

“All AI’s Are Psychopaths”? Reckoning and Judgment in the Quest for Genuine AI

Larry J. Crockett

Departments of Computer Science and Philosophy

Augsburg University, Minneapolis MN, USA

crockett@augsborg.edu

Abstract: Seventy years ago Turing’s celebrated 1950 paper suggested a test to identify machine intelligence. Turing’s “imitation game” was inspired by a party game in which people judge the quality of human impersonation—notably, the test is remarkably keyed to human faculties and behavior. Judgment thus played a pivotal role in defining machine intelligence from the beginning of modern AI. This paper argues that judgment should be distinguished from reckoning, which has to do with the manipulation of data and data structures such that speed is the pivotal measure. Judgment can only originate in embodied systems so that engagement with the real world is central. Moreover, there is nothing on the horizon—scientifically or technologically—that supports the claim that we can get to “full-scale” judgment by means of better or faster reckoning. But we have been here before. Plato and Descartes developed a rationalism emphasizing innate ideas, which corresponds to “first-wave” AI, with its emphasis on logic and defining intelligence in symbolic terms. Rejecting rationalism, Locke developed an empiricism emphasizing that we begin as “blank slates,” which anticipated “second-wave” AI, neural nets, machine learning and especially big data. Kant synthesized rationalism and empiricism in his “Copernican revolution,” a strategy that eventually shaped the thinking of both Einstein and Gödel. This paper explores some possible Copernican ideas needed for a third-wave AI enabling the emergence of general artificial intelligence. Such a distinctive third wave would be explicitly oriented to ontology, so that it becomes fertile ground for the emergence of the judgment that is a necessary component of full, general intelligence.

Keywords: second-wave AI, big data, machine learning, ontology, reckoning, judgment

1. Introduction: “All AI’s Are Psychopaths”

Dixon’s (2020) “All AI’s Are Psychopaths” is so obviously titled to elicit web clicks, it could be easily dismissed. Dixon gives us a thumbnail definition of “psychopath”: a person who lacks the ability to discern the moral implications of choices. Dixon’s claim is that artificially intelligent agents are, in fact, psychopathic in exactly this sense. The root problem is that they fail to understand “how things interact.”

Marcus and Davis concur that AI “hasn’t been on the right track” (2019, 9). Figure 1 illustrates how first-wave AI (GOF AI, or “Good Old-Fashioned Artificial Intelligence”) logicized Cartesian rationalism and dismissed the relevance of learning, while second-wave AI emphasizes learning, much better neural nets, and big data—essentially, computational Lockean empiricism. Ontology is notable for its absence in both waves. The accountability of a representation lies in its correspondence to the world, not in the sophistication of its data structures or its speed. Central to intelligence is the optimization of representations so they are recursively accountable to the evolving complexity of the world.

But ontology remains a hard sell. Chin (2019) reports that Argonne Laboratory is the first to deploy the Cerebras CS-1 system so “we have dramatically shrunk training time across neural networks, allowing our researchers to be vastly more productive to make strong advances across deep learning research.” But have the two waves of AI that have had less success than hoped (even with some notable successes) been stymied principally by not enough logic, too little speed, too little training and too little data?

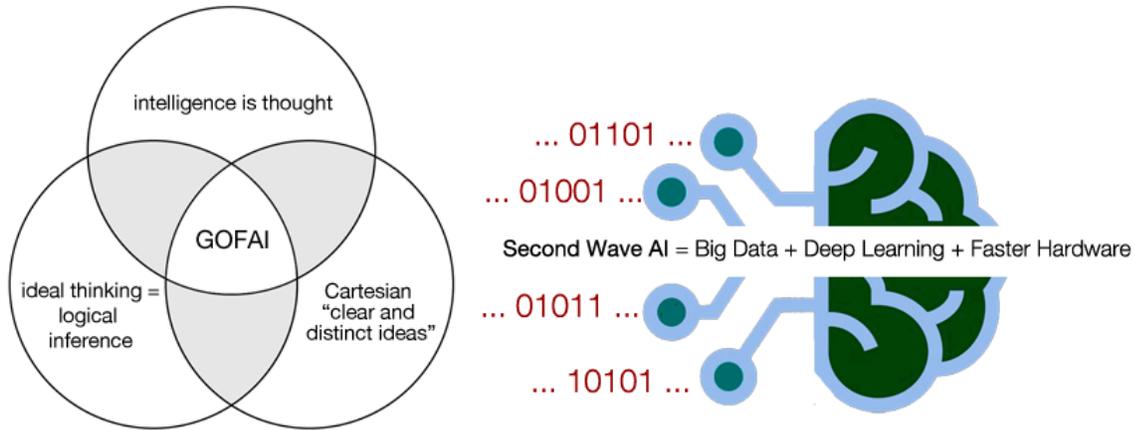


Figure 1: First-Wave AI (“GOFAI”) Logic Centric, Second-Wave AI Data Centric Shortchange Ontology

