

Artificial Intelligence and Theology

A Bibliographic Essay

by Brady Alan Beard

ABSTRACT: This bibliographic essay provides a starting point for theological librarians to engage the history, anthropology, ethics, theology of artificial intelligence. What it shows is that despite seeming to appear ex nihilo, artificial intelligence for writing has a long history, profound ethical consequences beyond academic integrity, and latent theology. While theological librarians may or may not have the computer science background necessary to engage the technicalities of AI, they certainly have the disciplinary knowledge and skill to engage the topic in a meaningful and theological way.

In November 2022, ChatGPT 3.0 was released on the Internet. Within days, pundits and talking heads in fields as diverse as fulfillment and education talked about how generative artificial intelligence (AI) and large language models (LLMs) would revolutionize every field. In higher education, the reception was no different. Within the first six months, articles with titles like “I’m a Student. You Have No Idea How Much We’re Using ChatGPT” and “It’s Not Just Our Students—ChatGPT is Coming for Faculty Writing” appeared in *The Chronicle of Higher Education*. Sessions like “After ChatGPT: Religion and Pedagogy” were presented at the 2023 annual meeting of the American Academy of Religion. I saw similar titles appear in webinars, blogs, and frantic text messages from friends who were trying to figure out how to (once again) adapt their pedagogy at the graduate and undergraduate levels.

The future role of generative AI in theological education and theological librarianship remains open and complex. The tendency in many circles is either to over- or underestimate the realities and impact of such technology in the domain of religious life and information landscapes. This essay identifies and annotates a selection of materials that, I think, may aid theological librarians in approaching generative AI in their libraries from a moderated, historical, and theological position. I hope the essay and the resources discussed will fill in the gaps and provide a starting point for our guild. Even if we are not experiencing the results of AI-enabled search strategies—though I suspect that is not the case—I think it is safe to assume that in the next three to four years, the information landscape and researchers’ strategies will look very different from the way they have for the past two decades.

In what follows, I have organized the resources into four main areas that interest theological librarians: history, anthropology, ethics, and theology/religious studies. Many of the materials below are scholarly, but several are not, and that is by design. Some of the most salient frameworks can be found in popular literature on this issue. That may be due to the brevity of such pieces, but I suspect it has more to do with the clarity that writing for a broad audience requires. The AI learning curve can feel as steep as learning Koine Greek or Ecclesiastical Latin for those without a computer science background. As a result, I have compiled clear and concise works. Each section discusses items in reverse chronological order unless thematic organization dictates otherwise.

Brady Alan Beard is Head of Research and Instruction at Pitts Theology Library, Candler School of Theology, Emory University.

This bibliographic essay concludes with a reflection on the three places theological librarians can intervene in discussions of AI in theological and higher education. Those three areas are theological anthropology, ethics, and information technology.

ARTIFICIAL INTELLIGENCE: INTRODUCTIONS AND HISTORIES

To understand the relationship between theology and AI, we must know something about the history of artificial intelligence. AI goes back farther than Stanley Kubrick's Hal in *2001: A Space Odyssey*, and, indeed, the drive for a nonhuman intellect goes, perhaps surprisingly, beyond the 19th-century. Additionally, the drive for AI relies on assumptions built on religio-philosophical understandings of personhood, intelligence, and the soul. Dennis Yi Tenen's *Literary Theory for Robots: How Computers Learned to Write* (2024) connects the technical aspects of chatbots and other AI systems to the Turing machine; Wittgenstein's lectures; Persian mathematician and coiner of the term "algorithms," Muḥammad ibn Musa al-Khwarizmi; and Ada Lovelace, one of the first programmers and daughter of Lord Byron. Tenen argues that language and worldview are bound together in such a way that engineering software from soft AI, like autocorrect and spell-check, to universal AI requires a philosophical outlook that masks the crowds of programmers, technicians, and specialists—humans who labored to create a subservient machine mind. These humans necessarily have philosophical, ethical, and religious outlooks that inform their work.

