

Artificial Intelligence and Robotics: Contributions from the Science and Religion Forum

with Gillian K. Straine, "Artificial Intelligence and Robotics: Contributions from the Science and Religion Forum"; Michael S. Burdett, "Personhood and Creation in an Age of Robots and AI: Can We Say "You" to Artifacts?"; Mohammad Yaqub Chaudhary, "The Artificialization of Mind and World"; and David Hipple, "Encounters with Emergent Deities: Artificial Intelligence in Science Fiction Narrative."

THE ARTIFICIALIZATION OF MIND AND WORLD

by *Mohammad Yaqub Chaudhary*

Abstract. The rapid advancement of artificial intelligence (AI) has led to renewed ambitions of developing artificial general intelligence. Alongside this has been a resurgence in the development of virtual and augmented reality (V/AR) technologies, which are viewed as "disruptive" technologies and the computing platforms of the future. V/AR effectively bring the digital world of machines, robots, and artificial agents to our senses while entailing the transposition of human activity and presence into the digital world of artificial agents and machine forms of intelligence. The intersection of humans and machines in this shared space brings humans and machines into ontological continuity as informational entities in a totalizing informational environment, which subsumes both cyber and physical space in an artificially constructed virtual world. The reconstruction of mind (through AI) and world (through V/AR) thus has significant epistemological, ontological, and anthropological implications, which constitute the underlying features in the artificialization of mind and world.

Keywords: anthropology; artificial intelligence; epistemology; metaphysics; mind; neuroscience; ontology; worldview

The recent resurgence of artificial intelligence (AI) development has infiltrated virtually every sphere of discourse. Alongside this, there has recently been a resurgence of interest in the development of augmented reality (AR) technologies with the aim of providing seamless continuity between the physical world and digitally augmented overlays.

Mohammad Yaqub Chaudhary is Research Fellow at Cambridge Muslim College, Cambridge, UK; e-mail: yc@cambridgemuslimcollege.ac.uk.

This article argues that both developments should be considered together. The development of artificial agents with mental agency involves the creation of artificial worlds in which these agents exist, and these artificial worlds are subsuming increasing portions of physical space, which is leading to new shared spaces of mediated interactions between natural and artificial agents. Here, we begin by providing an overview of why AI should be considered a key area in the field of science and religion, which raises fundamental issues of interest to theologians, philosophers, and anthropologists. This will be followed by extending the discussion to what it means for an artificial agent to exist in a virtual environment, and how this environment is discontinuous with the physical world, even in the case of agents that are physically embodied in robots. We will finally discuss how we are coming to share these artificial worlds via connected technologies, which we will gradually share fully as artificial agents have more activities transposed into the physical world, and as virtual and augmented reality (V/AR) technologies, which transpose our activities into virtual worlds, become more ubiquitous.

RECENT DEVELOPMENTS IN AI

Since the victory of AlphaGo over its human counterpart in the game of Go in 2015, there have been many further technical achievements from AI research including developments in machine vision, natural language processing, speech synthesis, and learning.

The latest version of AlphaGo, named AlphaGo Zero, beat its predecessor 100 to 0 by learning the game from self-play (Silver et al. 2017), whereas the original AlphaGo was trained on a library of over 100,000 human games before it achieved superhuman level play. Here, it is worth noting how this technology has been adopted commercially with significant social implications by considering the fact that one of the first projects the AlphaGo system was applied to by Google was YouTube's video recommendation algorithm, whereby, "*The same intelligence behind the system that defeated the human world champion at the game of Go is sitting on the other side of your screen and showing you videos that it thinks will keep you viewing as long as possible*" (Williams 2018, 90). In view of such applications and the way in which AI and machine learning are finding new applications in every industry sector, there has been widespread public discourse on the social, political, and ethical implications that will emerge from the integration of autonomous intelligent machines in society.

Besides these social implications, another key area where AI is being applied is to enhance scientific research itself, as highlighted by an additional vision for DeepMind's AI systems, according to its founder, Demis Hassabis, which is to master disciplines such as "cancer, climate change, energy, genomics, macro-economics [and] financial systems," (Burton-Hill

2016). Successful applications have already been achieved in areas such as planning chemical synthesis (Segler, Preuss, and Waller 2018), and in late 2018, DeepMind announced AlphaFold, which has been shown to be able to predict the 3D structure of proteins more effectively than any previous approach (Senior et al. 2020).

The use of AI in science therefore has significant implications for the philosophy of science. The impetus for this use of AI techniques in scientific research is to “drive and accelerate new scientific discoveries,” (Senior et al. 2018), especially since in an increasing number of disciplines, the problems have become too complex to handle by traditional computational simulations. However, with new AI techniques and high-performance computation, such complexity seems to become tractable as AI offers a telescope into unfathomably large arrays of data, expanding the vision of scientists to a new universe of complexity. With this new instrument of science, AI is fundamentally altering methodologies across virtually every scientific discipline and leading to a revolution in the modern worldview no less than that inaugurated by Galileo’s use of the telescope to view the celestial sphere. A prominent example of this is the recent image of a black hole, which was an output of a project where the lead researcher was from the field of computer vision based in MIT’s Computer Science and AI Department (The Event Horizon Telescope Collaboration, 2019).

AI, PHILOSOPHY OF MIND, AND THE MIND AND BRAIN SCIENCES

Similarly, AI, artificial neural networks (ANNs), and vast quantities of data from new brain imaging techniques have become the digital optics to offer a new microscope for brain scientists to peer further inward to the anatomy of the mind, which is leading to significant implications for the philosophy of mind. The success of deep learning is based on a revival of the connectionist stream of AI research, and together with advances in the mind and brain sciences, new scientific accounts about the nature of intelligence, rationality, intentionality, volition, attention, memory, and imagination are being proposed, as well as possible solutions to the most ancient questions in philosophy about the nature of the human mind and its connection with the body, other minds, and the world.

In particular, AI research raises ontological questions about the nature of mind, world, autonomy, agency and action for humans and machines, and epistemological questions about the nature of epistemic contact between mind and world for humans on one hand, and mind and world, or rather, program and environment for machines, on the other, and the questions of what machines know, how they know, and what we may know through them.

Central questions in AI, philosophy of mind, and the mind and brain sciences that intersect directly with theological questions include the nature

of the mind and soul, how humans reason, the nature of perception, whether machines can think and exercise free will and agency, and whether a computer can simulate consciousness.