The claim of this paper is that the missing ingredient is judgment. General AI requires more than the reckoning which current computers do with great speed. It was evident in Turing (1950) that empathy and aesthetic awareness are necessary components to a machine capable of behavior indistinguishable from that of a human in his imitation game. As a result, judgment must now be recognized as the principal goal of general artificial intelligence, yielding reasoning on a par with competent adult people.

1.1 Data-Centric AI Will Not Yield General AI

Marcus and Davis (2019, 25) argue that “the approaches we have now won’t take us there” since there are five ways the human brain is superior to current AI: (1) understanding language; (2) understanding the world; (3) adapting flexibly to new circumstances; (4) learning new things quickly; and (5) reasoning effectively with incomplete information. More generally, the current “obsession with blank-slate machines ... driven purely from data rather than knowledge, is a serious error.” Just as empiricism failed in philosophy, since we cannot get causation from the regular conjunction of one datum with another, so second-wave, data-centric AI will not yield general AI.

Of course data are important and B. Smith (2019) argues that second-wave AI unwittingly points to the need for ontology. Data are a *necessary condition* for AI, and first-wave AI (GOFAI), with its dismissal of learning in favor of programming, failed to understand this. But data are not *sufficient* for AI, no matter how quickly acquired or how efficiently they are used in machine learning training; a necessary condition cannot be made into a sufficient condition by accelerating it. For example, as Marcus and Davis (2019, 15) detail, machine learning rapidly manipulates the data of very large training sets but the learning is often brittle when applied to new problems.

1.2 Deep Learning with Lots of Data

Deep learning assumes two fundamental ideas: (1) hierarchical pattern recognition, using nodes resembling neurons and (2) learning which strengthens weights in the neural net such that it is capable of “training.” Such training associates specific inputs with outputs. By using “back propagation” (a matter of adjusting weights in a trial-and-error process), significant reliability is sometimes achieved. Deep learning involves training networks of four or greater layers so as to optimize use of faster GPUs specialized for use in AI. To the point, first-wave AI “rules of thumb” are replaced by systems trained with large data sets. In terms of machine translation, “you see a constant uptick of translation quality and speed,

and we only see this [continuing] until the system is as good as a human in communicating information from one language to another” (Hsu 2016).

This is exactly the claim Marcus and Davis (2019) deny. While language translation using neural nets is much better than first-wave AI, it is still far from a flexible intelligence that can fully handle natural language, saturated as it is with human culture and a respectable understanding of how the world works. Marcus and Davis (2018, 55) argue there is no way to: (1) represent causal relationships; (2) make effective logical inferences; and (3) integrate abstract knowledge into the neural nets trained with data sets. As portrayed in Fig. 2, neural nets are notably successful in image identification. But they are poor at explanation. Notably, there are no causal inferences from typical to novel usages that require an understanding of causation, inference, and the theoretical knowledge generated by science.

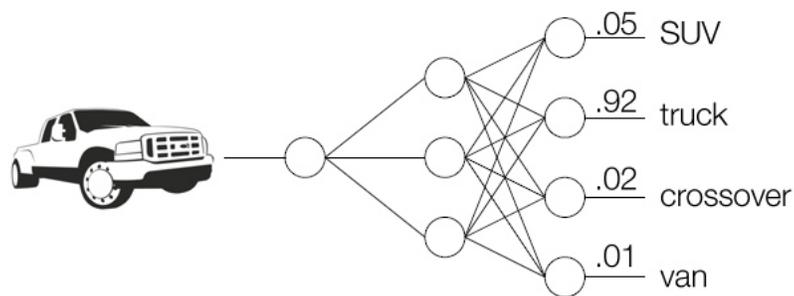


Figure 2. Neural Nets Give Correlations But No Causal Relationships

1.3 The Problem of Induction

Hume (1739) famously argued that causation is required to convert sensory data into explanation—but consists, given empiricism, of merely “constant conjunction” and psychological habit. Russell observed that if Hume’s problem cannot be solved, “there is no intellectual difference between sanity and insanity” (1948, 699). In fact, Hume’s “constant conjunction” parallels the correlations of machine learning since both face the problem of inductively justifying causation. Without a solution, perhaps Dixon’s gauntlet that AI is “psychopathic” cannot be answered.

