paradigms and the Future of Learning: What a Partial Review of Half a Century of AI Conceptualiza-

- tion Suggests. *2021 ASEE Virtual Annual Conference Content Access Proceedings* [Online]. Virtual Conference: ASEE Conferences. <https://doi.org/10.18260/1-2--36700>.
- Marcus, G., 2022. *Artificial General Intelligence Is Not as Imminent as You Might Think*. Available at: <<https://www.scientificamerican.com/article/artificial-general-intelligence-is-not-as-imminent-as-you-might-think1/>> [visited on 5 January 2023].
- Minsky, M.L., 1991. Logical Versus Analogical or Symbolic Versus Connectionist or Neat Versus Scruffy. *AI Magazine* [Online], 12(2), pp.34–34. <https://doi.org/10.1609/aimag.v12i2.894>.
- Mitchell, M., 2019. *Artificial Intelligence: A Guide for Thinking Humans*. New York: Farrar, Straus and Giroux.
- Neapolitan, R.E. and Jiang, X., 2018. *Artificial Intelligence: With an Introduction to Machine Learning, Second Edition* [Online]. 2nd ed. New York: Chapman and Hall/CRC. <https://doi.org/10.1201/b22400>.
- Newen, A., Bruin, L.d. and Gallagher, S., eds., 2020. *The Oxford Handbook of 4e Cognition*. First published in paperback. Oxford: Oxford University Press.
- Newen, A., De Bruin, L. and Gallagher, S., 2018. 4E Cognition: Historical Roots, Key Concepts, and Central Issues. In: A. Newen, L. De Bruin and S. Gallagher, eds. *The Oxford Handbook of 4E Cognition* [Online]. Oxford: Oxford University Press, pp.3–15. [visited on 5 January 2023].
- Norvig, P., 1992. *Paradigms of Artificial Intelligence Programming: Case Studies in Common Lisp*. San Francisco, Calif: Morgan Kaufman Publishers.
- Norvig, P., 2002. *A Retrospective on 'Paradigms of AI Programming'*. Available at: <<https://norvig.com/Lisp-retro.html>> [visited on 5 January 2023].
- Oberheim, E. and Hoyningen-Huene, P., 2018. The Incommensurability of Scientific Theories. In: E.N. Zalta, ed. *The Stanford Encyclopedia of Philosophy* [Online]. Fall 2018. Stanford, CA: Metaphysics Research Lab, Stanford University. Available at: <<https://plato.stanford.edu/archives/fall2018/entries/incommensurability/>>.

- Osaba, E. et al., 2021. *Artificial Intelligence - Latest Advances, New Paradigms and Novel Applications* [Online]. <https://doi.org/10.5772/intechopen.87770>.
- Paul, A.M., 2021. *The Extended Mind: The Power of Thinking Outside the Brain*. Boston: Houghton Mifflin Harcourt.
- Perzanowski, J.W., 2015. *Rozprawa ontologiczna i inne eseje*. Ed. by J. Sytnik-Czetwertyński. Toruń: Wydawnictwo Adam Marszałek.
- Pfeifer, R. and Iida, F., 2004. Embodied Artificial Intelligence: Trends and Challenges. In: D. Hutchison et al., eds. *Embodied Artificial Intelligence* [Online]. Vol. 3139. Berlin, Heidelberg: Springer Berlin Heidelberg, pp.1–26. https://doi.org/10.1007/978-3-540-27833-7_1.
- Pouw, W.T.J.L., van Gog, T. and Paas, F., 2014. An Embedded and Embodied Cognition Review of Instructional Manipulatives. *Educational Psychology Review* [Online], 26(1), pp.51–72. <https://doi.org/10.1007/s10648-014-9255-5>.
- Propper, D., 2022. *San Francisco police proposal could allow cops to kill suspects with robots*. Available at: <<https://nypost.com/2022/11/24/san-francisco-police-proposal-could-allow-cops-to-kill-suspects-with-robots/>> [visited on 5 January 2023].
- Quine, W.V., 1960. *Word and Object*. Cambridge, MA: The MIT Press.
- Rayo, A., 2007. Ontological Commitment. *Philosophy Compass* [Online], 2(3), pp.428–444. <https://doi.org/10.1111/j.1747-9991.2007.00080.x>.
- Richardson, F., 2018. *The AI Paradigm Shift*. Available at: <<https://becoming-human.ai/the-ai-paradigm-shift-53fa07ae3ab2>> [visited on 5 January 2023].
- Roitblat, H.L., 2020. *Algorithms Are Not Enough: Creating General Artificial Intelligence*. Cambridge, MA: The MIT Press.
- Romero, A., 2021. *Unpopular Opinion: We'll Abandon Machine Learning as Main AI Paradigm*. Available at: <<https://towardsdatascience.com/unpopular-opinion-well-abandon-machine-learning-as-main-ai-paradigm-7d11e6773d46>> [visited on 5 January 2023].
- Rowlands, M., 2010. *The New Science of the Mind: From Extended Mind to Embodied Phenomenology* [Online]. The MIT Press. <https://doi.org/10.7551/mitpress/9780262014557.001.0001>.

