VOL. 63 NO. 2 | SPRING 2021

Southwestern
JOURNAL OF THEOLOGY

The Doctrine of Humankind

FAKE AND FUTURE "HUMANS":
Artificial Intelligence, Transhumanism, and the Question of the Person

Jacob Shatzer*

Today, we are pulled in many different directions on what it means to be human. On one hand, a radical constructivism rules: I choose and build my identity, and for you to use any category to describe me that I have not chosen is an offense and affront. From this perspective, there really is not anything solid that determines what it means to be human, and we can build ourselves into whatever we want. Yet, on the other hand, we construct arguments and movements based on a shared humanity. Furthermore, as we develop more and more sophisticated technology, we cannot help but begin to refer to these technologies as though they bear some marks of what it means to be human. Our digital assistants have names, we use smart robots to provide companionship to the elderly, and we exult at how "intelligent" (a human-oriented trait) our systems are becoming, whether it is the artificial intelligence (AI) built into a thermostat or a robot. When it comes to ourselves, we want to determine our humanity, but when it comes to our machines, we are quick to use static human traits in order to describe the greatness of the works of our hands.

Our technological creations pull in multiple directions at our doctrine of humanity. A robust doctrine of humanity will give us a foundation from which to address these challenges, but these challenges will also affect—or perhaps infect—our understanding of what it means to be human. A basic understanding of AI ("fake humans") and transhumanism ("future humans") will press a variety of challenges onto our theological anthropology, both in what it means to be human and how we might consider and pursue human flourishing in light of these developments.

The history of technology is certainly complex, and there is some

* Jacob Shatzer serves at Union University as associate professor of theological studies and associate dean of the school of theology and missions.

debate as to whether technology is neutral or not. However, even if technology on its own is thought of as neutral, it is actually impossible for any of us to ever engage technology "on its own." We always encounter technologies embedded within human cultures, which do carry and cultivate values and ethics.[1] Not only do we always encounter technologies as embedded within cultures, we also struggle to be able to notice the ways that these devices impact our ability to see and desire the good,[2] or in simpler language, to avoid sin and honor Christ.

This issue is particularly important because of the age in which we live. Byron Reese argues in his book *The Fourth Age* that while we think we have seen change in the last 100 years, we really have not. We're still basically the same as people 5,000 years ago. He sees three main ages of humanity so far: fire (100,000 years ago); agriculture, cities, war (10,000 years ago); wheel and writing (5,000 years ago). We're on the cusp of the fourth: AI and robots.[3] Reese provides a perspective not present in many other treatments: he emphasizes that many proponents of different futures depend on unexamined assumptions about what it means to be human. We have to answer that question before we can understand the way to direct AI and robotics, and before we can really decide if these changes will be positive or not.

In other words, the other articles in this issue on theological anthropology have just as much to do with our response to AI and transhumanism as this article does! The questions Reese raises, from a secular perspective, show the fundamentally theological nature of the issue: "The confusion happens when we begin with 'What jobs will robots take from humans?' instead of 'What are humans?' Until we answer that second question, we can't meaningfully address the first."[4] With that in mind, in what follows we will look at AI and transhumanism in order to get a better view of the touchpoints and challenges that they raise for Christian theological anthropology.

[1]For more on the history of technology and understanding the connection between technology and ethics, see Eric Schatzberg, *Technology: Critical History of a Concept* (New York: Oxford, 2019).

[2]For an interesting take on this from a secular philosophical angle, see Shannon Vallor, *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting* (New York: Oxford, 2016).

[3]Byron Reese, *The Fourth Age: Smart Robots, Conscious Computers, and the Future of Humanity* (New York: Simon & Schuster, 2018).

[4]Reese, *The Fourth Age*, xi.

By going this route, we will begin to see ways that our doctrine of humanity is informed by these challenges and also forms our response to them.

I. ARTIFICIAL INTELLIGENCE

Artificial intelligence is a large and changing field that has also had a broad and varied history, both in reality and in pop-culture expressions. To consider how AI might develop and impact our thinking about what it means to be human, we will have to clear the ground a bit to make sense of what we are talking about.