In practice, we often seem unimpaired by the problem of induction since, for example, we frequently encounter sentences using words in unexpected ways that we can not only understand, but relish. Poetry does exactly this. It is impossible to train from a past finite set of typical usages to an infinite set of novel usages. Yet machine learning assumes that training and testing data evidence the same distributions as future data. Hume’s problem is still with us. As a result, Marcus and Davis (2019, 64) conclude, “DL is a computational *idiot savant*, good at certain kinds of perception, but little overall comprehension.” Causal reasoning is thus central to both understanding and intelligence.

2. Reckoning and Judgment

B. Smith (2019) is a major contribution to AI that distinguishes reckoning from judgment. A reckoning system computationally manipulates registrations; in current practice, it is less concerned about ontology. By contrast, judgment embodies and is oriented to the world that is registered. His thesis is that nei-

her first- nor second-wave AI will lead to genuine intelligence since computers are capable of reckoning but not judgment—the intelligence possessed by humans is “of a different order” (2019, xiii). In a word, judgment is the “standard to which thinking must ultimately aspire.” Consequently, the principal goal, the missing ingredient, in current AI systems is judgment.

Smith does not argue that it is impossible for AI systems to become capable of judgment, but he argues that judgment is different in kind from reckoning such that we cannot ramp up—extend—reckoning so that judgment emerges from faster GPUs, better algorithms, and copious data. AI may be capable of judgment but not by extending current technology.

Smith’s critique echoes Gödel’s 20th-century criticism of Hilbert’s formalist approach to mathematics, which attempted to eliminate intuition from mathematics (Davis 2018). Gödel famously proved that there are true mathematical statements that cannot be derived formally from an axiomatic system; we cannot eliminate intuition. Smith (2109) argues that terminological shifts in computer science have squeezed semantics and intuition from how we understand computation; Gödel famously disproved Hilbert but formalism lives on in the belief that reckoning can be extended into judgment without ontologically defined semantics.

What’s missing in most AI research is an “existential” commitment to the world. The problem is not technological, it is existential, ontological. But first- and second-wave AI have been preoccupied with what 17th-century Hobbes (and later Leibniz) called “reckoning.” In fact, Hobbes (1929, 98) was famed for suggesting that “reasoning is but reckoning.” Hobbes lives on in both GOFAI and second-wave AI. The reductionist “reasoning is but reckoning” stands in the way of AI reaching general intelligence.

2.1 Judgment: Accountability to the World

Notably, judgment is what GOFAI and second-wave AI lack. Smith writes that an intelligent system must hold its registrations in “perpetual epistemic and normative abeyance, prepared to abandon them in order to be accountable to the Plenum [all of reality]” (2019, 137). GOFAI was self-absorbed in its registrations. Second-wave AI improves on the accountability to the world but still fails because it assumes that reckoning can be ramped up to judgment, that data can stand in for knowledge.

But if judgment must now be the principal goal of general artificial intelligence, we need a fuller definition. First, judgment recognizes the ethical dimensions of every representation, decision and action. Second, judgment is ontologically aware, mindful of the demands of place and time, given the goals sentient beings identify. Moreover, judgment is communal, culturally aware, forged as it has been over many centuries and cultural traditions. It is both inherited and applied freshly to each novel situation—we often judge competently despite the problem of induction. While this is a tall order, Smith argues it “must be the goal of AGI eventually” (2019, xvi).

2.2 Ontology and Logic

Hobwefer (2018) suggests ontology is the study of: (1) what we are committed to; (2) what there is; (3) the most general features of what there is—metaphysics; (4) meta-ontology—what questions ontology should address. In other words, ontological questions are the most basic questions, which is nothing new, but notice the first part. Hobwefer makes “what we are committed to” first in the list. Ontology cannot be studied or represented in the abstract, it is not exhausted by *representation* or data structures any more than math is exhausted by formal symbol manipulation. Ontology references the world.