A view that has been gaining acceptance in contemporary neuroscience, is that our perception of the world is a “controlled hallucination” (Seth 2017), and neuroscientists are coming to believe that the brain creates our mental world as a model. This idea of perceiving the world through a model has also been adopted by AI researchers, and a widely accepted understanding about intelligence is that, “intelligence is just knowing a lot about the world and being able to use that knowledge flexibly to achieve goals” (Li 2018, 39).

As a consequence, intelligence is coming to be viewed as a phenomenon that may be decoupled from the phenomena of consciousness and life, and isolated for implementation in a nonbiological substrate. According to the CEO of NVIDIA, Jen-Hsun Huang, drawing inspiration from Pedro Domingos’s 2015 work titled “The Master Algorithm” (Domingos 2015), the brain may be understood as being similar to a vastly parallel computing architecture, similar to the graphics processing units (GPU) developed by the company (Huang 2016) and which have been crucial to the current revival of AI and deep learning by accelerating the matrix operations required for simulating deep neural networks (Krizhevsky, Sutskever, and Hinton 2017).

From a theological perspective, intelligence has more often been tightly coupled with consciousness and life, and monotheistic religions, such as Islam, hold intelligence to be essential in affirming the existence of God. A sound intellect is prerequisite for legal responsibility and revelation frequently urges mankind to use the intellect in contemplation, thought, reflection, investigation, and so on.

From an Aristotelian perspective, which was inherited and deliberated upon by Ibn Sina, the *mutakalimūn* and scholastic theologians, the rational soul is the basis for the distinctive human powers of intellect and will, and the power to grasp abstract concepts and universals. In this view, rationality has as its final cause the attainment of truth and knowledge, while free will has its final cause in choosing actions that accord with the truth about human purpose, nature, and essence. Hence, higher intellect means to more fully realize ultimate truth, that is, the existence of God, and the highest application of the intellect is thus to know God. Accordingly, the highest application of the will is therefore to act in accordance to this knowledge. All other human powers are subordinate to this end.

To further understand the theological significance of this turn toward a “science of mind” we must consider the historical formation of the modern concept of the mind, which has its roots stretching back over 400 years of Western philosophy.

In the book *Soul Machine*, George Makari shows how the modern concept of the mind was constructed over several centuries from Hobbes' invention of automata, Cartesian dualism, with the bifurcation of body and mind and later on, Locke's thinking matter. The understanding of the mind we have today is thus inherited from this philosophical deconstruction of the Aristotelian-scholastic tradition. Prior to this, the functions of intellect, thought, memory, reasoning, understanding, and so on were attributed to the soul itself. Hence, the invention of the mind represented the undoing of the "unifying link between nature, man and God," (Makari 2017, 7).

It has been argued that the philosophy of Descartes contributed most significantly to overhauling the earlier philosophical tradition, and hence our modern conceptions of mind and world, in three ways. The first element of Descartes' philosophy is the emergence of the modern subject and the beginning of subjectivism, through the *res cogitans*, which led to a particular type of exaggerated anthropocentrism, and a theory of knowledge in which man became the absolute foundation for the representation of reality.

According to Heidegger, this shift represented a change in the idea of truth from being a revelation to merely an adequate representation of reality or "world picture" (Heidegger 1997). In this case, the subject representing the world reduces it to a collection of calculable, controllable objects, which initiated a technological orientation that reduced the world to manipulable objects. The use of mathematics was such as to only reveal the world insofar as it is controllable. That is, the way of representing the world in the modern sciences "entraps" nature (Heidegger 1977, 21; Kureethadam 2017, 307).

The second element is the mechanistic worldview of Descartes, which introduced a new conception of the physical world, including animate beings that displaced the Aristotelian-scholastic hylomorphic conception of matter. In this case, physical entities have mechanistic properties alone, and alongside this was a mechanistic physiology for animals, such that they could be subsumed under the category of *res extensa* as beast-machines. According to the mechanistic ontology, the physical world is regarded as nonagentic, passive, noncreative, and inert, with no internal principle of agency or movement, and admits no teleology, or final cause (Kureethadam 2009, 16). All of reality, including mental reality is seen as reducible to matter when nature is reconceived of in mechanistic terms, which opened the way for technological manipulation of the animate and inanimate as "objects of detached analysis, observation and experimentation" (Broszimmer 2002, 56).

The third element was the dualistic divide between humans and the physical world. This represented a deep metaphysical dualism and an ontological division of reality into two realms, thus establishing a sharp discontinuity between humans and the rest of the physical world.

What made scientists retain an implicit dualism was the scientific advantage of dealing with nature as pure *res extensa*, which provided one substance with the essential attribute of extension, which could be known exclusively in the mode of mathematical description.

It did not matter that the other substance, *res cogitans*, with its essential attribute being awareness, could not be so clearly described; instead, its inconvenient features could be isolated. The isolation of *res cogitans* meant, “the complete detachment of external reality from what was not extended and measurable” (Jonas 2001, 54). *Res extensa* was thus a self-contained portion of reality for the universal application of mathematical analysis and provided the metaphysical justification for the mechanical materialism of modern science. Dualism did not deny the reality of nonextended entities; it merely reassigned them to a separate domain. However, life and mind, pose significant theoretical difficulties for Cartesianism and other forms of dualism.

Here, we argue that each of these three areas is currently undergoing reconfigurations in the transition from modernity to postmodernity. First, there has been a radical shift from the exaggerated anthropocentrism of Descartes, which is leading toward a mechanocentric worldview. Second, the reduction of the physical world including animate beings to pure extended matter based on a mechanistic conception of the natural world has undergone several transformations via a succession of metaphysical paradigms, most notably via cybernetics and now an informational ontology. Finally, the dualistic divide between humans and the physical world has been eliminated through the monism of information (after a succession of other monist paradigms) which has become a secular form of “spiritualist monism” (Dupuy 2013, 68) in postmodern and posthuman discourse. In what follows, we discuss how these reconfigurations are key features contributing to the artificialization of mind and world.

First, the anthropocentrism of Descartes and its replacement by a mechanocentric worldview represents a significant upheaval in epistemology and philosophy of science. We have already mentioned how machine learning and machine vision are providing the digital optics to aid scientists to uncover new insights in the vast quantities of complex high-dimensional data they are able to generate across numerous spatial and temporal orders.

For example, the relatively recent branch of neuroscience known as connectomics, which seeks to map the neural connections of brains and nervous systems to identify the relationship between structure and function in these systems (Lichtman, Livet, and Sanes 2008; Bargmann and Marder 2013), has recently been significantly boosted by whole brain electron microscopy (EM) data of the *Drosophila* brain. The acquisition of this new data set in 2017, which amounts to approximately 21 million images and 106TB of EM data, was made possible by a custom high-throughput EM platform (Zheng et al. 2017).

