Science Fiction's Imagined Technologies

with Emanuelle Burton, "The Nuts and Bolts of Transformation: Science Fiction's Imagined Technologies and the Civic Imagination"; Michelle A. Marvin, "Memory Altering Technologies and the Capacity to Forgive: Westworld *and Volf in Dialogue"; Nathan Schradle, "In Algorithms We Trust: Magical Thinking, Superintelligent AI, and Quantum Computing"; and Zhange Ni, "Reimagining Daoist Alchemy, Decolonizing Transhumanism: The Fantasy of Immortality Cultivation in Twenty-First Century China"*

IN ALGORITHMS WE TRUST: MAGICAL THINKING, SUPERINTELLIGENT AI AND QUANTUM COMPUTING

by Nathan Schradle

Abstract. This article analyzes current attitudes toward artificial intelligence (AI) and quantum computing and argues that they represent a modern-day form of magical thinking. It proposes that AI and quantum computing are thus excellent examples of the ways that traditional distinctions between religion, science, and magic fail to account for the vibrancy and energy that surround modern technologies.

Keywords: artificial intelligence; magic; quantum computing; religion and science; technology

---

The eschatological enthusiasm that surrounds artificial intelligence (AI) in the twenty-first century is often headline grabbing. Silicon Valley entrepreneur Anthony Lewandoski's Way of the Future Church, officially registered with the IRS as a religious organization, is a prime example. Lewandoski claims to be devoted to, "the realization, acceptance, and worship of a Godhead based on Artificial Intelligence (AI) developed through computer hardware and software" (Harris 2017). Though Lewandoski's church has faded from the headlines since its creation, similar eschatological enthusiasm is invested in AI by other wealthy tech elites like futurist and Google Director of Engineering Ray Kurzweil, who asserts that the relatively imminent arrival of the Singularity, a single global networked consciousness composed of all individual organic and carbon-based intelligences, will be tantamount to the creation of a God and signal the end of human suffering (Kurzweil 1999, 2005, 2012). The undeniably

Nathan Schradle is a PhD candidate at the University of North Carolina at Chapel Hill, Chapel Hill, NC, USA; e-mail: schradle@unc.live.edu.

religious idiom employed by many influential figures at big tech firms is so prevalent that it has long since caught the eye of scholars seeking to demonstrate the affinity between such utopian technological prognostications and centuries-old Christian millenarian movements and monastic traditions (Noble 1999; Geraci 2010). Erik Davis, for one, insists on the spiritual significance of such endeavors, writing that "regardless of how secular this ultramodern condition appears, the velocity and mutability of the times invokes a certain supernatural quality that must be seen, at least in part, through the lens of religious thought" (Davis 2004, 4).

In this essay, I take up Davis's injunction to examine the enthusiasm that surrounds AI and quantum computing "through the lens of religious thought," but in a way that I have not encountered in scholarship on the issue. Specifically, I propose that popular scientific thinking about AI and other algorithmically afforded technologies (most recently quantum computing) represents a modern version of what religious studies scholars have long associated with magic and magical thinking. I am wary that equating AI with magic may seem facile, given that Arthur C. Clarke's famous "third law" stating that "any sufficiently advanced technology is indistinguishable from magic" has become a well-known, almost clichéd reference point in this kind of work (Clarke [1962] 2000, 36). Still, the following analysis opens creative pathways for reconsidering what scholars have made of the intersection of magic, religion, and science in modernity, especially as it pertains to the theoretical and conceptual delineations made between the three spheres of thought.

Of course, the boundaries between religion, science, and magic are not the only ones being blurred in contemporary society. Especially when it comes to prognostications about the future of technology, the easy distinctions between fact and fiction start to similarly collapse. This is especially true in the contemporary mediasphere, where movies, books, and television shows are constantly discussed and dissected by critics and fans. As one example, a simple Google search for "why is everyone so obsessed with Zombies" will turn up endless hits for online articles published on behalf of entities ranging in seriousness and political leaning from Mashable to NPR and the Federalist analyzing the sheer volume of zombie-based entertainment and tying it to deeply serious concerns about the state of politics, climate change, and other large-scale threats that are percolating in the contemporary zeitgeist. The marketing for movies and television shows often plays on this modern development. For example, the creator of the dystopian technology anthology television series *Black Mirror* initially promoted his now popular show as a set of one-off stories about, "the way we live now—and the way we might be living in 10 minutes' time if we're clumsy" (Brooker 2011). *MIT Technology Review* publishes stories about the impending climate disaster by science fiction authors like Paolo Bacigalupi alongside its more straightforward reportage about technological innovations meant to combat climate change (Bacigalupi 2019). While the idea that art and works of fiction might have something to say about the real world is certainly not a new concept or only true of the twenty-first century, it is undeniable that modern mass media entertainment such as television is more available to a broader public and more intensely analyzed by a larger mass of critics and scholars than ever before. It is also undeniable that AI has currently captured the cultural imagination to such an extent that it appears in nearly every form of media constantly, even relentlessly. Focusing on the current cultural enthusiasm for AI opens up a means of examining the blurring boundaries between fiction and fact and the production of a collaborative fictive disposition toward AI, quantum computing, and the many technologies that announce, in the apocryphal words of acclaimed cyberpunk author William Gibson, "the future is already here."