In his recent book *2084*, John Lennox defines key terms related to AI, and we will rely on his explanation of "robot," "AI," and "algorithm." First, "a robot is a machine designed and programmed by an intelligent human to do, typically, a single task that involves interaction with its physical environment, a task that would normally require an intelligent human to do it."[5] This definition is pretty straightforward and unsurprising. Second, Lennox defines AI in two ways: "The term is now used both for the intelligent machines that are the goal and for the science and technology that are aiming at that goal."[6] Third, Lennox expands on "algorithm" using the *OED*: "a precisely defined set of mathematical or logical operations for the performance of a particular task."[7] He points out that such concepts can be found as far back as Babylonia in 1800–1600 B.C., though obviously not coded into digital technology. The key feature of the algorithm is that "once you know how it works, you can solve not only one problem but a whole class of problems."[8] Lennox follows up with some mathematical examples, such as instructions for various steps to arrive at, say, the greatest common denominator of two numbers. You can follow the steps for any set of two numbers and it will work. Algorithms, then, are embedded within software that uses them to interact with and evaluate different data inputs.[9] This type of system can take any input that can be digitized—sound, text,

[5]John C. Lennox, *2084: Artificial Intelligence and the Future of Humanity* (Grand Rapids: Zondervan, 2020), 16.

[6]Lennox, *2084*, 17.

[7]Lennox, *2084*, 19.

[8]Lennox, *2084*, 19.

[9]For a deeper explanation of algorithms, see Kartik Hosanagar, *A Human's Guide to Machine Intelligence: How Algorithms Are Shaping Our Lives and How We Can Stay in Control* (New York: Viking, 2019).

images—apply a set of steps to that data, and come up with some sort of conclusion. That conclusion can include or lead to action. Algorithms are vital to understand because they are at the center of how AI works.

There are four main categories of algorithms. First, prioritization algorithms make an ordered list, say, of items you might want to buy or shows you might want to watch. Classification algorithms take data and put it into categories, perhaps automatically labelling photos for you, or isolating and removing inappropriate content from social networks. Association algorithms finds links and relationships between things. Filtering algorithms isolate what is important (say, eliminating background noise so a voice-enabled assistant like Siri can "hear" what you're saying).[10]

Let's take a look at a quick example. A smart thermostat can take in pieces of information, such as the current temperature in a room, the time of day, and the weather forecast for the day, run that data through a series of steps, and determine how long and how high to run the furnace to reach a certain temperature. (Smart thermostats can also take in data on household inhabitants over a period of time to determine what that certain temperature should be.) To incorporate our definitions above, the thermostat would be a type of robot, an example of AI, running an algorithm to achieve climate bliss.

Typically, experts divide AI into "narrow" AI and "general" AI, and our thermostat serves as an example of "narrow." A "narrow" intelligence can be taught or programmed to do something. A "general" intelligence can be taught or programmed to do anything.[11] For example, a robot vacuum is able to do one thing: clean up. Now, it certainly relies on various elements, including reading data on mapping a room, and even things like whether its bin is full. But it basically does one thing; no one is worried about their Roomba running away from home and joining the circus.

A general intelligence, on the other hand, is able to adapt and learn a variety of actions. Some thinkers describe artificial general intelligence (AGI) as being able to do anything that humans can do, but that is primarily because humans serve as the standard for the

[10]Hannah Fry, *Hello World: Being Human in the Age of Algorithms* (New York: Norton, 2018), 8–10.

[11]Lennox, *2084*, 13.

ability to adapt and adopt different ways of doing things and viewing the world. In all likelihood, an AGI would quickly surpass human abilities in many areas, thus rendering this comparison less useful.

We must take one more step in understanding AI to see the complexity and potential growth of this field. My first introduction to robotics occurred when I was in second grade. My class went on a field trip to North Idaho College. We worked with some simple robots that could move and pick up items. Our challenge: program the robots to navigate a course and retrieve an item. How far forward before a right turn? How much more before the next turn? Etcetera. Early advances in AI were made with this same method. Humans were creating algorithms, steps of instructions (vastly more complicated than my second-grade robot example) to allow robots and other smart machines to interact with their environment in desired ways. This is what most of us think about when we think about AI: human programmers teaching robots to do amazing things.