Compared to other techniques, EM is able to resolve all neurons and synapses in brain tissues; however, it has been technically challenging to generate significant EM volumes. Hence, a second-generation Transmission EM (TEM) camera array system was developed by the authors, incorporating automatic high-speed sample handling and high frame rate TEM imaging (by exposing samples to a higher electron dose), which yielded a two orders of magnitude improvement in complete brain volume imaging capabilities, taking sixteen months to complete and generating imaging data to construct a volume of $8 \times 10^7 \mu\text{m}^3$. The reason for elaborating on this technique for brain imaging here is to highlight the complex technical challenges involved in imaging the relatively simple *Drosophila* brain, which has approximately 100,000 neurons, compared to the human brain with 87 billion neurons (Azevedo et al. 2009) and a volume of approximately $1,200 \text{ cm}^3$ (Cosgrove, Mazure, and Staley 2007, 848).

It is hoped that new capabilities in tracing neural circuits at synaptic resolution will allow for deeper exploration of the neural circuits involved in memory, learning, and behavior and so on, leading to insights which can also be used as inspiration for AI research. Historically, key strands of AI development have been based on insights from neuroscience, for example, deep learning in convolutional neural networks for machine vision was directly inspired by the mammalian visual cortex (Hubel and Wiesel 1959) and the technique of reinforcement learning was inspired by behavioral experiments in animal learning (Hassabis et al. 2017).

ARTIFICIALIZATION OF THE MIND

At present, contemporary AI and the mind and brain sciences are converging in an increasing number of areas as machine learning techniques are encroaching into the modeling toolkits of researchers in the mind and brain sciences, and neuroscience is increasingly being used to validate AI results. This convergence of AI and neuroscience has been described as a “virtuous circle” of shared insights between AI and neuroscience, where first, neuroscience provides inspiration for AI, second, neuroscience and psychology help to validate AI models and finally, AI models begin to be used to solve problems in neuroscience (Hassabis et al. 2017).

In the past three decades, AI, ANNs, and neuroscience have coalesced in the work and thought of Paul and Patricia Churchland, for whom the nature and possibility of human epistemic contact with the world is the central problem of their neurophilosophy, which is a program of research that is considered a highly plausible basis for understanding human intelligence and making progress toward artificial general intelligence. The approach of neurophilosophy is to provide models for understanding the brain as well as inspiration for developing AI tools for specific tasks and applications.

A key work by Paul Churchland at the intersection of philosophy of science and philosophy of mind is *Plato's Camera - How the Physical Brain Captures a Landscape of Abstract Universals* (2012). In this work, he begins by contrasting “a Kantian portrait of our epistemological situation” in which the two faculties of intuition and judgment constitute a canvas on which human cognition draws the empirical world we perceive, with his own proposal of many hundreds of high-dimensional internal cognitive maps through which our perceptions unfold (Churchland 2012, 1). His proposal is that there are numerous abstract spaces of representation—thousands of “cognitive spaces” on which human cognition is continually unfolding and embedded in the collective activities of ensembles of neurons. According to Paul Churchland, the fundamental unit of cognition is the activation pattern of ensembles of neurons; hence knowledge is represented in this activation space, and sculpted by years of learning through sensory impressions.

This constitutes the basis of the Churchlands’ “eliminative materialism” which regards “folk” psychology as flawed and must be eliminated in favor of mature cognitive neuroscience. The Churchlands’ program of eliminative materialism proceeds on the basis of the success in neuroscience in explaining phenomena related to brain states, flaws in folk psychology in explaining the same phenomena, and the incommensurability of the two approaches.

Eliminative materialism also has significant implications for the philosophy of science. First, Paul Churchland views the theories that yielded scientific revolutions from Galileo, Kepler, Descartes, Huygens, Newton, Boyle, and so on as progressively better reconceptualizations of various empirical domains, which provided us with a body of knowledge that allows us to identify the epistemological features they shared in making their discoveries. From this, he attempts to construct a neurally grounded approach to theory making in which previous semantic views of theories are to be viewed as instances of dynamical learning for which we can identify a neurocomputational basis. Hence, scientific research and theorizing are reconceived of as modifying and amplifying conceptual maps in the neuronal activation space, which, it is believed, may be simulated in computer hardware without the constraints of human neurobiology.

As a consequence, eliminativism has significant epistemological implications in all domains of knowledge, for example, in the philosophy of religion, according to the eliminativist position, witnessing divine action is simply to have one’s high-dimensional conceptual map vectorially indexed by sensory systems differently from nonbelievers. The problem with eliminativism, as Hilary Putnam explained, is that if all norms are explained away, there are no standards by which to assess competing explanatory claims, that is, there is nothing we can be right or wrong about. Hence, all attempts to naturalize epistemology question the very notion that we are

thinkers. The notion of “true” must go as well, and Putnam asks, “what are our statements but noise-makings? What are our thoughts but *mere* subvocalizations?” (Putnam 1982, 20).

In line with the program of neurophilosophy, many projects and research groups are seeking to “reverse engineer” the mind. For example, the Intelligence Advanced Research Projects Activity (IARPA) has a program of Machine Intelligence from Cortical Networks (MICrONS) which, “seeks to revolutionize machine learning by reverse-engineering the algorithms of the brain,” (“MICrONS” n.d.) and the DiCarlo Lab at MIT seeks to account, “for each ability of the mind (namely, intelligence) using components of the brain (neurons and their connections) in the language of engineering (computational models)” (DiCarlo 2018). The impetus behind such research is to recover models of specific mental processes that may be redeployed across a wide range of disciplines and domains of knowledge, especially scientific knowledge.

It is in this context that we can understand DeepMind’s ambition for developing AI systems to advance scientific knowledge and its founder’s aim to, “solve intelligence and use it to solve everything else,” (Burton-Hill 2016). Likewise, AlphaGo’s lead researcher David Silver said about AlphaGo Zero that, “we’ve removed the constraints of human knowledge and it is able to create knowledge itself” (Sample 2017), which raises epistemological questions about the status of such knowledge and what we can know through machines.

Regarding the status of such knowledge, the scientific ambition of reverse engineering nature requires remaking the portion of nature under investigation according to a particular metaphysical framework. The epistemological consequences for human knowledge from recasting the world in this way may be understood in reference to Giambattista Vico’s principle of *verum factum* (Dupuy 2009, 28). According to Vico, “The true and the made are convertible,” by which he meant that human beings can only rationally know of what they are themselves the cause, that is, what they have fabricated themselves. This was originally meant to say that we will never know nature as God does. Instead of seeking to understand the being of things, modern scientific knowledge is concerned with the how of processes by attempting to imitate the process of the coming into being of things and hence to see them from the standpoint of being their maker. Consequently, nature becomes artificial nature and it is no longer nature that is known, but that which we ourselves have made (Dupuy 2013, 68).