This fictive disposition collapses boundaries between fact and fiction, reason and imagination, procedure and creativity. By exploring the notion of cutting-edge technology as "magical," I hope to demonstrate how public perceptions of the brave new world we are creating through AI and other innovations chafe against longstanding conceptions of the relationship between modernity and rationality in spite of the pseudo-scientific vernacular in which they are couched.

Magic and Modernity: A Brief Genealogy from Religious Studies

The myth of a disenchanted modernity was based not on banishment of religion from the world but rather of magic. Religion was certainly thought more and more to be a private matter of one's personal belief, but magic and belief in magic were outright expelled from what was supposed to be the modern worldview. Many scholars have pointed out that Max Weber's famous statement about the "disenchantment of the world" is probably more accurately translated from the original German as "the de-magic-ing of the world." As Jason Josephson-Storm (2017, 4) writes, "if there is one thing we've been taught to take for granted, it is that the contemporary, industrial, capitalist societies of Western Europe and North America have lost their magic, and that it is this absence that makes them modern." Josephson-Storm is not alone in this analysis. Randall Styers has previously argued that, "one common feature throughout these debates [about modernity] is that magic is an archetypically nonmodern phenomenon. Magic has offered scholars and social theorists a foil for modern notions of religion and science and, more broadly, a foil for modernity itself" (Styers 2004, 8).

This line of argument is compelling. In the decades leading up to Weber's famous declaration, scholars in the nascent social sciences seeking to delineate between religion and science often saved their most sneering criticism for what they considered to be magic or magical thinking. In the intellectualist tradition, early anthropologists Edward B. Tylor (1871) and James G. Frazer ([1900] 1951) considered magic to be a primitive mode of understanding, one that was superseded by religious belief and ultimately by scientific knowledge as cultures evolved into more and more complex forms. The evolutionist model these thinkers employed has obviously long since been discredited, but I want to focus on Frazer's articulation for a moment, since we will return to it briefly in the next section. Magic, Frazer argued, was based on the magician's faulty assumption that "things act on each other at a distance through a secret sympathy," one that was typically referred to as "an invisible ether" (Frazer [1900] 1951, 54). Of course, in Frazier's view such talk of "secret sympathies" gave way to the more sophisticated explanations of theology, and then ultimately to scientific description and explanation.

Many other early twentieth century anthropologists and sociologists sought to develop their own delineations between science, magic, and religion, including such landmark figures as Emile Durkheim and Sigmund Freud. However, the particularities of those delineations do not really concern us here, since by the latter half of the twentieth century and extending into the present day discerning religious studies scholars do not really consider the boundaries between these categories to be so hard and fast. This acknowledgment can be said to begin with Claude Lévi-Strauss, who argued that the distinction between magic as a fundamentally subjective practice and science as objective did not hold since, as Styers explains, "magic is based on the fundamental belief that humanity can intervene in the natural world to modify or add to its system of determinism" (Styers 2004, 9). Or, as Lévi-Strauss (1966) himself wrote:

> Religion consists in a *humanization of natural laws and magic in a naturalization of human actions*—the treatment of certain human actions *as if* they were an integral part of physical determinism. The anthropomorphism of nature (of which religion consists) and the physiomorphism of man (by which we have defined magic) constitute two components which are always given, and vary only in proportion. As we noted earlier, each implies the other. There is no religion without magic any more than there is magic without at least a trace of religion. The notion of a supernature exists only for a humanity which attributes supernatural powers to itself and in return ascribes the powers of its superhumanity to nature. (220–21)

Ultimately, it is this insight of Lévi-Strauss's that leads me to assert that our modern relationship to technology is inherently magical. As we shall see in the next section, what are the popular science advocates of AI discussing if not "the physiomorphism of humanity" or "adding to nature's system of determinism?"