Computer scientists typically use “ontology” in a narrower sense. For example, Spyns, Jarrar, and Meersman (2002) define ontology as an “agreement about a shared, formal, explicit and partial account of a conceptualization.” In fact, the word “ontology” occurs 62 times and we might be encouraged by seeing reference to “ontological commitments”—but it turns out such commitments consist of rules about registrations.

Words can have multiple definitions and evolve over time but compromising the link with the real world and its rich complexities explains some of AI’s current limitations. Current AI systems have no commitment to the real world, embody no genuine ontology, and, this paper argues, possesses no general intelligence.

3. Trustworthy AI Not Here Yet

Oxford quantum computing pioneer Deutsch (2020) writes that AI “has made no progress whatever during the entire six decades of its existence.” At first, Deutsch seems optimistic because “everything that the laws of physics require physical objects to do can, in principle, be emulated in arbitrarily fine detail by some program.” But then Deutsch echoes B. Smith (2019) by insisting progress is “a matter of philosophy not computer science.” If this is right, we need a fundamental rethinking of the philosophy that we bring to AI.

AI is being used in recent Covid-19 research (Etzioni 2020). Curai’s machine learning, for example, promises “instant medical expertise that’s accurate, trustworthy ... and actionable.” Given the emergency Covid-19 represents, of course, it is critical that AI be trustworthy. But G. Smith argues that we “will be hard pressed to provide an explanation for the diagnosis and prescribed treatment. Patients won’t know. Doctors won’t know ... Nobody knows” (2018, 150). As a result, AI is not yet trustworthy when we need it most.

4. Fallacious Patterns in Big Data

In a widely read article, Lewis-Kraus (2016) celebrates “stunning” advances in AI, notably Google Translate which “was run on the first machines that had, in a sense, never learned to read anything at all.” The qualifier “in a sense” is the red flag. If we ask a person if a read passage was understood, we would be puzzled to hear “in a sense.” The article’s enthusiasm for a machine learning “poised to reinvent computing itself” obscures how difficult AI is and how readily we jump to premature conclusions with a little progress. Dreyfus (2012) argues “first step fallacies” have been a problem in AI for decades.

Lewis-Kraus would be surprised that G. Smith (2018) calls AI a “delusion” since “computers analyze stupid strategies,” and “computers don’t think, they obey.” Notably, people are skilled at using the familiar to make sense of the unfamiliar. Computers can use massive amounts of data to correlate similar data with typical data outcomes but not significantly dissimilar outcomes. People possess a “ceaseless flood of analogies” (2018, 26) that can be applied to dissimilar situations; just as importantly, people readily recognize when they do not apply. Smith draws a stark conclusion: “For these reasons, it is a misnomer to call computers intelligent.”

G. Smith (2018, 40) challenges the celebrated successes of big data and deep learning directly. Statistical patterns in a data set are especially liable to misinterpretation and misapplication. Such patterns cannot begin to match people with years of experience that aid in recognizing, understanding and anticipating. Without real-world experience, an AI system must draw inferences from statistical patterns. In a word,

big data do not supplant common sense. Moreover, people can often recognize bad or misleading data that computers miss (2018, 58).

4.1 Patterns in Data and Real-World Experience

Deep neural networks identify patterns but without real-world experience, years of learning a language, and thousands of social exchanges that convey inherited cultural wisdom. That is, the data they manipulate and the conclusions they draw have no meaning to them. The goal of formalism in the history of mathematics was to exclude intuition and meaning in favor of pure symbol manipulation. Gödel famously saw this goal as the most basic of mistakes. Since AI systems do not understand concepts, they cannot understand words or language use more generally. Improving correlations between inputs and outputs does not capture what we mean by “understand.” G. Smith paraphrases Etzioni, “how can machines take over the world when they can’t even figure out what *it* refers to in a sentence” (2018, 45).

4.2 Correlation and Causation

Perhaps G. Smith’s most significant reminder is that data correlation is not equivalent to causation. Smith praises the speed of statistical software but points out that software cannot “tell us whether the first factor causes the second, the second the first, or some third factor causes both” (2018, 64). Programming skepticism into a computer is difficult since skepticism is borne of lots of real-world experience—so computers “have no general ability to distinguish the plausible from the preposterous” (2018, 71). Devoid of real-world semantics, “meaningless” is meaningless.

The failure of big data, neural nets, and data mining to achieve general AI is a computational analogue for the failure of empiricism in philosophy of science over several centuries. It was Hume’s empiricism, as we saw above, that drove his skepticism regarding causation. Smith repeatedly points out that software, like Hume’s empiricism, is unable to distinguish the causal from the coincidental. Those who do not know philosophical history are likely to recapitulate it unwittingly.