The problem with such an approach to scientific knowledge, as Floridi explains, is that science becomes increasingly artificial since complexity requires models that rely on artificial forms of understanding, which can lead to methodological mistakes from allowing oneself to become “enchanted by the affordances provided by the data” by ignoring “the constraints provided by the same data” (Floridi 2017, 284).

On the question of what we can know through machines, Paul Churchland offers a model of the mind as a recurrent neural network (RNN) with feedback to all parts of the cortex and a vision for ANNs to change the way we do science. At the basic level, this means pattern recognition in data, or in detection, measurement, and classification stages, much as we are seeing sweeping through the main areas of mind and brain research now. At another level, neural networks are envisaged as combining sensory modalities to capture high-dimensional phenomena for study beyond human capacities.

The height of Churchland's vision for ANNs to change science is to dispatch ANNs against the massive information drawn from the activity of the brain for insights into the character of human cognition, to understand principles of brain function, and the nature of our reasoning (Churchland 1987). The basic units of cognition in this vision are thus reduced to activation vectors, brain computations are vector-vector transformations and the basic unit of memory is simply synaptic weight configurations (Churchland 1996).

A major problem with such an approach is that ANNs are as opaque as the natural systems they attempt to model, as exemplified by problems of explainability, interpretability, transparency, and understandability of outputs (Doshi-Velez and Kim 2017). The nature of intelligence and cognition remains opaque and what is really gained is knowledge about the generalizability of an AI model, not knowledge corresponding to the real nature of the systems under investigation.

In remaking the mind in an artificial digital mold, contemporary attempts to "reverse engineer" the brain are reviving the AI fallacy of cognitive science from several decades ago, whereby "if a functional account can be given to a particular cognitive system, that *must* be how humans do it," which Chris Mortensen calls the AI fallacy (Slezak and Albury 1989). In addition, there is no reason from neuroimaging to suggest the patterns of activity seen in the brain are indicative of the communication channels used by the brain, that is, whether the activations recorded by experimenters are of any consequence for the brain itself (de-Wit et al. 2016).

Ed Feser argues that Paul Churchland does not succeed in using neuroscience to show abstract universals are represented in neuronal activities. Instead, what he does is change the subject by describing judgment, understanding, and knowledge as activation patterns, activation vectors, and high-dimensional activation spaces, and ignoring cognition by "talking about physiology instead." The question Feser raises is, "how does a pattern of neural activity constitute cognition any more than flexing of a tendon or the secretion of bile?" (Feser 2013, 31). Feser thus argues that Churchland is not explaining cognition in terms of physiology, he is simply ignoring (i.e., eliminating) cognition and discussing physiology in the vocabulary of cognition. He does nothing to justify why a neural

process counts as representation any more than any other physiological process.

According to Ed Feser, Paul Churchland's arguments may help with elucidating the material aspects of thought, but the immaterial aspects cannot be refuted by neuroscience (Feser 2013, 32) and dualism follows necessarily if one wants to maintain a mechanistic picture of the physical world while avoiding eliminative materialism. A mechanistic worldview entails eliminative materialism since science is an activity involved in assertions, theories, explanation, and knowledge, each of which is suffused with intentionality, which is central to defining the mind. Since all these activities point toward something, they are as intentional as the mind is. If the mind is eliminated, so too are the processes of science and reason in general. In effect, the eliminative materialist "saws off the branch on which he is seated" (Feser 2019, 123).

On the one hand, materialists would like to maintain concepts of truth, beliefs, desires, the mind, and intentionality. On the other hand, they would like to avoid dualism, even though it is taken for granted by the mechanical conception of the world within which technologists and engineers work.

AI IN THE SOCIAL SCIENCES AND HUMANITIES

In other domains, such as the social sciences and humanities, the use of AI raises similar epistemological issues that need careful attention and has led to societal concerns as these techniques are adopted by corporations, institutions, and governments. For example, AI-based facial recognition and psychographic profiling are already in widespread use and are actively being used to entrench the power of authoritarian regimes to turn entire countries into detention camps through algorithmic control, which amounts to a veneration of the algorithmic gaze and belief in AI as a magical tool for a new form of digital physiognomy.

Liberal democracies are also faced with new challenges arising from this technology as the same tools for algorithmic control are altering the future of education, work, and healthcare. For example, in work, AI is being integrated into the hiring process or for workplace monitoring as companies are under increasing pressure to compete for the most productive employees and to push their workers to achieve maximum output. In all such cases, human knowledge, understanding, reasoning, and judgment are gradually being ceded to machines, which is another aspect of the increasing orientation toward mechanocentrism.

ARTIFICIALIZATION OF THE WORLD

The second area undergoing reconfiguration is the mechanistic ontology, which is being replaced by an informational ontology. According to the mechanistic ontology, the physical world is regarded as nonagentic, passive,

and inert, with no internal principle of agency or movement, and admits no teleology, or final cause (Kureethadam 2009). All of reality, including mental reality, is seen as reducible to matter within the worldview of modernity given by scientific materialism, and nature is reconceived of in mechanistic terms, which opens the way for technological manipulation of the animate and inanimate as “objects of detached analysis, observation and experimentation” (Broschwimmer 2002, 56).

Successive metaphysical paradigms, namely, naturalistic, vitalistic, mechanistic, cybernetic, and now informational have been the basis through which scientists and engineers have attempted to force nature, life, and mind into the conceptual boxes supplied by these paradigms, to use a phrase from Kuhn. As discussed, there is a trend in modern science toward the elimination of mind and hence, the portion of reality to which it is assigned according to Cartesian dualism, and in place of mind, is a mathematical construction that appears to achieve an isomorphic mapping between inputs and outputs.

Furthermore, in place of the mind’s home in *res cogitans* is the creation of new hyperdimensional spaces in the form of the data structures that not only mediate between the physical and virtual world but are the microworlds for the artificial minds of artificial agents. In effect, AI represents a program of reinstalling what is notionally considered mind into *res extensa*, after the collapse of its former home in *res cogitans*.