AI Innovators' Magical Thinking: The Promises of ai and "The Religion of Technology"

Beginning in the mid-1970s, as personal computers first became available and the technologies that would form the backbone of the original "world wide web" were being developed, enthusiasm for this new-found digital technology reached new heights. The blossoming of information theory and cybernetics reshaped the physical and biological sciences, as well as providing a communicative and ontological bridge between the natural world and the engineered world of machines. Computers could be encoded via algorithm the way that humans and all living things were encoded via their DNA. One could be either constructed or treated as a facsimile of the other. Inspired by an earlier generation of cybernetic thinkers such as W. Ross Ashby and Norbert Wiener, researchers began to act upon the idea that machines (especially computers) were both the means of better understanding and improving upon human biology, especially as it pertained to human intelligence. Computer models would enable more accurate modeling of the brain and advances in technology would enable scientists to build artificial brains that lacked the foibles and imperfections of the actual organs. A vocal cadre of AI researchers cast themselves as the harbingers of a new era in which humans could escape the fleshly entombment of the body via mind uploading, the construction of artificial replacements for failed organs, and other postulated techniques.

Hans Moravec was one of the first to prognosticate that these digital technologies would fundamentally reshape the human experience both biologically and intellectually. He argued that the ability to transfer human minds into computers would result in the practical attainment of immortality, through a process in which human intelligence learned to "constantly [improve] and [extend] itself, spreading outwards from the solar system, converting non-life into mind" (Moravec 1978, 6). Scientists' ability to discern the patterns of the brain, he said, would allow them to perfectly (albeit artificially) replicate those patterns. This would theoretically lead to what Moravec called the "Age of Mind," in which humans occupy a perpetually networked mental space known as the "Mind Fire." This process, Moravec argued, would "leave a subtler world, with less action and even more thought, in its ever-growing wake" (Moravec 1999, 163). Most significantly, Moravec promised that the era of human existence overseen by the Mind-Fire would be one in which humans finally have control over their evolutionary future (Moravec 1999, 158–59).

Already, one can easily see how this type of thinking might be described as "magical" in the sense highlighted above. "Nature's system of determinism" had been discovered to be informational and thus subject to the kinds of rules of feedback, regulation, and communication that likewise governed the virtual world of computers. Thus, the knowledgeable could use these new computerized machines to unlock the mysteries of both the cosmos and the human mind. To translate this directly into Levi-Strauss's terms, the human action that constituted an "integral part of physical determinism" was the capacity to shape and model previously uninvestigable natural processes using computerized technology. The reward for this process would be the "physiomorphism of man," as human increasingly integrated with machines and joined Moravec's "Mind Fire."

It is therefore unsurprising that Moravec argued that the "inhabited portions of the universe will be rapidly transformed into cyberspace…where every bit will be part of a relevant computation or will be storing a useful datum" (Moravec 1992, 15; 1999, 166). "Once boring old Earth is swallowed by cyberspace," Moravec asserted, it will be free to "host astronomically more meaningful activity" (Moravec 1999, 167). To follow Moravec's argument to its conclusion, the Mind-Fire will converge into a "final bubble of Mind," such that our traditional notion of time will collapse, and "entire world histories," and other, more traditionally religious promises, specifically the resurrection of the dead, will be made possible by the global mental-collective-as-computer-processor (Moravec 1999, 167).

Ray Kurzweil likewise emphasizes mind-uploading and melding via technological means (what he has referred to throughout his career as "the Singularity") as the ultimate goal of technological innovation in his own work (Kurzweil 1999, 260; 2005, 487). In an explicitly religious idiom, he imagined it will "ultimately infuse the universe with Spirit," and, in so doing, make life "truly meaningful" (Kurzweil 2005, 389). This, of course, leads to any number of happy circumstances that humans will experience once their minds are allowed to roam free as part of the Mind-Fire, including the end of all need, the end of nationalism and war, immortality, and the infinite expansion of intelligence (Geraci 2010, 36). In other words, the Mind-Fire will bring about an end to the most directly felt limitations and hardships experienced by modern bodies, both individually and collectively. For these reasons, Moravec argued that roboticists were and are leading us in the final phase of evolution (Moravec 1988, 2). Danny Hillis (2001) refers to the AI revolution as, "one of those rare times in history when humanity transforms from one type of human society to another"—Levi-Strauss' "physiomorphism of man" once again (29–30). Kurzweil (with his typical bluster) goes even further, arguing that, "the emergence in the early twenty-first century of a new form of intelligence on Earth that can compete with, and ultimately significantly exceed, human intelligence will be a development of greater import than any of the events that have shaped human history" (Kurzweil 1999, 5).