Such outlooks are not only implicit. In *The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do* (2021), Erik J. Larson demonstrates that the contemporary conception of AI depends on the late 19th-century and early 20th-century works of mathematicians like David Hilbert, Kurt Gödel, Alan Turing, and Jack Good. These men and their colleagues, Larson argues, defined intelligence in a specific and problematic way, laying the groundwork for the worldviews of scholars who say machine intelligence will one day surpass human intelligence. Larson suggests the reality may be otherwise. By providing a narrow definition of what counts as intelligence, Turing and others limit the possibilities for intelligent machines. For instance, a computer that can play a game like chess well can be described as intelligent even if it cannot do anything else. In other words, the faster a computer wins a chess game, the more intelligent it is, even if it cannot do other tasks, such as offering an analysis of *A Tale of Two Cities*.

The history Larson describes demands that AI proponents and critics alike reconsider the nature of personhood and the definition of intelligence when applied to machines. It also requires that critics reevaluate the "AI mythology" (275) in popular discourse. Such reevaluation can only be done when the worldviews of individuals like Ray Kurzweil, Elon Musk, and others are deconstructed and analyzed philosophically and religiously. This, Larson writes, requires "an exploration of spiritual isolation" to probe the questions embedded in the AI myth (280).

Meredith Broussard's book *Artificial Unintelligence: How Computers Misunderstand the World* (2018) covers some of the same issues Larson raises, while providing an imminently approachable introduction to machine learning and AI that critically examines unbridled technological enthusiasm and assumptions that technology is always right—a phenomenon she refers to as "technochauvinism." This unexamined enthusiasm feeds assumptions that technological advancement is inevitable and that it will solve complex problems, a form of technochauvinism. Knowing the limits of technology can help us make better decisions about how to employ that technology.

While dated by the speed at which computer technology advances, Kevin Warwick's 2012 book, *Artificial Intelligence: The Basics*, provides practical chapters that outline definitions of intelligence,

classical AI, the philosophy of AI, and future developments. Warwick's definitions are a helpful starting point for the slippery terminology at work. Consider, for instance, his definition of classical AI as "based more on trying to get machines/computers to copy humans in tasks that, when humans do those tasks, *we deem them to be intelligent acts*" (58).¹ Who deems such acts to be intelligent, and why?² Nevertheless, at the end of the next chapter, he writes, "In some ways, computers have been able to outperform humans for many years, whereas in other ways—human communication, for example, as witnessed by the Turing Test—computers are perhaps not quite yet able to perform in exactly the same way as humans" (86). Once again, readers might ask who selects the tasks and defines success.

While ChatGPT and other chatbots have significantly improved since the time of Warwick's claim, the underlying point is well taken. AI is limited by outside factors, just as human activity is limited. Understanding fundamental aspects of AI, like those above, is crucial to resisting the overstated panic and praise common in the cultural discourse of AI. While the book is in some ways outdated, its approachability is still helpful for the nonspecialist.³ Warwick's book should be supplemented with the conference paper titled "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" by Emily Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell (2021). Bender et al. cover the history and technological developments of ChatGPT 2/3, bridging the gap between Warwick's book and recent developments in AI. The authors maintain a critical eye toward the environmental, societal, and cultural costs of AI programs based on LLMs and algorithms.

It is important to note that much AI theorizing comes from a select few labs at Massachusetts Institute of Technology (MIT) and other institutions. The works of these theorists and scientists provide excellent primary source material for exploring the fundamental philosophical motivations of some AI and machine learning work. Principally among these thinkers are Ray Kurzweil and Marvin Minsky. In many of these works, such as in Kurzweil's *The Age of Spiritual Machines: When Computers Exceed Human Intelligence* (1999), the spiritual, theological, and technological are intertwined and uncritically assumed.

The materials in this section lay a foundation for the following two sections. Two essential points should be laid out here. First, contrary to certain popular accounts, ChatGPT and generative AI did not appear *ex nihilo*. Generative AI has a long history that goes beyond the early foundations of programming and has been connected with information, knowledge creation, and writing for a long time. The explicit questions librarians must now consider have been latent throughout most, if not all, technological development for the last 50 years. Second, the development of generative AI was not inevitable but resulted from desire, research, and mechanics. Despite the tenor of many conversations around AI today, the trajectory is not inevitable, even if certain powerful voices want it to be. Shaping the future of generative AI requires understanding its past and clearly evaluating its usefulness. The history of AI is a history of people.