In the case of key AI agents, such as an agent based on reinforcement learning, AI programmers use models to provide a representation of the environment to an artificial agent. Two possible approaches are first, model-free methods, where the agent learns its own model from training over a large number of samples by trial and error in unknown dynamical environments, such as the virtual world of a computer game, and second, model-based methods, where the agent learns according to a given model of the environment and a predefined value function or policy.

In contemporary AI discourse, much of the attention has been given to the nature of the agent but little is being said about the environment that represents the *umwelt*, or “self-centered world” of the agent. It is important to note that this environment is constituted of data input to the machine, which are cached in memory or a data structure, and can therefore be replaced by any other data. The point here is that an AI agent in any embodiment does not encounter our world at all; its world is a purely mathematical construction, implemented in a digital computational architecture. This is an issue discussed by authors such as Katherine Hayles during an earlier wave of AI and artificial life (AL) research. Hayles highlights that “the *material* space” of a computer’s interior architecture is differentiated from the lifeworld of the AL “creatures” which exist in “the *imagined* space that, in actuality, consists of computer addresses and electronic polarities on the computer disk” (Hayles 2000, 229).

One reason for the resurgence of robotics is because it is now possible to construct detailed virtual worlds with simulated real-world environments and law of physics in which robots may explore and learn, thus avoiding the high costs involved in assembling prototype robots. For example, NVIDIA's Isaac Platform, "lets developers train and test their robot software using highly realistic virtual simulation environments" ("Introducing NVIDIA Isaac" n.d.) with no constraints on the number of virtual robots that can be trained simultaneously. After training in these artificial worlds ("alternate universes" as described by NVIDIA's CEO, Jen-Hsun Huang) that look, sound, and behave like the real world, the "virtual brain" of the simulated robot may be transferred to a physical version of the robot to act in the physical world (*GTC 2017* 2017).

However, this disjuncture between the physical analog world and the agent's digital world is at the root of the problem of why AI systems often fail catastrophically when confronted with basic manipulations of their inputs, such as adversarial attacks. It may be argued that many problems in the emerging field of AI Ethics can be advanced by understanding the virtual world of the artificial agents, how the underlying ontology of digital systems prefigures the affordances of this rapidly developing artificial digital ecology, and how the artificial world of AI agents relates to the physical world.

One of the key areas in which contemporary AI has achieved success is in machine vision and object recognition, which are key requirements for AR technologies which require efficient scene recognition and image registration systems. Hence, in recent years, there has also been a surge of interest in the development of AR technology. As we discuss elsewhere (Chaudhary 2019), AR is the intermediate realm between the totalizing experience of virtual reality (VR) and the concrete physical world. VR, by definition, provides the most encompassing immersion of humans in cyberspace through the integration of audio, visual, tactile, and motile modalities and seeks to place the individual in a new virtual space.

There are different degrees of immersion in cyberspace in increasing degrees of sensory and bodily integration with digital information processing systems. The lowest degree is in the digital environments of web interfaces viewed through screens, which act as portals into the domain of cyberspace, the intermediate degree is the digital layer discussed earlier that maps onto and renders new entities in physical space and time (the case of AR and mixed reality), and the highest degree is complete immersion in virtual simulated environments (the case of VR).

Digital avatars, which are three-dimensional representations of a person or autonomous agent that embody actions, gestures and emotions in a virtual environment, are being constructed to provide presence for humans in cyberspace. Avatars may be controlled directly or autonomously and are

a key step toward the transposition of human activity from the physical world to the artificial digital world.

A recent demonstration by Facebook at their 2019 conference provides an initial illustration of the avatarization of human presence in V/AR (*F8 2019 Day 2 Keynote* 2019). The demonstration begins showing two avatars playing football in a realistic virtual environment. However, it is revealed that both (human) players are actually in a lab facility wearing VR headsets with full-body tracking, which replicates their presence and motions in the virtual environment. In reality, there is no grass or goal, and there is even no ball, yet one of the players jumps to block a return pass at chest height.

The ambition for AR, on the other hand, is the convergence of virtual and real space to achieve greater degrees of perceptual continuity between the virtual and real (Avram 2016, 35). AR platforms and applications are now becoming widespread, which is leading to the existence of “*a hidden data layer that you access through your devices – phones today, glasses tomorrow*” (Constine 2018).

With these advances in AR technology, augmented features, objects, and entities no longer appear as static digital overlays. Instead, according to our perceptions, which are mediated through screen-based devices, these features appear seamlessly blended and persistent within dynamic physical environments. Companies such as Amazon, Apple, Facebook, and Google are aiming to create a new computing paradigm where physical and digital objects are blended to the extent they become indistinguishable (Slater 2018).

The Google subsidiary, Sidewalk Labs, envisages the future smart city as a computing platform with a “digital layer” that is coextensive with the physical environment and is analogous to the operating system of a computer, which runs applications using the computer’s hardware subsystems. Except, in this case, the various subsystems are the constituent parts of the city and urban environment, such as, street lighting, traffic lights, and waste management, and other municipal services (Goodman and Powles 2019, 22–23). Data flowing from a wide range of sensors embedded throughout the built environment of the future smart city and connected via high-speed fiber optic and wireless networks to high-performance computing clusters will enable AI-based predictive analytics and automated decision making for management and governance of the city and its inhabitants (Goodman and Powles 2019, 24).

A key aspect of this technology is the virtualization of all people, places, and things as artificial constructs in the hidden digital data layer. For example, Replica, a Sidewalk Labs affiliate, is developing a full city model to simulate the city for future planning, which includes a synthetic population of virtual individuals generated from personal data of the city’s inhabitants (Goodman and Powles 2019, 27). Goodman and Powles describe the breaking down of “the material and social world into data flows”

as datafication (Goodman and Powles 2019, 33), which we describe more broadly here as part of the process of the “artificialization of the world.”

Connected technologies provide access to cyberspace, and AR is a means of rendering the salient features that exist in cyberspace in a form visualizable by human perception, such as providing augmented overlays for directions and detailed information cards attached to buildings and sights. However, the ambition goes further than providing informational overlays since cyberspace and AR represent new frontiers for capitalist expansion as many companies and start-ups are seeking to achieve a one-to-one mapping of the world covering everything, including interior spaces. This is leading to the creation of a new commercially controlled realm that is not only coexistent with the material world, but increasingly *envelopes*, a key word from industrial robotics, physical space. In this usage, “an envelope is the three-dimensional space that defines the boundaries that a robot can reach” (Floridi 2011, 228).