Despite this affinity with magical thinking, these thinkers reach for the idiom of religion, long associated with notions of salvation and transcendence, to describe their ideas. Given the monumental and virtually miraculous nature of these imagined shifts, it is little wonder. One of the few (and certainly the most thorough and forceful) academic works to highlight the religious elements of these prognostications, Robert Geraci's *Apocalyptic AI* (2010), deals with exactly these issues. Geraci's genealogy of "Apocalyptic AI" draws on the work of Moravec, Kurzweil, and so on to compare these AI researchers' claims about transcending the body and entering into a state of "pure mind" with the mind-body dualism of first and second century Christian and Jewish apocalypticists. Specifically, he suggests that "Apocalyptic AI promises to resolve the problems of dualism and alienation in a radically transcendent future where we forsake our biological bodies in favor of virtual bodies that will inhabit an omnipresent and morally meaningful cyberspace" (Geraci 2010, 9).

In so doing, he argues, "Apocalyptic AI sets up values and practices designed to transport the human being from a state of ignorance, embodiment, and finitude to a state of knowledge, immateriality, and immortality" (Geraci 2010, 139). Geraci also locates this scientific pursuit in a broader pop-cultural frame, particularly in the pages of *Wired* magazine, in which founding editor Kevin Kelly and frequent contributor Margaret Wertheim have advocated a "Nerd Theology" that imagines cyberspace as the means of fulfilling a "psychological, religious void" (Wertheim 1999, 30) and the programmers of computerized worlds as the gods of their own theological systems (Kelly 1999, 389).

Geraci's argument is compelling precisely because it identifies religion in a movement that is not always quick to embrace that label even as it coopts its thematic language. William Sims Bainbridge, for example, imagines that improved cognitive sciences based on AI innovations will squeeze out the last of our religious superstitions (Bainbridge 2006, 207–08), even as he sends National Science Foundation funding to Hans Moravec and the Carnegie Mellon Robotics Institute. However, Kurzweil, a former employee of the Institute, has not only proclaimed the spiritual significance of the Singularity, but likewise written that he believes intelligent machines will in fact be *more* spiritual than human beings, and that, in the future, these machines will gather in both real and virtual houses of worship for religious purposes (Kurzweil 1999, 153). Hence, Kurzweil argues, we "need a new religion" to enhance morality and encourage the spread of knowledge (Kurzweil 2005, 374–75).

The broad spectrum of attitudes toward traditional religion and AI's relationship to it is best represented by the juxtaposition between Hugo de Garis' work in *The Artilect War* (2005) and the incredibly influential work of Marvin Minsky, the founder of the Computer Science and Artificial Intelligence Lab at MIT. De Garis imagines "the Artilects" (his phrase for the super-intelligent robots connected via a hive mind that humans are on the verge of building) as literal gods, worthy of reverence and worship. He describes faith in the "Artilect mission" as a "powerful new religion" (de Garis 2005, 105) that should, in due course, supplant the "superstition of older religious systems by providing its faithful with incontrovertible evidence (de Garis 2005, 94). Minsky (1985), on the other hand, is perpetually dismissive of religion, claiming that the various beliefs regarding souls manifest in religious modes of thought throughout history are, "all insinuations that we're helpless to improve ourselves" (41). In his science fiction novel, *The Turing Option*, Minsky reiterates that souls do not exist, and that belief in gods of any kind is an "invalid" superstition (Minsky 1992, 386).

These two positions toward religion could not be more diametrically opposed. However, these views are not really so different, but instead are using the term "religion" for different rhetorical purposes. Whether it is imagined that AI will bring about a "new religion" to supplant the false superstitions of the old one or eliminate the need for religion entirely, "religion" as a category is being employed in all these views as an epistemological mode that addresses the existential concerns of people who lack the intellectual resources (cultural or personal) to fully recognize and grapple with the reality of the situation. The falsity of these traditional religions is proven in AI researchers' minds by their failures to make good on any of their promises—the end-times have not come, there is no proof that Heaven exists, much less of eternal life through an immaterial soul, and so on. When these AI apocalypticists go so far as to describe the 'new religion' that will emerge as a result of the blossoming of the Mind-Fire/Singularity/Artilects, it is proposed in the sense that, at long last, science's capacity to radically re-frame human experience and intelligence through artificial means will bring about the revelation long hoped for but never received. Even the dupes who practice religion, it seems, will be forced to recognize the vastly improved human circumstances produced by these technologies in the face of the monumental, landscape-shifting evidence to come. However, by setting themselves up as the progenitors of a science so powerful it will fulfill the previously empty promises of religion, these men either blithely ignore or remain unaware of the resemblance of their ideas about the promise of deliverance through technology to earlier magical notions in the Levi-Straussian vein.