That is the way that it worked for a while. The history of AI provides helpful context in understanding what we should come to expect. Most people are aware of Moore's Law, which relates to the (generally accurate) rule of thumb that computer power doubles every 18 months (this being related to the construction of microchips). Many assume that AI, since it relies on computing power, has increased at a similar, steady rate, for the last 50 years. That is simply not the case.

Artificial intelligence hit a series of walls—what is referred to as "AI winters"—for two main reasons. First, creating algorithms is really complicated, and some tasks were just too complex for humans to "crack" with the instructions they could embed in an algorithm. Second, computing power, speed, and storage are not infinite. In other words, we reached the outer limit of our ability to "write" complex instructions, and we didn't have the computing power to process them quickly and at scale. But this "AI winter" came to an end in the early 2000s.

Recent advances in AI—its emergence from "winter"—have occurred because of changes in these two areas. The second one is obvious: computers are faster and more powerful, and data storage is exponentially larger now. But the problem of creating algorithms wasn't as simple as waiting for Moore's Law to catch up. The advent

of "machine learning" has led to the great growth in AI in the last ten to fifteen years. The "rule-based algorithms" that humans can create directly are being replaced by "machine-learning algorithms."[12] Basically, AI scientists have gone from creating algorithms for desired outcomes to creating learning algorithms: ways to set up an AI to learn for itself.[13] This occurs by "training" the AI on a set of real-world data. Through machine learning, the AI is able to identify patterns and create algorithms that match those patterns.[14] Once that is done, the AI can be fed pieces of data, and it will use its newly created algorithm to determine the relevant action or outcome. It is predicting what is most likely the proper outcome based on the dataset it used to determine the pattern and algorithm.[15]

Some argue that we should view AI not as intelligence but as prediction: an algorithm takes inputs and, based on patterns recognized within the data, makes a prediction on what the output should be. This could be a prediction about the answer to a question, or a prediction about whether to turn or brake, or a prediction about consumer behavior. AI will make prediction cheaper, which will mean businesses can do other things better. At some point, cheap prediction might change business models drastically.[16] One example of this is Amazon's work in "anticipatory shipping." There could come a point when Amazon's AI is so good at predicting what consumers want that it is more beneficial for them to simply ship things before people shop. It knows what you want; it sends it. Sure, sometimes it would be wrong, but once its correct predictions cross a certain threshold, it is actually more financially feasible for Amazon to ship and then allow returns on what it gets wrong. Their profit would be so high based on the increased number of items people would buy from them rather than elsewhere that it would be worth eating the costs of returns the times when it gets it wrong.[17]

---

[12]Fry, *Hello World*, 10.

[13]I am of course simplifying here. To get a better grasp of the different types of algorithms and approaches to this aspect of AI, see Pedro Domingos, *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World* (New York: Basic, 2015).

[14]Darrell M. West, *The Future of Work: Robots, AI, and Automation* (Washington, D.C.: Brookings Institution Press, 2018), 24.

[15]See Ajay Agrawal, Joshua Gans, and Ayi Goldfarb, *Prediction Machines: The Simple Economics of Artificial Intelligence* (Boston: Harvard Business Review Books, 2018).

[16]See Agrawal, Gans, and Goldfarb, *Prediction Machines*.

[17]Agrawal, Gans, and Goldfarb, *Prediction Machines*, 16.

We will look at some challenges below, but one jumps out here immediately. Machine learning is powerful, but part of its genius is that it works without having human programmers setting it up. In many cases, we're not really sure how these algorithms work. This can lead to biases or other problems. In other words, bad data can lead to bad machine learning, which can then perpetuate the same problems. As one scholar puts it, "When a new technology is as pervasive and game changing as machine learning, it's not wise to let it remain a black box. Opacity opens the door to error and misuse."[18] But the very nature of machine learning makes transparency difficult. Scientists often are not sure exactly how the AI has trained itself on the data set, or whether the data set itself harbors problematic assumptions.