The ambition of fields such as cyber-physical systems (CPS) is thus to bring hundreds of billions of edge devices such as smartphones autonomous vehicles, drones, and robots, as well as artificial and human agents into ontological contiguity as informational entities embedded in a new informational environment. A recent paper by Couldry and Mejias on “Data Colonialism” says, “If successful, this transformation will leave no discernible ‘outside’ to capitalist production: everyday life will have become directly incorporated into the capitalist process of production” (Couldry and Mejias 2018, 8).

This new informational environment subsumes both cyber and physical space into a unified artificially constructed virtual world, which is being transposed over the physical world. This superimposed simulated model of the world is where the activities of artificial agents in various embodiments occur. That is, artificial agents do not encounter our analog world directly but rather the analog world is being “reformatted” according to the logic and ontology of the digital for the benefit of machines. Hence, an autonomous vehicle, for example, does not see or encounter the physical environment directly but instead processes streams of data to simulate roads, signs, and pedestrians in a virtual world which is mapped to the physical world.

In these virtualized worlds, artificial agents now outperform the most highly ranked human players in challenging games that require long-term planning, strategic decision making, and reasoning based on the imperfect knowledge of the microworlds of the games, as recent progress in AI has demonstrated. However, artificial agents are not intended to remain confined to the virtual microworlds of the games and simulations in which they are gestated and trained, and their digital worlds, which are better suited to the capacities of machines than humans, are gradually being unified and transposed over increasing portions of the physical world.

The philosophical implication of the intersection of humans and machines in the shared space of the “information sphere” entails what Luciano Floridi has described as a “re-ontologization of our environment and of ourselves” (Floridi 2010, 12).

Rather than cyborgs, we are becoming informational organisms, or in-forgs, as described by Floridi. A stark illustration of what it means to be an in-forg is the phenomenon of so-called microworkers, who are being used as *artificial* artificial intelligence (Atanasoski and Vora 2019; Gray and Suri 2019). Most notable is Amazon’s Mechanical Turk, which offers a marketplace for “human intelligence tasks” such as labeling data or images. These repetitive tasks are used to train machine learning algorithms such that the tasks can eventually be fully automated and taken over by artificial agents. For example, in 2019 it was first reported that behind Amazon Alexa’s voice system, and other similar systems, is a vast network of microworkers, each listening to and labeling thousands of voice samples to improve the quality of speech recognition (Day, Turner, and Drozdiak 2019). In cases such as microwork, the deeply metaphysical research program underlying AI research becomes apparent with the aim, “to place humankind in the position of being divine maker of the world, the demiurge, while at the same time condemning human beings to see themselves as out of date” (Dupuy 2009, xiv).

This leads us to the third reconfiguration, which is the elimination of the dualistic divide between humans and the physical world through the monism of an informational worldview, which entails ontological continuity between humans and artificial agents in a new informational environment.

The informational ontology may be considered on two levels, first as a strong claim that information constitutes the nature of ultimate reality, and a weaker claim that at a particular level of abstraction, human and machine agents can be understood informationally. In the former claim, according to physicist Paul Davies, information has a fundamental ontological status (Davies and Gregersen 2014) and is the substratum out of which matter, life, nature, and mind arise. An informationalist conception of the universe provides a unitary framework at the interface of biology, physics, chemistry, computing, and mathematics, and from this framework, it is patterns of information flows in cells, brains, and ecosystems that give rise to agency.

The latter claim entails the conceptual anthropomorphization of the machine and mechanomorphization of the human toward a central point of abstraction where both humans and machines are construed as informational entities. The form of human presence in this artificial digital world is through an abstraction into informational entities by hollowing out human nature to leave only what can be peeled off as quantifiable data such as biometrics, psychometrics, behavioral patterns, preferences, purchase history, and so on, which are reaggregated and reinscribed on

digital twins to form informational representation of ourselves (Haggerty and Ericson 2000, 606).

As inforgs alongside other informational entities, humans and machines become interchangeable as the distinction between the two is erased. According to Floridi, future generations are coming to live in a new condition known as the onlife, in which these “digital natives” will cease to appreciate any ontological difference between artificial agents and themselves as information entities. Previously, Hayles argued that, “envisioning humans as information processing machines with fundamental similarities to other kinds of information-processing machine, especially intelligent computers,” is the deeper sense in which it means to become posthuman (Hayles 2000, 246). Furthermore, the infosphere has blurred any distinction between the online and offline, and the digital world and the physical world, although the former is privileged over the latter.

CONCLUSION

The development of AI and AR, and the information communication technologies on which they are contingent, is the basis for a conceptual revolution that affects “how we understand the world, how we relate to it, how we see ourselves, how we interact with each other” (Floridi 2019, 208). Furthermore, these technologies “are increasingly ‘artificializing’ or ‘denaturalizing’ the world, human experiences, and interactions, as well as what qualifies as real” (Floridi 2019, 53).

Here, we have highlighted three interrelated philosophical reconfigurations that pertain to the epistemological, ontological, and anthropological implications of the research and development of AI and AR, which together constitute the underlying features of the process we have described as the artificialization of mind and world.

First, as new instruments in the hands of researchers in the natural and social sciences, AI and machine learning are raising profound epistemological issues that remain underexamined as scientific research and technological development outstrip the pace of philosophical inquiry. Second, more of the world is being recast into this artificial world or digital layer as artificial models come to stand in place of the buildings, places, and objects that exist in the physical world, leading to an artificial ontology of virtualized digital objects, actions, locations, and interactions. Finally, humans and machines, as well as the artificial worlds of cyberspace and the real world, are being brought into ontological continuity, as the natural world is being progressively subsumed in an artificial world, which is itself formed from the same digital substratum out of which the artificial minds of artificial agents come into being.

ACKNOWLEDGEMENTS

An earlier version of this article was presented at the 2019 Science and Religion Forum conference entitled “AI and Robotics: The Science, Opportunities, and Challenges.” The conference was held at St John’s College, Durham, UK, April 11–13.

I would like to thank the John Templeton Foundation for their financial support, Cambridge Muslim College (CMC) for hosting the broader project, and Gillian Straine and the Science and Religion Forum for inviting this contribution, which gathers together work presented throughout 2019, including at the 2019 CMC Religion and Science Conference on “Mind & World for Humans & Machines”.