ARTIFICIAL INTELLIGENCE AND ANTHROPOLOGY

One of the most critical questions artificial intelligence raises is about the nature of humanity. In many theological institutions, such a question would be addressed under theological anthropology, which seeks to understand the theological meaning of human existence. Just as the theological anthropologist must confront the difference between humanity, God, and the rest of creation, so too, in a theologically informed bent, AI invites theologians to consider the nature of the human and human experience in contrast to machines.

In his 2021 book, *In Defence of the Human Being: Foundational Questions of an Embodied Anthropology*, Thomas Fuchs sets out to develop a humanism of embodiment that resists the reductionism of humanity so often found in discussions around artificial intelligence and machine learning. In Fuchs's terms, the human comprises embodiment and aliveness. These qualities mark it as distinct from machines (5). The human is a complex creature, not easily divided into body and mind, or in the technochauvinist metaphor, hardware and software. The human's ability to embody, in the fullest sense, a space distinguishes humanity from the machine (11). When terms like "intelligence" and "learning" are applied to machines, boundaries between human and nonhuman slip. The dangers of this slippage arise when machines are tasked with making increasingly substantive decisions about the daily lives of humans (44). Only by reconsidering the category of anthropology can humanity resist the danger.

What needs do AI and robots fill for humans? The answer to this question reveals more than simply the desire to make a subservient task-doer (the Slavic root for "robot," *robota*, relates to the word "work"). The psychological need such innovations fill relates to humans' ability to form attachments with one another. This psychological need, Kathleen Richardson argues, stems from a (not-so) latent annihilation anxiety that stems from a flattening of ontological distinctions between the mechanical and the human (2015, 5). This flattening has implications for social structures.

Any examination of anthropology would be remiss to exclude additional categories that frame human experience. In *Race after Technology: Abolitionist Tools for the New Jim Code* (2019), Ruha Benjamin argues that racism and technology have always been codependent. Algorithmic AI and machine learning codify racism by permeating the technology that underlies these services (17). Technology cannot disrupt social structures like race and racism, in part because it codifies them, thereby contributing to social ills like racism. AI, like anthropology and other social-scientific disciplines, requires theologians to take a critical stance toward the assumptions, language, and functions such technology assumes and obscures. By investigating the human alongside the machine, rather than just poking at technology, theologians can begin to make connections between AI and humanity.

Jonathan Tran (2018) is one theologian who has taken up the question of AI and theological anthropology. In his essay "The Problem Artificial Intelligence Poses for Humans," he refers to the assumption that machines will replace humans as an "empty proposition." In his brief essay, Tran simplifies what makes machines intelligent and why it matters that machine learning is said to replicate human thought processes. Such learning depends on the information intake. In short, the issues that prevent effective human learning will also apply to machines. Naturally, these two features lead to several ethical questions about what it means to be human.

Up to this point, most references in this essay feature academics. My first exposure to the importance of a theological approach to AI came in the form of a memoir, *God, Human, Animal, Machine: Technology, Metaphor, and the Search for Meaning* (2021), by Meghan O'Geiblyn, essayist and tech journalist. While reflecting on her relationship with a robot in the form of a small dog, O'Geiblyn provides a powerful yet nuanced history of AI, our technological age, and the religious questions that our desire for ever more effective technology pushes to the fore. O'Geiblyn deftly shows how technology has changed societal assumptions about humanity and religion.