Geraci argues that the researchers and advocates he analyzes, whether consciously or not, adopt the language of religion as a kind of appeal to the public. Geraci writes, "apocalyptic AI advocates have been more successful… in part because they use religious categories to heighten the allure of their subject matter…[it] will not likely appeal to the traditionally religious faithful, but it finds a ready audience among the religiously disaffected who might find a 'powerful new religion' and a new kind of god to worship in the movement's promises" (Geraci 2010, 61–62). He moves from there to point out how drumming up religious levels of enthusiasm for AI research is a clever way to drive the enthusiasm that leads to increased levels of funding, which he suggests may be the ultimate goal.

AI and Cutting-Edge Technology as Magic

This is where Geraci's otherwise brilliant and articulate argument goes slightly astray. On a practical level, it elides the realities of the broader careers of these AI researchers. Put simply, none of these men are desperately in search of funding given the current financial enthusiasm for AI, nor will they be should the more bombastic elements of their prognostications fail to come to pass. Rather than being a central, defining element of their careers, most of their bold prognostications have little to do with the work for which they have been formally recognized. In Kurzweil's case, he has been lavished with public praise for his more practical research. He was both named one of PBS' "Revolutionaries Who Made America" and hailed as "Edison's rightful heir" as one of the "most fascinating" entrepreneurs in the United States by *Inc.* magazine. In both instances, his work with optical- and speech-recognition patterns leading to machine learning innovations and not his popular scientific writings earned him such praise. The story with Minsky is quite similar—he was the inventor of the first functioning neural network (the Stochastic Neural-Analog Reinforcement Computer, or SNARC) in 1951, as well as some of the first programmable robots. He was also a major innovator in machine learning, and founded and ran the MIT Artificial Intelligence Lab for 15 years. As such, both men have engaged in continuously self- and investor-enriching work during nearly every step of their careers. In so doing, he aims to call into question the distinct boundaries between the categories of "religion" and "science" and to challenge prevalent narratives that associate the modern world with disenchantment. This is a worthy goal, but one that can be even more productively explored by re-installing the category of "magic" alongside those of "religion" and "science."

To further explore this link, a detour through the dawning era of quantum computing is necessary alongside a reexamination of AI There is no more curious example of the collapsed boundaries between the categories of magic and science than quantum theory. The quantum realm forms a kind of "invisible ether" through which the observable world of classical physics moves. Quantum entanglement dictates that things, in this case electrons or photons, act on one another at even unimaginable distances. Even the world's most knowledgeable scientists are still sorting out the process, what one might call a "secret sympathy" if Einstein had not so famously dubbed it "spooky action at a distance." In short, Frazer's theory of magic developed at the turn of the twentieth century maps onto one of the most advanced scientific theories that modern science has devised. At the same time, the leaders of major companies and research teams working in the field invoke exactly the kind of magical thinking that the earlier generation of AI researchers engaged in.

Pulling Rabbits out of Black Boxes: The Current State of AI and Quantum Computing

On October 23, 2019, as I prepared the original draft of this article, the Google AI Quantum team published an article in the journal *Nature* announcing the achievement of "Quantum Supremacy," the act of building and programming a quantum computer that can quickly complete a task that a classical computer simply cannot complete in any practical amount of time. In this particular instance, the quantum computer, dubbed "Sycamore," took a mere 200 seconds to accomplish what the Google team estimated would have taken 10,000 years on a supercomputer composed of roughly 100,000 classical central processing units (Arute et al. 2019). Google's claims may be a bit overblown—IBM claimed on their research blog on the very same day that the Summit supercomputer it installed at the Oak Ridge National Laboratory last year could actually store all of the relevant data generated on its immense hard drives and simulate Google's results in a mere two and a half days. However, whether or not Google has actually achieved "quantum supremacy," this announcement is undoubtedly a major milestone in the history of computing.

Computer scientists and technology experts alike have been effusive when describing the potential applications of quantum computation, especially when combined with ongoing advances in machine learning and AI. Perhaps self-evidently, increases in quantum computing power could finally realize Richard Feynman's dream of accurately simulating the complex quantum behavior of physical systems, currently impossible given the limitations in representing the probabilistic superpositions of electrons, and so on via the typical bits of data that classical computers process (Feynman 1982, 5). But the truly awesome applications don't stop there: potential revolutions in current practices range from turning the lengthy and costly process of drug discovery into a matter of an afternoon's work (Cao, Romero, and Aspuru-Guzik 2018, 16) to the near-perfect optimization of the evaluative processes involved in the pricing of financial derivatives and supply chain management (Orús, Mugel, and Luzaso 2019).