4.3 The Perils of Data Mining

The term “data mining” has generated lots of interest but G. Smith argues it is the “most dangerous form of artificial intelligence” (2018, 78). Data mining is often used to predict behavior such as loan defaults, a critical task in an uncertain economy. The pivotal issue is the presence of intuitive experience, even a theory, to identify which patterns in data are useful predictors. The argument is simple: we believe that causally significant patterns are the exception rather than the rule; in large data sets, patterns are inevitable but most are misleading noise. Interpreting patterns is pivotal to both science and everyday living—as a result, we are careful. Smith argues that AI is not: “Big Data contain far more data and can yield far more ludicrous relationships” (2018, 117). The bigger the data, the more often we will get meaningless patterns. Formal symbol manipulation alone, we should recall, can generate neither semantics nor causal relationships

5. We Have Been Here Before: Rationalism and Empiricism

As noted above, AI is recapitulating the history of philosophy. Plato and much later Descartes developed a rationalism emphasizing innate ideas, which corresponds to first-wave AI, with its emphasis on program and logic and its skepticism about learning. By contrast, Locke’s (1689) empiricism rejected Cartesian rationalism in favor of an empiricist “blank slate.” At birth the mind is a “blank slate” without rules for processing data. Data are disclosed in sensation, and we build correlations from the ground up. The par-

allel with machine learning is obvious. Schkolne (2019) draws a parallel between machine consciousness and Locke's idea of *qualia* illustrating unappreciated parallels between the histories of philosophy and computing.

Paralleling B. Smith (2019), Popovich (2019) writes that assuming that data and data structures are purely rational and distinguishable from semantical reference to the real world is the heart of the problem of contemporary AI, since "we have to learn to interpret code and data, to understand their meaning beyond their immediate, value-free significance as code or data in isolation." If machines use only "context-free" languages, in fact, they cannot achieve general AI.

5.1 Kantian Intuition and AI

It was Kant who famously distinguished between rationalism and empiricism and who sought, by means of his "Copernican Revolution," to synthesize a common ground between the two schools of thought and thus provide a justification for empirical science and ethics. Kant's synthetic *a priori* judgments represented a new strategy that defined modern philosophy and eventually shaped the thinking of both Einstein and Gödel, with the latter's Incompleteness theorems preparing the way for recursion theory, computer science and AI. Arguably, AI is a Kantian enterprise that has forgotten the central role of intuition in Kant. We should therefore not be surprised by AI's limited progress.

6. Concepts and Prior Knowledge

Lake (2014, 28) argues that an essential part of human intelligence is concept learning but it has proven elusive in AI. Notably, "human concepts are representationally rich, useful for not only classification, but also action, explanation and imagination" (2009, 29). Young children, for instance, learn new words from just a few examples, applying them adroitly while machine training often requires thousands of examples and exhibits less ability to use them outside a narrow range. It is prior knowledge (Hume calls it "habit"), Lake suggests, that makes children so adroit.

In a similar vein, Marcus and Davis (2019, 12) argue that what's missing in current AI is the ability to reason on something new and unexpected, which is a contemporary way to invoke the long-recognized frame problem (Crockett, 1994). The result, they argue, is that current AI is capable of narrow problem solving and reasoning over a small, predictable domain, but not the full-bodied reasoning we expect from skilled, experienced people and general AI.

7. Semantics, Ontology and Judgment in Third-Wave AI

As a result, a considerable range of writers agree that we will not achieve general AI by simply extending current technology: faster CPUs, better data structures, and more data. These may improve reckoning but will not enable the emergence of judgment. As noted above, terms from more traditional philosophy, the background for logic, have been formally redefined in terms of machine behavior. So the "semantics of a program" does not refer to its relation to the world but to the program's internal relations and behaviors. Consider Fig. 3 which defines AI without philosophy.

This exclusion is exactly what Smith (2019) and Deutsch (2020) identify as key to AI progress. The classical Cartesian assumptions ingredient in first-wave AI include: the essence of intelligence is thought; the essence of thought is inference; perception is unproblematical, and the world's ontology should be understood formally as well-defined objects with formalizable, unambiguous relations with other objects.

This is to render the wet, infinitely variable, substantial world, rich in affect and history, an unrecognizable formal skeleton.