Whatever the future of generative AI holds, it will undoubtedly require that theological librarians investigate the human part of human information activity. By necessity, librarians confront the differences between machine and human processing daily. Cataloging, systems design, and reference services require thoughtful respect for the differences between what a human and a machine can do. Theological librarians also know of the rich literature behind the question, "What is the human?" AI allows us to expand our expertise into conversations beyond our daily workflows.⁴

ARTIFICIAL INTELLIGENCE AND ETHICS

Imagine conversations ten years ago about algorithms. Most people probably would not know what they are or how they impact their lives. Relevance was simply assumed in search results. Since then, however, algorithms have burst into popular discourse. We hear about them everywhere: in the news, on social media, and at the dinner table. In her book, *Algorithms of Oppression: How Search Engines Reinforce Racism* (2018), Safiya Umoja Noble describes the dangers of participating in an algorithmic social and media environment. Chief among these dangers is the “black box” quality of proprietary algorithms that impact daily lives. Noble’s work is an acute diagnosis of the dangers of algorithms in search engines. Her warning is framed by the importance of algorithmic transparency and justice as technology advances and AI are normalized. Some of the most critical questions users can ask of artificial intelligence and machine learning are as follows: “Where does this data come from?” “Who controls it?” and “Who evaluates it” (Widder, West, and Whittaker 2023).

Like Noble’s work, Cathy O’Neil’s book *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (2017) outlines the multiple ethical and societal dangers that arise from an unregulated algorithmic culture. In contrast to Noble’s book, which outlines how algorithms impact search and reimpose bias in knowledge production, O’Neil’s book illustrates how algorithms, data, and AI impact our lives in ways we may not even know. One of the best examples comes from her chapter on the role of algorithms in auto insurance rate setting. This data often has little to do with the individual driver and more to do with the “data bucket” into which drivers are sorted. When, where, and how long a driver drives are as important to these systems as a safe driving record. The failure of transparency around data and its applications sets a dangerous precedence for AI.

Artificial intelligence’s ethical implications include how algorithms and data feed opaque systems. When discussing AI with students, I find it often useful to ensure they understand the concrete costs of using artificial intelligence in their coursework. Kate Crawford’s book *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* (2021) spells out these costs—the chapters in this book detail the impact of AI on Earth, labor, and society. Chapter 1, for instance, shows that chatbots are significantly water- and mineral-reliant—becoming one of the most natural resource-intensive industries and potentially surpassing the water usage of entire countries.⁵ Chapter 2 examines labor and shows that much of what we call “artificial intelligence” is simply human intervention removed from our direct view. Social media content moderators must be trained to recognize violent or harmful content. These programs require humans to code the content in ways machines can recognize. Often, these laborers work psychologically taxing jobs in countries with few worker protections and minimal compensation, as contractors for major U.S. corporations. Crawford deftly demonstrates that AI is simply power by another name in an increasingly digital world.⁶

As the above items demonstrate, the data environment for AI and machine learning is complex, and there are multiple opportunities for bias and harm to enter these systems. In “A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle,” Harini Suresh and John Gutttag (2021) outline a concise overview of machine learning and identify seven sources of harm in machine learning: historical, representation, measurement, aggregation, evaluation, learning, and deployment biases. Each stage requires mitigation. Most important, however, the authors clearly articulate that data is not simply the only source of harm in machine learning and that human choice must also be addressed through the life cycle of machine learning.

With an understanding of the ethics and morality raised by automation, theological librarians can begin to engage their researchers to meaningfully integrate AI into the life cycle of their work and institutions. Undoubtedly, AI has already burrowed into the research practice of students and

faculty alike, but most have likely not integrated AI into their work meaningfully. By “meaningfully,” I mean that users integrate AI into their work in ways aligned with their own and their guild’s ethics. I also suggest they responsibly incorporate AI into their work, checking for errors, accounting for bias, and critically engaging with generative AI.

ARTIFICIAL INTELLIGENCE AND RELIGION

To this point, I have touched on topics and disciplines related to theology and religious studies, but these works do not treat those fields explicitly. This was necessary to demonstrate that the questions the theologian, religion scholar, or librarian brings to artificial intelligence are not entirely out of place. In contrast, the above sections show that data scientists, librarians, and literary scholars are concerned with questions that motivate much of theological education (Dorobantu 2022). These questions relate to the nature and future of humanity, history, and ethical engagement with others and natural resources. Finally, this essay turns to an explicit discussion of works that frame AI and technological development as religious and theological questions. In other words, the works below ask, “What does AI mean theologically?” The answer to that question requires that theological librarians apply methods and techniques of theology and religious studies to “the myth of AI.”