This sanguine outlook on the future of quantum computing has been accompanied by dire warnings. Though an exhaustive survey of all the prognostications currently being bandied about relative to quantum computing in both academic and popular outlets would be both impractical and impossible, one of the most common refrains following Google's announcement has been that it spells the end of data encryption as we currently conceive of it. Moreover, this outcome also appears to be relatively easily realized. Konstantinos Karagiannis, Director of the Penetration Testing Team at the New York labs of security consulting firm Protiviti, is one of many experts who are certain that current encryption systems will be rendered obsolete in the very near future. Karagiannis (2019) says "I have no doubt that within five years we will have a machine that we really should be protecting against…a machine that at a minimum could be stealing cryptocurrency, at a maximum reading everybody's messages."

In short, quantum computing stands poised to radically reform the technological apparatuses employed in all of our major institutions. Or does it? While estimates for the number of qubits that a quantum computer must employ in order to achieve any of the tasks mentioned above are very much the subject of debate, estimates for something like Feynman's simulation of the quantum behavior of real-world physical systems range into the millions. Getting a quantum computer to run Shor's algorithm, devised in the 1990 as a potential way to break and exploit all RSA encryption and hence render the vast majority of digital encryption systems easily crackable, would only require upward of 10,000 qubits. Google's Sycamore quantum computer used all of its 53 qubits to perform a task that was selected explicitly as a kind of proof of concept trial for quantum computing, and which has no practical applications beyond said proof of concept. Setting aside the complicated disagreements over precise numbers of qubits needed for particular applications and ideas about how to increase the reliability of these systems, it is obvious that there is a major leap in complexity when moving from a system that counts its requisite processing parts in the tens or hundreds to one that does so in the millions. Many knowledgeable and consequential figures in the field are apprehensive about the difficulty of this leap. The report put out by the National Academy of Sciences on the state of quantum computing from 2019, for its part, concludes on the note that it "is not yet clear if or when" such a leap will be made, stating that, "while the component technologies and baseline protocols…have already been demonstrated, system-scale demonstration with practical levels of performance remains a major challenge" (Grumbling and Horowitz 2019, 130). Another major concern is the current cost associated with such projects—even if the appropriate architecture existed to build a large-scale quantum computer, one capable of the kinds of simulations that so inspired Richard Feynman and other physicists would cost somewhere in the neighborhood of $10 billion (Hossenfelder 2019).

However, whether or not quantum computing of practical problems is economically or scientifically feasible is not really the point I wish to raise here. John Preskill, the computer scientist responsible for coining the phrase "quantum supremacy," has commented on Google's announcement in a way that clarifies what I see as the more significant point, one that acknowledges the "magico-scientific" paradox that lies at the heart of the quantum computing enterprise. He writes:

> In a sense, merely looking at a quantum system unavoidably disturbs it, a manifestation of Heisenberg's famous uncertainty principle. So if we want to use such a system to store and reliably process information, we need to keep that system nearly perfectly isolated from the outside world. At the same time, though, we want the qubits to interact with one another so we can process the information; we also need to control the system from the outside and eventually measure the qubits to learn the results of our computations. It is quite challenging to build a quantum system that satisfies all of these desiderata, and it has taken many years of progress in materials, fabrication, design and control to get where we are now. (Preskill 2019)

Here, Preskill puts an optimistic spin on the most fundamental challenge that faces quantum computation: how to control a computational process whose promise lies in its capacity to function beyond our ability to monitor and intervene in it, and thus beyond our ability to address the errors and faults that inevitably arise in a probabilistic quantum system. Preskill (2018, 16) writes of these systems that "the truly transformative quantum technologies of the future are probably going to have to be fault tolerant, and because of the very hefty overhead cost of quantum error correction, the era of fault-tolerant quantum computing may still be rather distant. No one really knows how long it will take to get there."

The preceding discussion of quantum computing makes it plain that both the promise of employing and the challenge posed by interpreting the use of the technology seem to lie in its capacity to perform operations beyond the purview of human perception. Thus, quantum computing expands the spectrum and depth of human knowledge of the physical world beyond what humans acting on their own are capable of, and it does so in ways that humans are not able to fully apprehend. What is that if not a formula for magic in the traditional scholastic sense?