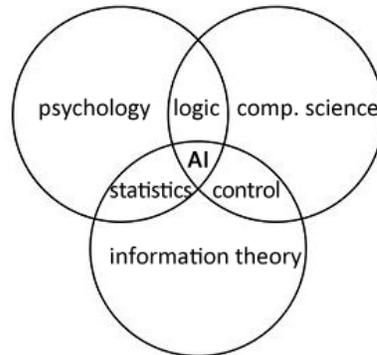


Figure 3: Defining AI Without Philosophy (Hughes and Hughes (2019))

7.1 Semantic Relations and Intelligence

“Semantics” in the formalist computer science understanding names mechanically individuated behavior of a program or a data structure. Likewise, “interpretation” can be redefined in terms of *internal* operations—not surprising since computer science has been shaped by the formalist tradition in logic. But a good case can be made that these words, as well as others, should retain their traditional denotations referring to the world. Moreover, any AI system is going to be assessed semantically—for instance, medically (“the virus is spreading”) or mathematically (“there are more integers than prime numbers”). When computer science redefines “semantics” in a mechanist, formalist way—with no necessary connection to the world, to intuition—the seeds of AI’s failure are irretrievably sown.

7.2 Semantic Relations Are Not Effective

Notably, semantic relations between a sentient being and the world are not what computer scientists would call *effective*—that is, the semantic idea of “being referred to” is certainly an important property but it is not objective in the way prized by the formal sciences. Specifically, it does not generate a signal that any instrument can detect. Daydreaming of an idyllic isle, for example, sends no detectable signals.

This point is critical because it enables us to understand how integral semantics is to intelligence. The reason we need intelligence to achieve goals, illustrated in Fig. 4, is because semantics is non-effective. That is, they make no causal contact between the world and the intentions embodied in the intentional states. Intentional acts (and semantic states) cannot make or receive effective, causal contact with those objects of the world they are about. A semantic relation, therefore, is signally detectable at neither end. Effective reach is only local. We are in contact, as B. Smith notes, with a $1/\text{radius}^2$ “spatio-temporal envelope.” The world is largely beyond our effective reach which human intelligence strives to extend so we

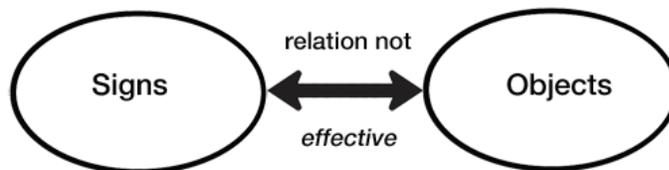


Figure 4. Semantic Relations Between Signs and Objects are not Effective

risk our ideas rather than our lives. At the heart of intelligence, as portrayed in Fig. 5, is thinking at a distance, semantically referencing a real world, which forces us to confront questions of ontology.



Figure 5. Non-effectiveness of Semantics Drives Intelligence

7.3 The Path to Genuine AI Is through Ontology

Davis, Shrobe and Szolovitz (1993) prominently anticipated B. Smith (2019), yet their prescient suggestions did not shape AI history. Their question: “How do we define *representation*?”

- (1) a surrogate structure enabling thinking before acting so that such structures die in our stead;
- (2) ontological commitments detailing what comprises the world;
- (3) a theory of reasoning consistent with these commitments;
- (4) an ability to organize information so effective inferences can be drawn;
- (5) a language capable of speaking about the world, not just our representations.

Proposition 1, “surrogate structures,” which emerged early in GOFAI, reflects Cartesian influence. Propositions 2 and 5 anticipate B. Smith’s call to re-integrate ontology into AI. Proposition 4 runs counter to early AI and anticipates big data and deep learning. Proposition 5 embodies the importance of both ontology and semantics: AI must learn from and make judgments about the world—not just its internal operations. We need a way to reference the world that reaches beyond our representations and is ontologically accountable. The way to the judgment genuine AI requires, therefore, is not around but directly through ontology.

7.4 Third-Wave AI

The goal for third-wave AI is therefore judgment, but the path to that goal is linguistic. AGI systems must learn to use language as effectively as we do since language enables the imagination needed to identify accountability to the world. Language enables ideas and not just representations. Ideas depend on data, on perception, on representations, but, contra Locke and second-wave AI, are not limited to them.