Robert M. Geraci addresses the technochauvinism latent in many optimistic perspectives of how AI will impact humanity, in his 2022 work *Futures of Artificial Intelligence: Perspectives from India and the U.S.*, Geraci demonstrates that notions of apocalypse, sometimes referred to as “the Singularity,” often accompany AI development. Singularity is the predicted point at which machines will outpace humans and become like gods to them (165). This line of reasoning is often assumed and popular. Geraci argues that it is neither inevitable nor natural (167). In his work, Geraci explores AI in an Indian context and revisits the AI apocalypse from a Hindu and global perspective. Science and religion intersect at many points, AI among them.⁷

In 2010, Geraci published a more fulsome discussion of apocalyptic AI that anticipated his work in 2022. In *Apocalyptic AI: Visions of Heaven in Robotics, Artificial Intelligence, and Virtual Reality* (2010), Geraci argues that in addition to being a religious movement, Apocalyptic AI is “a strategy for enhancing the social power of technoscientific researchers” (140).⁸ Apocalyptic AI, according to Geraci, is the genre of science writings—especially those by Marvin Minsky, Hans Moravec, Kevin Warwick, Hugo de Garis, and Ray Kurzweil—that bring robots and artificial intelligence into public discourse (1).⁹ It has scholars and evangelists, and is concerned with life on Earth. Apocalyptic AI can be found in works by theorists who seek to overcome the human condition through technology. Like religion, Apocalyptic AI influences technology and culture, distinguishes between the human and the nonhuman, and attempts to transcend dualistic divisions of human experience.

In what is perhaps one of the most important analyses of religion, theology, and AI, Damien P. Williams (2022) argues in his unpublished dissertation, “Belief, Values, Bias, and Agency: Development of and Entanglement with ‘Artificial’ Intelligence,” that artificial intelligence cannot be understood correctly without reference to its logics, which function like religious and occult logics. These logics obscure and mystify the actual technical processes and, as a result, create harm through the digital reification of data bias and machine decision-making. Only when these logics are unmasked through religious and theological analysis can we demythologize their uses, potentials, and harms. To do this, we must rely on the skills and insights of religious studies and theology: critical evaluation of belief, careful reading of myth, and understanding of power dynamics. Recontextualizing AI within social knowledge is one of the tasks Williams argues can best be accomplished by scholars of religion.¹⁰

Calvin Mercer and Tracy J. Trothen’s 2021 *Religion and the Technological Future: An Introduction to Biohacking, Artificial Intelligence, and Transhumanism* examines artificial intelligence in the broader context of technological advancement and transhumanism. The possibility of “strong AI” (i.e., Artificial General Intelligence [AGI]) as opposed to “weak AI” (i.e., narrow computation of select tasks like playing chess) stands to radically alter discussions of religion and AI. However, we currently have weak AI. The point at which machine intelligence surpasses human intelligence, “the Singularity” of Ray Kurzweil, is hotly debated. The question of AGI for the religionist and theologian is what AGI’s relationship would be to conceptions of God (186–189). This question is especially pertinent since the language used about the Singularity is routinely spiritual and theological. As a result, the questions raised are definitional and ethical, and are concerned with ritual from a theological and anthropological perspective.¹¹

In “The AI Creation Meme: A Case Study of the New Visibility of Religion in Artificial Intelligence Discourse,” Beth Singler (2020) argues that the discourse of artificial intelligence resounds with religious continuities. These continuities can be found in the cultural artifacts—textual, verbal, and imagistic—and discourses of AI. Despite the post-secular assumptions of contemporary technological discussions, religious tropes appear regularly. In doing so, they offer social value. Far from pushing spirituality and enchantment from cultural discourse, the perceived potential of AI is assumed and put to work in the zeitgeist.