- Rupert, R.D., 2009. *Cognitive Systems and the Extended Mind* [Online]. 1st ed. Oxford; New York: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195379457.001.0001>.
- Russell, B., 1912. *The Problems of Philosophy* [Online]. New York : H. Holt. Available at: <<http://archive.org/details/problemsofphilo00russuoft>> [visited on 5 January 2023].
- Russell, S.J. and Norvig, P., 2020. *Artificial Intelligence: A Modern Approach*. Fourth edition, global edition, *Pearson series in artificial intelligence*. Harlow: Pearson.
- Ryan, J.G., 2002. *Feyerabend and incommensurability* [Online]. Master's thesis. Durham University. Available at: <<http://etheses.dur.ac.uk/3754/>> [visited on 5 January 2023].
- Sánchez, D.M., Cavero, J.M. and Martínez, E.M., 2007. The Road Toward Ontologies. In: R. Sharman, R. Kishore and R. Ramesh, eds. *Ontologies: A Handbook of Principles, Concepts and Applications in Information Systems* [Online]. New York: Springer, pp.3–20. Available at: <http://archive.org/details/springer_10.1007-978-0-387-37022-4> [visited on 5 January 2023].
- Sankey, H., 1993. Kuhn's Changing Concept of Incommensurability. *The British Journal for the Philosophy of Science* [Online], 44(4), pp.759–774. <https://doi.org/10.1093/bjps/44.4.759>.
- Sarosiek, A., 2021. The role of biosemiosis and semiotic scaffolding in the processes of developing intelligent behaviour. *Philosophical Problems in Science (Zagadnienia Filozoficzne w Nauce)* [Online], (70), pp.9–44. Available at: <<https://zfn.edu.pl/index.php/zfn/article/view/535>>.
- Schopman, J., 1986. Artificial Intelligence and Its Paradigm. *Zeitschrift für allgemeine Wissenschaftstheorie / Journal for General Philosophy of Science* [Online], 17(2), pp.346–352. Available at: <<https://www.jstor.org/stable/25170750>> [visited on 5 January 2023].
- Searle, J.R., 1998. *Mind, Language and Society: Philosophy in the Real World*. New York: Basic Books.
- Searle, J.R., 1969. *Speech Acts: An Essay in the Philosophy of Language* [Online]. 1st ed. Cambridge University Press. <https://doi.org/10.1017/CBO9781139173438>.