This conjuncture of what once was called magic with what is still called science coalescing around cutting-edge technology is only exposed as wider and deeper than previously imagined when one turns their attention to current events in the realm of AI. Machine learning algorithms are not so much programmed as they are unleashed. In fact, the entire premise of developing such algorithms hinges on the notion that they will be able to process data and enact simulations and plans at far faster than human speeds. The term for this kind of technology has been, since the postwar rise of cybernetics, "black boxing," referring to any system where an observer can only interact with it based on known inputs and recorded outputs. What goes on inside of the box is a mystery (Ashby [1952] 1960, 7–13). Such algorithms, and the neural networks built using them, must necessarily be "black boxes" in the same way that quantum computers are—the entire purpose of the enterprise is defeated if human users are able to observe the process.

These black box algorithms have remarkable real-world effects. Medical diagnostic algorithms developed using deep learning processes have proven to be more effective than their human counterparts, as in the Deep Patient algorithm developed at the Mount Sinai Hospital in New York City. Deep Patient proved to be more effective than any previous method of diagnosing a host of ailments and illnesses, from liver cancer to schizophrenia. Yet no one is certain what patterns it has learned to recognize in the patient data it was given (Knight 2017).

This style of engaging in scientific enquiry has become so prevalent that Daniel Dennett devoted a chapter of his most recent book to exploring the ramifications of "black box science" (Dennett 2017, 385–88). In spite of all that we do not and all that we cannot know about the functioning of these technologies, enthusiasm for them is only growing, and the bulk of AI researchers are beginning to sound more and more like Kurzweil in their predictions as to when potential applications of such black box technologies might reshape society. In a survey of 352 leading researchers presenting at major AI conferences, the participants on average concluded that such deep learning systems would be able to outperform humans at a myriad of tasks in the near future. Noteworthy prognostications include writing a bestselling novel by 2049 and working as a surgeon by 2053. The same sample of researchers also estimated there was a 50 percent chance that AI would be able to outperform humans at any conceivable task within 45 years (Grace et al. 2018, 729).

The enthusiasm that surrounds black box technologies, and the way they stand poised (whether in the public imagination or in reality) to radically reshape society shows that modern disenchantment and the denial of magic were always at least as much myth as reality. Of course, the abandonment of that myth also involves the abandonment of the codified distinctions between religion, science and magic which it justified. By calling these technologies "magical," I do not wish to revivify the old classificatory systems but instead to point to the aporias they produced. As Styers (2004, 215) notes, "positioned in stark counterpoint to the norms of modernity, magic holds great allure for cultural theorists with its intimations that the subterranean operations of modernity might be uncovered." One might say the same of the fictive disposition at play in science fiction and other modern media that insist on not only the enchanted nature of technology, but also it multiplicity. The more alternative fictive possibilities presented by science fiction, the more we might collectively be able to stave off the most insidious of the "subterranean operations of modernity": an all-too-easy faith in the inevitability of positive forward societal progress, guaranteed by the eager adoption of technological innovations. But regardless of what alternatives might be proposed or fictive possibilities explored, the process of fully appreciating the transformative role that AI and quantum computing stand poised to play must necessarily begin with an acknowledgement of the magical thinking that surrounds these technologies. For their part, scholars working on the intellectual history of religion, science, and magic must continue asserting that the notion that humans and their societies have become "modern" in some clean break with the past is as much fantasy as fact. Our very relationship to modern technology proves as much.

Acknowledgment

An earlier version of this article was presented in November 2019 at the annual American Academy of Religion (AAR) conference in San Diego, California as part of a panel cosponsored by *Zygon: Journal of Religion and Science* and two of AAR's program units, the Religion and Science Fiction unit and the Religion, Science and Technology unit. The panel addressed "The Nuts and Bolts of Transformation: Science Fiction's Imagined Technologies and the Civic Imagination."

References

Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J., Barends, R., Biswas, R. et al. 2019. "Quantum Supremacy Using a Programmable Superconducting Processor." *Nature* 574: 505–10. https://doi.org/10.1038/s41586-019-1666-5

Ashby, Ross W. (1952) 1960. *Design for a Brain: The Origin of Adaptive Behavior*. London: Chapman and Hall.

Bacigalupi, Paolo. 2019. "A Full Life." *MIT Technology Review*, April 24, 2019. https://www.technologyreview.com/s/613349/a-full-life/

Bainbridge, William Sims. 2006. "Cognitive Technologies." In *Managing Nano-Bio-Info-Cogno Innovations: Converging Technologies in Society*, edited by William Sims Bainbridge and Mihail C. Roco, 203–26. Dordrecht: Springer.

Brooker, Charlie. 2011. "The Dark Side of Our Gadget Addiction." *The Guardian*, December 1, 2011. https://www.theguardian.com/technology/2011/dec/01/charlie-brooker-dark-side-gadget-addiction-black-mirror

Cao, Yudong, Jhonathan Romero, and Alán Aspuru-Guzik. 2018. "Potential of Quantum Computing for Drug Discovery." *IBM Journal of Research and Development* 62 (6): 1–20.