In another essay from 2020, “‘Blessed by the Algorithm’: Theistic Conceptions of Artificial Intelligence in Online Discourse,” Singler argues that discourse around AI and algorithms indicates that our modes of thinking about AI are implicitly religious. This thinking is revealed by online phrases, such as, “I’ve been blessed by the algorithm,” on social media. The insights of Singler’s study show one of numerous ways AI fits within and alongside conceptions of God. This notion is especially prevalent in overtly religious spaces like the Turing Church and requires that scholars engage AI with approaches informed by religious studies and theology.¹²

Behind the façade of algorithms, AI, and technological production, humans employ manual labor to build complex systems and tools. In the same way, “the cloud” is simply other people’s computers, algorithms, big data, and automation; it is just other people who tag, train, record, maintain, and organize algorithmic systems. Social change because of automated systems is not inevitable or omnipotent. While the technologies that make up generative AI are referred to in theological language and hushed religious tones, these systems are not inevitable or necessarily natural. They can be valorized, employed, resisted, adapted, ceased, ignored, and forgotten.

CONCLUDING THOUGHTS: WHY SHOULD THEOLOGICAL LIBRARIANS CARE?

In many respects, theological librarians find themselves at the front of emerging technology as it impacts their organizations. As mediators of information, faculty, administrators, students, and other staff often ask us for insights or recommendations on the newest technological developments, academic integrity (i.e., avoiding plagiarism), and education. Artificial intelligence has not broken this expectation. Additionally, some librarians are seeing their services and staffing cut while chatbots and generative AI interfaces replace essential services. Michael Hanegan and Chris Rosser (2023) have argued that theological librarians have much to say about this new development in theology. Even librarians not concerned about implementing AI in their libraries have, no doubt, experienced working with researchers who integrate ChatGPT or Gemini into their search strategies.

Theological librarians should care about AI and its impact on our discipline, for several reasons (Campbell and Cheong 2023). First, by the nature of its inquiry, AI and machine learning raise questions of what it means to be human. The nature of the human and the rest of the natural and supernatural world is a central question of theology. Second, AI raises ethical and moral questions that also appear in the core curriculum of our institutions. Third, AI is here to stay. It has already been here and will continue to frame our information landscapes (Iacovitti 2021).

By asking fundamental questions about what intelligence is, whether it can be artificially manufactured, and how it can be deployed, computer science asks a question that has been at the heart of many, if not all, religions and certainly part of the foundation of Christian theology. On one hand, the very phrase “artificial intelligence” presupposes two things: that we agree on a definition of “intelligence” and that we agree on a definition of “artificial.” As a close reading of any of the works mentioned in this essay would suggest, neither question can be answered quickly or simply. In some cases, it seems the question of the soul has just been repackaged as the question of intelligence in the 21st-century.

The phrase “artificial intelligence,” like “the Internet,” masks a complex series of interrelated mechanisms and humans attempting to maximize communication and information networks. In its innocuous use, the phrase simplifies complex functions. It causes us to forget that on the other side of the screen are programmers, trainers, developers, cables, servers, minerals, buildings, and a myriad of other physical attributes that make AI possible. In its more disingenuous use, “artificial intelligence” masks the human and resource costs to promote the next big thing in tech. In the same way automated content moderation and self-driving vehicles simply remove the laborer from sight, AI can prevent us from seeing the actual cost of something as straightforward as querying ChatGPT.

As a field, theology is replete with ethical and moral convictions. Thus, it is only reasonable that theological researchers should consider the ethical implications of deploying AI in their research methods in addition to concerns about academic integrity. Since it is impossible to retreat behind the walls of the seminary library and not engage AI, we can at least reflect on our use of new technologies in ways consistent with the values of our various disciplines, in both theology and library sciences.

Although it may seem AI and algorithmic machine learning arrived from out of nowhere with the first announcements of ChatGPT 3, that is simply not the case. The desire for improved communication and algorithmic writing has existed throughout human history. There is even an argument to be made that automated writing has long been a literary dream. Regardless of one’s attitude toward chatbots and automated writing, they are here to stay and are indeed already present in our email, cell phones, and autocomplete capabilities. Like word processors and laptop computers, this technology is not going anywhere. The students and researchers we serve will certainly come to rely on this technology in their research strategies, if they do not already.