- Šekrst, K. and Skansi, S., 2022. Machine learning and essentialism. *Philosophical Problems in Science (Zagadnienia Filozoficzne w Nauce)*, (73), pp.171–196.
- Shapiro, L., 2010. *Embodied Cognition* [Online]. London: Routledge. <https://doi.org/10.4324/9780203850664>.
- Shapiro, L. and Spaulding, S., 2021. Embodied Cognition. In: E.N. Zalta, ed. *The Stanford Encyclopedia of Philosophy* [Online]. Winter 2021. Stanford, CA: Metaphysics Research Lab, Stanford University. Available at: <<https://plato.stanford.edu/archives/win2021/entries/embodied-cognition/>>.
- Sharman, R., Kishore, R. and Ramesh, R., eds., 2007. *Ontologies: A Handbook of Principles, Concepts and Applications in Information Systems* [Online]. New York: Springer. Available at: <http://archive.org/details/springer_10.1007-978-0-387-37022-4> [visited on 5 January 2023].
- Shin, D., 2021. Embodying algorithms, enactive artificial intelligence and the extended cognition: You can see as much as you know about algorithm. *Journal of Information Science* [Online], p.016555152098549. <https://doi.org/10.1177/0165551520985495>.
- Skuza, A., 2020. *What is artificial general intelligence and who builds it?* Available at: <<https://arekskuza.com/the-innovation-blog/artificial-general-intelligence/>> [visited on 4 January 2023].
- Smith, B.C., 1998. *On the Origin of Objects*. 1st paperback, A Bradford book. Cambridge, MA: The MIT Press.
- Smith, B.C., 2019. *The Promise of Artificial Intelligence: Reckoning and Judgment*. Cambridge, MA: The MIT Press.
- Staab, S. and Studer, R., eds., 2009. *Handbook on Ontologies* [Online]. 2nd ed., *International Handbooks on Information Systems*. Berlin; Heidelberg: Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-540-92673-3>.
- Stróżewski, W., 2004. *Ontologia, Kompendia Filozoficzne*. Kraków: Aureus; Znak.
- Studer, R., Benjamins, V.R. and Fensel, D., 1998. Knowledge engineering: Principles and methods. *Data & Knowledge Engineering* [Online], 25(1), pp.161–197. [https://doi.org/10.1016/S0169-023X\(97\)00056-6](https://doi.org/10.1016/S0169-023X(97)00056-6).

- Swar, A., Khoriba, G. and Belal, M., 2022. A unified ontology-based data integration approach for the internet of things. *International Journal of Electrical and Computer Engineering (IJECE)* [Online], 12(2), p.2097. <https://doi.org/10.11591/ijece.v12i2.pp2097-2107>.
- Textor, M., 2021. States of Affairs. In: E.N. Zalta, ed. *The Stanford Encyclopedia of Philosophy* [Online]. Summer 2021. Stanford, CA: Metaphysics Research Lab, Stanford University. Available at: <<https://plato.stanford.edu/archives/sum2021/entries/states-of-affairs/>>.
- Thomasson, A.L., 2015. *Ontology Made Easy*. Oxford: Oxford University Press.
- Turner, J., 2014. Metaontology. In: Oxford Handbooks Editorial Board, ed. *The Oxford Handbook of Topics in Philosophy* [Online]. Oxford University Press, pp.1–18. <https://doi.org/10.1093/oxfordhb/9780199935314.013.25>.
- Van Inwagen, P., 1998. Meta-ontology. *Erkenntnis* [Online], 48(2/3), pp.233–250. <https://doi.org/10.1023/A:1005323618026>.
- Van Inwagen, P., 2009. *Metaphysics* [Online]. 3rd. Boulder, CO: Westview Press. Available at: <<http://site.ebrary.com/id/10271914>> [visited on 5 January 2023].
- Varela, F.J., Thompson, E. and Rosch, E., 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge; London: The MIT Press.
- Wheeler, M., 2011. Embodied cognition and the extended mind. In: J. Garvey, ed. *The Continuum Companion to Philosophy of Mind, Continuum companions*. London; New York: Continuum Press, pp.220–238.
- Wooldridge, M., 2021. *The Road to Conscious Machines: The Story of AI*. London, UK: Penguin.
- Xu, Y. et al., 2021. Artificial intelligence: A powerful paradigm for scientific research. *Innovation* [Online], 2(4), p.100179. <https://doi.org/10.1016/j.xinn.2021.100179>.
- Yong, E., 2022. *An Immense World*. London: The Bodley Head.