Clarke, Arthur. (1962) 2000. *Profiles of the Future: An Enquiry Into the Limits of the Possible*. London: Orion Publishing.

Davis, Erik. 2004. *Techgnosis: Myth, Magic, and Mysticism in the Age of Information*. New York, NY: Harmony Books.

de Garis, Hugo. 2005. *The Artilect War: Cosmists vs. Terrans: A Bitter Controversy Concerning Whether Humanity Should Build Godlike Massively Intelligent Machines*. Palm Springs, CA: ETC Publications.

Dennett, Daniel. 2017. *From Bacteria to Bach and Back*. New York, NY: WW Norton and Co.

Feynman, Richard. 1982. "Simulating Physics with Computers." *International Journal of Theoretical Physics* 21 (6–7): 467–88.

Frazer, James G. (1900) 1951. *The Golden Bough*. London: MacMillan & Company.

Geraci, Robert. 2010. *Apocalyptic AI: Visions of Heaven in Robotics, Artificial Intelligence, and Virtual Reality*. New York, NY: Oxford University Press.

Grace, Katja, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans. 2018. "Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts." *Journal of Artificial Intelligence Research* 62: 729–54.

Grumbling, Emily, and Mark Horowitz. 2019. *Quantum Computing: Progress and Prospects*. (National Academies of Sciences, Engineering, and Medicine). Washington, DC: The National Academies Press. https://doi.org/10.17226/25196.

Harris, Mark. 2017. "Inside the First Church of Artificial Intelligence." *Wired*, November 15, 2017. https://www.wired.com/story/anthony-levandowski-artificial-intelligence-religion/

Hillis, Danny. 2001. "A Time of Transition/The Human Connection." In *True Names and the Opening of the Cyberspace Frontier*, edited by James Frenkel, 27-32. New York, NY: Tor.

Hossenfelder, Sabine. 2019. "Quantum Supremacy Is Coming. It Won't Change the World." *The Guardian*, August 2, 2019. https://www.theguardian.com/technology/2019/aug/02/quantum-supremacy-computers.

Josephson-Storm, Jason. 2017. *The Myth of Disenchantment: Magic, Modernity, and the Birth of the Human Sciences*. Chicago, IL: University of Chicago Press.

Karagiannis, Konstantinos. 2019. Interview on "The Future of Quantum Computing," Richard Clarke's Future State Podcast.

Kelly, Kevin. 1999. "Nerd Theology." *Technology in Society* 21 (4): 387-92.

Knight, Will. 2017. "The Dark Secret at the Heart of AI." *MIT Technology Review*, April 11, 2017. https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/

Kurzweil, Ray. 1999. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*. New York, NY: Penguin.

———. 2005. *The Singularity is Near: When Humans Transcend Biology*. New York, NY: Penguin.

———. 2012. *How to Create a Mind: The Secret of Human Thought Revealed*. New York, NY: Penguin.

Lévi-Strauss, Claude. 1966. *The Savage Mind*. Chicago, IL: University of Chicago Press.

Minsky, Marvin. 1986. *The Society of Mind*. New York, NY: Simon and Schuster.

Minsky, Marvin, with Harry Harrison. 1992. *The Turing Option: A Novel*. New York, NY: Grand Central Publishing.

Moravec, Hans. 1978. "Today's Computers, Intelligent Machines, and Our Future." *Analog* 99 (2): 59–84.

———. 1988. *Mind Children: The Future of Robot and Human Intelligence*. Cambridge, MA: Harvard University Press.

———. 1999. *Robot: Mere Machine to Transcendent Mind*. New York, NY: Oxford University Press.

Orús, Román, Samuel Mugel, and Enrique Lizaso. 2019. "Quantum Computing for Finance: Overview and Prospects." *Reviews in Physics* 4 (November): 1–12. https://doi.org/10.1016/j.revip.2019.100028

Preskill, John. 2018. "Quantum Computing in the NISQ Era and Beyond." *Quantum* 2 (79): 1–16.

———. 2019. "Why I Called It Quantum Supremacy." *Quanta*, October 2, 2019. https://www.quantamagazine.org/john-preskill-explains-quantum-supremacy-20191002/

Styers, Randall. 2004. *Making Magic: Religion, Magic, and Science in the Modern World*. Oxford: Oxford University Press.

Tylor, Edward B. 1871. *Primitive Culture: Researches into the Development of Mythology, Philosophy, Religion, Language, Art, and Custom*. London: John Murray.