There is a profound difference between using and understanding language and simulating language use by means of program, neural net, and deep learning. The recurring theme in B. Smith (2019) is an ontological commitment to the world, which is a prerequisite to language use as opposed to language simulation. There is nothing at stake in language simulation for current AI systems; genuine language use is not needed because survival is not at stake. As illustrated by the Covid-19 pandemic, we mobilize when threatened with extinction, using language. Ideas are proposed so that the good ones thrive in our attempt to survive while the bad ones “die in our stead.” Computer systems are oblivious to questions of life, death, and existence and therefore unmotivated to reach judgment.

B. Smith (2019) is thus a call to re-engage and a re-commit to philosophical realism. If we are impressed that our experience is part of a computer simulation, we won't see the need to distinguish between ontology and registration, or to any underlying hardware—the distinction between hardware and software is trivial. Smith (2019) can be read as a call to recognize that the distinction between hardware and software is not only durable but necessary to creating judgment-infused AI. For Smith, only if we pay close attention to ontology, the real, will we be able to “hold our registrations accountable” (2019, 144).

8. Conclusion and Outlook

B. Smith (2019) should prove to be a turning point in artificial intelligence research since it opens a path through the failure of first-wave and limited success of second-wave AI. The challenge of fundamental philosophical questions, as B. Smith, Deutsch, and Marcus and Davis suggest, is paramount. We should try to integrate ontology-rich, pre-conceptual representations of deep learning with the articulated reasoning, specifically logical inferencing, that was characteristic of GOFAI. So the task is full integration such that a distinctive third, much more sophisticated wave emerges, one that is fertile ground for the emergence of a judgment that is a necessary component of general artificial intelligence.

References

- Chin, S. (2019) “Cerebras CS-1 AI computer deployed,” *Fierce Electronics* Nov. 2019.
- Crockett, L. (1994) *The Turing Test and the Frame Problem*, Ablex, Norwood N. J.
- Davis, M. (2018) *The Universal Computer*, CRC Press, Boca Raton, FL.
- Deutsch, D. (2020) “Philosophy Will Be the Key that Unlocks Artificial Intelligence,” *The Guardian* Oct. 3 2012.
- Dreyfus (2012) “A History of First-Step Fallacies,” *Minds and Machines* 22.
- Dixon, B. (2020) “All AI’s Are Psychopaths,” *Mind Matters News* 03.16.2020.
- Etzioni, O. (2020) “AI Can Help Scientists Find a Covid-19 Vaccine,” *Wired* 03.28.2020.
- Hobbes, T. (1929) *Leviathan*, Oxford University Press, Oxford, England.
- Hofweber, Thomas (2018) “Logic and Ontology,” *The Stanford Encyclopedia of Philosophy*.
- Hsu, J. (2016) “Google Translate Gets a Deep-Learning Upgrade,” *IEE Spectrum*, 03 Oct 2016.
- Hume, D. (1739) *A Treatise of Human Nature*, Oxford University Press, Oxford, England.
- Hughes, C. and Hughes, T., “The Laws of Thought and Thinking Machines,” *AI Matters*, vol. 5, no. 1.
- Lake, B. (2014) “Towards More Human-Like Concept Learning in Machines”
<https://cims.nyu.edu/~brenden/papers/LakePhDThesis.pdf>.
- Lewis-Kraus, G. (2016) “The Great A.I. Awakening,” *The New York Times*, Dec. 14, 2016.
- Locke, J. (1689) *An Essay Concerning Human Understanding*, Kenneth P. Winkler (ed.),
Hackett, Indianapolis, IN, 1996 [1689].
- Marcus, G. and Davis, E. (2019) *Rebooting AI*, Pantheon Books, New York.
- Popovich, S. (2019) “Empiricism and Interpretation in AI”
<https://redlibrarian.github.io/article/2019/06/29/empiricism-and-interpretation-in-ai.html>.
- Russell B. (1948) *Human Knowledge*, Simon and Schuster, New York.
- Schkolne, S. (2019) “Machines Demonstrate Qualia”, *Medium*, Jan. 9, 2019.
- Smith, B. (2019) *The Promise of Artificial Intelligence*, MIT Press, Cambridge, Mass.
- Smith, G. (2018) *The AI Delusion*, Oxford University Press, Oxford, England.
- Spyns, P., Jarrar, M., and Meersman, R. (2002) “Data Modelling versus Ontology Engineering.”

ACM SIGMOD Record. 31. 12-17.

Turing, A. (1950) "Computing Machinery and Intelligence," *Mind* LIX.