REFERENCES

- Atanasoski, Neda, and Kalindi Vora. 2019. *Surrogate Humanity: Race, Robots, and the Politics of Technological Futures. Perverse Modernities*. Durham, NC: Duke University Press.
- Avram, Horea. 2016. “*The Visual Regime of Augmented Reality Art: Space, Body, Technology, and the Real-Virtual Convergence*.” PhD thesis, Montreal, Quebec, Canada: McGill University, Department of Art History and Communications Studies. http://digitool.library.mcgill.ca/R/?func=dbin-jump-full&object_id=143599&local_base=GEN01-MCG02.
- Azevedo, Frederico A. C., Ludmila R. B. Carvalho, Lea T. Grinberg, José Marcelo Farfel, Renata E. L. Ferretti, Renata E. P. Leite, Wilson Jacob Filho, Roberto Lent, and Suzanaerculano-Houzel. 2009. “Equal Numbers of Neuronal and Nonneuronal Cells Make the Human Brain an Isometrically Scaled-up Primate Brain.” *Journal of Comparative Neurology* 513 (5): 532–41. <https://doi.org/10.1002/cne.21974>.
- Bargmann, Cornelia I., and Eve Marder. 2013. “From the Connectome to Brain Function.” *Nature Methods* 10 (6): 483–90.
- Brosch, Franz. 2002. *Ecocide: A Short History of the Mass Extinction of Species*. London; Sterling, VA: Pluto Press.
- Burton-Hill, Clemency. 2016. “The Superhero of Artificial Intelligence: Can This Genius Keep It in Check?” *The Guardian/The Observer*, 16 February 2016, sec. Technology. <https://www.theguardian.com/technology/2016/feb/16/demis-hassabis-artificial-intelligence-deepmind-alpha>.
- Chaudhary, Mohammad Yaqub. 2019. “Augmented Reality, Artificial Intelligence, and the Re-Enchantment of the World”: with Mohammad Yaqub Chaudhary, “Augmented Reality, Artificial Intelligence, and the Re-Enchantment of the World”; and William Young, “Reverend Robot: Automation and Clergy.” *Zygon: Journal of Religion and Science* 54: 454–78. <https://doi.org/10.1111/zygo.12521>.
- Churchland, Patricia Smith. 1987. “Epistemology in the Age of Neuroscience.” *Journal of Philosophy* 84 (10): 544–53. <https://www.jstor.org/stable/2026917>.
- Churchland, Paul M. 1996. *The Engine of Reason, the Seat of the Soul: Philosophical Journey into the Brain*. New ed. Cambridge, MA: MIT Press.
- . 2012. *Plato’s Camera: How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge, MA; London, England: MIT Press.
- Constine, Josh. 2018. “Facebook Launches AR Effects Tied to Real-World Tracking Markers.” *TechCrunch* (blog). September 3. <http://social.techcrunch.com/2018/03/09/facebook-launches-ar-effects-tied-to-real-world-tracking-markers/>.
- Cosgrove, Kelly P., Carolyn M. Mazure, and Julie K. Staley. 2007. “Evolving Knowledge of Sex Differences in Brain Structure, Function, and Chemistry.” *Biological Psychiatry* 62 (8): 847–55. <https://doi.org/10.1016/j.biopsych.2007.03.001>.

- Couldry, Nick, and Ulises A. Mejias. 2018. "Data Colonialism: Rethinking Big Data's Relation to the Contemporary Subject." *Television & New Media*, September, 1527476418796632. <https://doi.org/10.1177/1527476418796632>.
- Davies, Paul C. W., and Niels Henrik Gregersen, ed. 2014. *Information and the Nature of Reality: From Physics to Metaphysics*. Cambridge; New York, NY: Cambridge University Press. <https://www.cambridge.org/pl/academic/subjects/physics/general-and-classical-physics/information-and-nature-reality-physics-metaphysics-1?format=HB&isbn=9781107684539>.
- Day, Matt, Giles Turner, and Natalia Drozdiak. 2019. "Amazon Workers Are Listening to What You Tell Alexa." *Bloomberg.com*, April 10.
- de-Wit, Lee, David Alexander, Vebjørn Ekroll, and Johan Wagemans. 2016. "Is Neuroimaging Measuring Information in the Brain?" *Psychonomic Bulletin & Review* 23 (5): 1415–28. <https://doi.org/10.3758/s13423-016-1002-0>.
- DiCarlo, James. 2018. "2018 Computational Neuroscience Workshop - Talk by Dr. James DiCarlo (MIT) and Open Discussion - YouTube." <https://www.youtube.com/watch?v=em8lPQVtFM>.
- Domingos, Pedro. 2015. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. 1st ed. New York, NY: Basic Books.
- Doshi-Velez, Finale, and Been Kim. 2017. "Towards a Rigorous Science of Interpretable Machine Learning." *ArXiv:1702.08608 [Cs, Stat]*, February. <http://arxiv.org/abs/1702.08608>.
- Dupuy, Jean-Pierre. 2009. *On the Origins of Cognitive Science*. Cambridge, MA: MIT Press. <https://mitpress.mit.edu/books/origins-cognitive-science>.
- . 2013. *The Mark of the Sacred*. Translated by M. B. DeBevoise. *Cultural Memory in the Present*. Stanford, CA: Stanford University Press. <http://www.sup.org/books/title?id=21129>.
- F8 2019 Day 2 Keynote. 2019. <https://www.youtube.com/watch?v=j48PqBP-OA0>.
- Feser, Edward. 2013. "Kripke, Ross, and the Immaterial Aspects of Thought." *American Catholic Philosophical Quarterly* 87 (1): 1–32. <https://doi.org/10.5840/acpq20138711>.
- . 2019. *Aristotle's Revenge: The Metaphysical Foundations of Physical and Biological Science*. Germany: EDITIONES SCHOLASTICAE.
- Floridi, Luciano, ed. 2010. *The Cambridge Handbook of Information and Computer Ethics*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511845239>.
- . 2011. "Children of the Fourth Revolution." *Philosophy & Technology* 24 (3): 227–32. <https://doi.org/10.1007/s13347-011-0042-7>.
- . 2017. "A Plea for Non-Naturalism as Constructionism." *Minds and Machines* 27 (2): 269–85. <https://doi.org/10.1007/s11023-017-9422-9>.
- . 2019. *The Logic of Information: A Theory of Philosophy as Conceptual Design*. Oxford, New York, NY: Oxford University Press.
- Goodman, Ellen P., and Julia Powles. 2019. "Urbanism under Google: Lessons from Sidewalk Toronto." SSRN Scholarly Paper ID 3390610. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=3390610>.
- Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston, MA: Houghton Mifflin Harcourt.
- GTC 2017. 2017. <http://www.youtube.com/playlist?list=PLZHnYvH1qtOZDhK77h7Gj3BDVUgDagh0L>.
- Haggerty, Kevin D., and Richard V. Ericson. 2000. "The Surveillant Assemblage." *British Journal of Sociology* 51 (4): 605–22. <https://doi.org/10.1080/00071310020015280>.
- Hassabis, Demis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. 2017. "Neuroscience-Inspired Artificial Intelligence." *Neuron* 95 (2): 245–58. <https://doi.org/10.1016/j.neuron.2017.06.011>.
- Hayles, N. Katherine. 2000. *How We Became Posthuman: Virtual Bodies in Cybernetics, Literature, and Informatics*. Chicago, IL: University of Chicago Press.
- Heidegger, Martin. 1977. *The Question Concerning Technology and Other Essays*. New York, NY: Cambridge: Harper & Row.
- . 1997. "The Age of the World Picture." In *Science and the Quest for Reality*, edited by Alfred I. Tauber, 70–88. *Main Trends of the Modern World*. London, UK: Palgrave Macmillan. https://doi.org/10.1007/978-1-349-25249-7_3.