REFERENCES

- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?” In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*: 610–623. FAccT ’21. New York: Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>.
- Benjamin, Ruha. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge: Polity.

- Bogost, Ian. "The Cathedral of Computation." *The Atlantic*. January 15, 2015. <https://www.theatlantic.com/technology/archive/2015/01/the-cathedral-of-computation/384300/>.
- Broussard, Meredith. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. Boston: MIT Press. <https://doi.org/10.7551/mitpress/11022.001.0001>.
- Campbell, Heidi A., and Pauline Hope Cheong. 2023. *Thinking Tools for AI, Religion & Culture*. e-book. <https://doi.org/10.21423/oak/1969.1/198190>.
- Checketts, Levi. 2022. "Artificial Intelligence and the Marginalization of the Poor." *Journal of Moral Theology* 11, Special Issue no. 1: 87–111. <https://doi.org/10.55476/001c.34125>.
- Copeland, Jack. 1993. *Artificial Intelligence: A Philosophical Introduction*. Oxford: Wiley-Blackwell.
- Crane, Tim. 1995. *The Mechanical Mind: A Philosophical Introduction to Minds, Machines and Mental Representation*. New York: Penguin.
- Crawford, Kate. 2021. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, CT: Yale University Press.
- Dorobantu, Marius. 2022. "Artificial Intelligence as a Testing Ground for Key Theological Questions." *Zygon: Journal of Religion and Science* 57, no. 4: 984–999. <https://doi.org/10.1111/zygo.12831>.
- Fetzer, James H. 2001. *Computers and Cognition: Why Minds Are Not Machines*. Studies in Cognitive Systems Vol. 25. Dordrecht, Netherlands: Kluwer Academic.
- Fuchs, Thomas. 2021. *In Defence of the Human Being: Foundational Questions of an Embodied Anthropology*. Oxford: Oxford University Press.
- Geraci, Robert M. 2010. *Apocalyptic AI: Visions of Heaven in Robotics, Artificial Intelligence, and Virtual Reality*. New York: Oxford University Press.
- . 2022. *Futures of Artificial Intelligence: Perspectives from India and the U.S.* New Delhi: Oxford University Press.
- . 2024. "Religion among Robots: An If/When of Future Machine Intelligence." *Zygon: Journal of Religion and Science*. <https://doi.org/10.16995/zygon.10860>.
- Hanegan, Michael, and Chris Rosser. 2023. "Artificial Intelligence and the Future of Theological Education (Version 2.0)." <https://iparchitecture.notion.site/Artificial-Intelligence-and-the-Future-of-Theological-Education-9e035aeb8710406c85f1144cf0e9d1e6>.
- Iacovitti, Giovanni. 2022. "How Technology Influences Information Gathering and Information Spreading." *Church, Communication and Culture* 7, no. 1: 76–90. <https://doi.org/10.1080/23753234.2022.2032781>.
- Kurzweil, Ray. 1999. *The Age of Spiritual Machines*. New York: Viking.
- Larson, Erik J. 2021. *The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do*. Cambridge: Belknap Press.