- Huang, Jen-Hsun. 2016. "GTC Europe 2016 - Keynote - YouTube." <https://www.youtube.com/watch?v=npzRyTmcZo>.
- Hubel, David H., and Torsten N. Wiesel. 1959. "Receptive Fields of Single Neurones in the Cat's Striate Cortex." *Journal of Physiology* 148 (3): 574–91. <https://doi.org/10.1113/jphysiol.1959.sp006308>.
- "Introducing NVIDIA Isaac." n.d. NVIDIA. Accessed 3 September 2019. <https://www.nvidia.com/en-us/deep-learning-ai/industries/robotics/>.
- Jonas, Hans. 2001. *The Phenomenon of Life: Toward a Philosophical Biology*. Northwestern University Studies in Phenomenology and Existential Philosophy. Evanston, IL: Northwestern University Press.
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. 2017. "ImageNet Classification with Deep Convolutional Neural Networks." *Communications of the ACM* 60 (6): 84–90. <https://doi.org/10.1145/3065386>.
- Kureethadam, Jostrom Isaac. 2009. "Banished behind the Curtain of Nothingness." 19pp. https://filosofia.unisal.it/images/PDF/convegni/2009/ambivalenza_nulla/10.%20Kureethadam%20Banished%20behind%20the%20curtain%20of%20nothingness.pdf.
- . 2017. *The Philosophical Roots of the Ecological Crisis*. 1st ed. Newcastle upon Tyne: Cambridge Scholars Publishing.
- Li, Yuxi. 2018. "Deep Reinforcement Learning." *ArXiv:1810.06339 [Cs, Stat]*, October. <http://arxiv.org/abs/1810.06339>.
- Lichtman, Jeff W., Jean Livet, and Joshua R. Sanes. 2008. "A Technicolour Approach to the Connectome." *Nature Reviews Neuroscience* 9 (6): 417–22. <https://doi.org/10.1038/nrn2391>.
- Makari, George. 2017. *Soul Machine: The Invention of the Modern Mind*. Reprint ed. New York, NY: W. W. Norton & Company. <https://www.bloomberg.com/news/articles/2019-04-10/is-anyone-listening-to-you-on-alexa-a-global-team-reviews-audio>.
- 'MICRONS'. n.d. Accessed August 26, 2019. <https://www.iarpa.gov/index.php/research-programs/microns>.
- Putnam, Hilary. 1982. 'Why Reason Can't Be Naturalized', *Synthese*. Springer, 52(1), pp. 3–23. Available at: <https://www.jstor.org/stable/20115757> (Accessed: 10 April 2020).
- Sample, Ian. 2017. "'It's Able to Create Knowledge Itself': Google Unveils AI That Learns on Its Own." *The Guardian*, October 18, sec. Science. <https://www.theguardian.com/science/2017/oct/18/its-able-to-create-knowledge-itself-google-unveils-ai-learns-all-on-its-own>.
- Segler, Marwin H. S., Mike Preuss, and Mark P. Waller. 2018. "Planning Chemical Syntheses with Deep Neural Networks and Symbolic AI." *Nature* 555 (7698): 604–10. <https://doi.org/10.1038/nature25978>.
- Senior, Andrew W., John Jumper, Demis Hassabis, and Pushmeet Kohli. 2018. "AlphaFold: Using AI for Scientific Discovery." *Deepmind* (blog). December. <https://deepmind.com/blog/article/alphafold>.
- Senior, Andrew W., Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Židek, Alexander W. R. Nelson, Alex Bridgland, Hugo Penedones, Stig Petersen, Karen Simonyan, Steve Crossan, Pushmeet Kohli, David T. Jones, David Silver, Koray Kavukcuoglu, and Demis Hassabis. 2020. "Improved Protein Structure Prediction Using Potentials from Deep Learning." *Nature* 577 (7792): 706–10. <https://doi.org/10.1038/s41586-019-1923-7>.
- Seth, Anil. 2017. "Transcript of 'Your Brain Hallucinates Your Conscious Reality'." 2017. https://www.ted.com/talks/anil_seth_how_your_brain_hallucinates_your_conscious_reality/transcript.
- Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. "Mastering the Game of Go without Human Knowledge." *Nature* 550 (7676): 354–59. <https://doi.org/10.1038/nature24270>.
- Slater, Michael. 2018. "Facebook for Developers - AR for Everyone." <https://developers.facebook.com/videos/f8-2018/ar-for-everyone/>.
- Slezak, Peter, and W. R. Albury. 1989. *Computers, Brains and Minds: Essays in Cognitive Science*. Dordrecht, Netherlands: Springer. <http://public.eblib.com/choice/publicfullrecord.aspx?p=3100842>.

- The Event Horizon Telescope Collaboration. 2019. "First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole." *The Astrophysical Journal Letters* 875 (1): L1. <https://doi.org/10.3847/2041-8213/ab0ec7>.
- Williams, James. 2018. *Stand out of Our Light: Freedom and Resistance in the Attention Economy*. 1st ed. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781108453004>.
- Zheng, Zhihao, J. Scott Lauritzen, Eric Perlman, Camenzind G. Robinson, Matthew Nichols, Daniel Milkie, Omar Torrens, John Price, Corey B. Fisher, Nadiya Sharifi, Steven A. Calle-Schuler, Lucia Kmecova, Iqbal J. Ali, Bill Karsh, Eric T. Trautman, John Bogovic, Philipp Hanslovsky, Gregory S. X. E. Jefferis, Michael Kazhdan, Khaled Khairy, Stephan Saalfeld, Richard D. Fetter, and Davi D. Bock. 2017. "A Complete Electron Microscopy Volume of the Brain of Adult *Drosophila Melanogaster*." *BioRxiv*, January. <https://doi.org/10.1101/140905>.