- McCorduck, Pamela. 2004. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. Natick, MA: A K Peters.
- Mercer, Calvin and Tracy J. Trothen. 2021. *Religion and the Technological Future: An Introduction to Biohacking, Artificial Intelligence, and Transhumanism*. New York: Palgrave Macmillan.
- Midson, Scott A. 2018. *Cyborg Theology: Humans, Technology and God*. Library of Modern Religion 56. New York: I.B. Tauris.
- Noble, Safiya Umoja. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: NYU Press.
- O’Gieblyn, Meghan. 2022. *God, Human, Animal, Machine: Technology, Metaphor, and the Search for Meaning*. New York: Anchor.
- O’Neil, Cathy. 2017. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- Richardson, Kathleen. 2015. *An Anthropology of Robots and AI: Annihilation Anxiety and Machines*. Routledge Studies in Anthropology. New York: Routledge.
- Shneiderman, Ben. 2022. *Human-Centered AI*. Oxford: Oxford University Press.
- Singler, Beth. 2017. “An Introduction to Artificial Intelligence and Religion For the Religious Studies Scholar.” *Implicit Religion: Journal for the Critical Study of Religion* 20, no. 3: 215–231. <https://doi.org/10.1558/imre.35901>.
- _____. 2020a. “The AI Creation Meme: A Case Study of the New Visibility of Religion in Artificial Intelligence Discourse.” *Religions* 11, no. 5: 253. <https://doi.org/10.3390/rel11050253>.
- _____. 2020b. “‘Blessed by the Algorithm’: Theistic Conceptions of Artificial Intelligence in Online Discourse.” *AI & Society* 35, no. 4: 945–955. <https://doi.org/10.1007/s00146-020-00968-2>.
- _____. 2022. “Origin and the End: Artificial Intelligence, Atheism, and Imaginaries of the Future of Religion.” In *Emerging Voices in Science and Theology*: 105–120. New York: Routledge.
- Suresh, Harini, and John Guttag. 2021. “A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle.” In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*: 1–9. EAAMO ’21. New York: Association for Computing Machinery. <https://doi.org/10.1145/3465416.3483305>.
- Tenen, Yi. 2024. *Literary Theory for Robots: How Computers Learned to Write*. Norton Shorts. New York: W. W. Norton & Company.
- Tran, Jonathan. 2018. “The Problem Artificial Intelligence Poses for Humans.” *The Other Journal* 29. <https://theotherjournal.com/2018/06/the-problem-artificial-intelligence-poses-for-humans/>.
- Widder, David Gray, Sarah West, and Meredith Whittaker. 2023. “Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI.” *SSRN Scholarly Paper*. <https://doi.org/10.2139/ssrn.4543807>.
- Williams, Damien Patrick. 2022. “Belief, Values, Bias, and Agency: Development of and Entanglement with ‘Artificial Intelligence.’” PhD diss., Virginia Polytechnic Institute and State University. <http://hdl.handle.net/10919/111528>.

- . 2023a “Bias Optimizers.” *American Scientist* 111 no. 4: 204–207. <https://www.americanscientist.org/article/bias-optimizers>.
- . 2023b. “Any Sufficiently Transparent Magic...” *American Religion* 5, no. 1: 104–110. <https://doi.org/10.2979/amerreli.5.1.06>.

NOTES

- 1 Emphasis added.
- 2 Levi Checketts (2022) addresses the problem of intelligence as defined by a small section of white, wealthy, able-bodied men in “Artificial Intelligence and the Marginalization of the Poor.”
- 3 Several other older yet foundational works should be mentioned here. They include Copeland (1993); Crane (1995); Getzer (2001); and McCorduck (2004). These titles and their publication dates demonstrate that the questions many of us are now asking of AI, machine learning, and philosophy of mind have been in discussion for the past several decades even if many theologians and librarians haven’t been paying attention to them.
- 4 See also Shneiderman (2022).
- 5 For more on the water usage of AI, see Pengfei Li, Jianyi Yang, Mohammad A. Islam and Shaolei Ren’s preprint (2023).
- 6 In addition to Crawford’s book, see Williams (2023a).
- 7 For more on Geraci’s work on the intersection of futurism, AI, and religion see his 2024 essay “Religion among Robots: An If/When of Future Machine Intelligence.”
- 8 For a popular level analysis of this point, see Bogost (2015).
- 9 For another example of Apocalyptic AI in action see the May 30, 2023, open letter published by the Center for AI Safety. Bill Gates, Sam Altman, and Ted Lieu are signatories. <https://www.safe.ai/work/statement-on-ai-risk>.
- 10 For a much-condescended version of Williams’s dissertation, see Williams 2023b.
- 11 Many of these questions are also taken up in Midson (2018). In his work, Midson concludes that cyborg theology must concern itself with context and actors, narratives and “(hi)stories” that shape cultural understanding of technology, center people, resist easy assumptions, remain critical, engage inclusive mythologies and anthropologies, and emphasize relationality (188–198).
- 12 See Singler’s (2017; 2022) two additional essays.