We have not even begun to consider at what point this "artificial intelligence" becomes something meriting a new category. Technologists are already dreaming and planning about creating consciousness. As one puts it, "*techno-optimism about machine consciousness*… is a position that holds that if and when humans develop highly sophisticated, general purpose AIs, these AIs will be conscious."[19] Schneider uses the "precautionary principle" to argue that if we have any reason to believe an AI to be conscious, we should extend the same rights to it that we would to other sentient beings.[20] In fact, she argues that we should be really careful *not* to create consciousness and should thus limit our development of AI. While some are concerned about AI developing to merit something mirroring human rights, others are looking to technology to change radically what humans are.

II. TRANSHUMANISM

If we imagine a Venn diagram, AI and transhumanism would be their own circles, but there would certainly be overlap. We need this image, because we do not want to assume that all AI is part of transhumanism, nor do we want to assume that transhumanism is only about merging humans with AI. Both are bigger, but they are related. And, as we will see later, they produce some of the same

[18]Domingos, *Master Algorithm*, xvi.

[19]Susan Schneider, *Artificial You: AI and the Future of Your Mind* (Princeton, NJ: Princeton University Press, 2019), 23.

[20]Schneider, *Artificial You*, 67.

existential quandaries for us.

At root, transhumanism is about harnessing a broad range of enhancement technologies in order to bootstrap humans to the "next step" in the evolutionary process. Lennox quotes a character in Dan Brown's novel *Origin*, who speaks this way about transhumanism: "New technologies… will forever change what it means to be *human*. And I realize there are those of you who believe you, as *Homo sapiens*, are God's chosen species. I can understand that this news may feel like the end of the world to you. But I bet you, please believe me . . . the future is actually much *brighter* than you imagine."[21] This quotation captures both the essence of transhumanism—changing what it means to be human—and also the inescapable religious dimension. Transhumanists, by-and-large, see all religions as opiates of the people distracting from pain and preventing or denigrating the very advances that provide the only "true" hope. (Yet this stance itself is a religious one!)

Even though religion is the *persona non grata* of transhumanism in most cases, more Christians are finding common cause with transhumanism. One group is more theologically progressive, proposing "post-anthropologies" that emphasize "posthuman subjectivity and relationality, multiple embodiments, and hybridity as its key components" and goes so far as to propose a "cyborg Christ" as the center of a posthuman Christology.[22] Most evangelical Christians will not find such proposals alluring due to their radical theological innovation. Theological engagement with such groups will invite further research and thought from evangelical theologians and ethicists, but this response remains mostly peripheral among Christian responses to transhumanism.

However, a growing number of Christians identify with the transhumanist movement and seek to support it theologically without going quite as far in theological innovation.[23] "Christian Transhumanists" have founded an organization and gather in an annual conference. Engaging their thought is more important at

[21] Lennox, *2084*, 45–46.

[22] Jeanine Thweatt-Bates, *Cyborg Selves: A Theological Anthropology of the Posthuman* (New York: Routledge, 2012), 13.

[23] For a recent edited volume that covers a variety of perspectives on the issue, see Steve Donaldson and Ron Cole-Turner, eds., *Christian Perspectives on Transhumanism and the Church: Chips in the Brain, Immortality, and the World of Tomorrow* (New York: Palgrave McMillan, 2018).

this stage because their arguments and thinking are more likely to gain traction in evangelicalism broadly.

A Christian Transhumanist is "someone who advocates using science & technology to transform the human condition—in a way consistent with, and as exemplified by, the discipleship of Christ."[24] They choose to use "transhumanism" intentionally, believing that it provides a touchpoint for conversation with leading-edge thinkers in science and technology. According to their website, "[Transhumanism] originates with Dante in 1320, winds through Christian history, and is picked up in the work of Jesuit priest and paleontologist Pierre Teilhard de Chardin. Teilhard's friend Julien Huxley uses the term in 1957 in attempt to define a philosophy of humanity's ongoing transformation. This leads to secular transhumanism, as it is understood today." Further, the group thinks that it can "promote positive engagement between Christianity and the leading edges of scientific & technological thought."[25]

The group's statement of faith is fairly short but important:

> As members of the Christian Transhumanist Association:
>
> **1. We believe that God's mission involves the transformation and renewal of creation** including humanity, and that we are called by Christ to participate in that mission: working against illness, hunger, oppression, injustice, and death.
>
> **2. We seek growth and progress along every dimension of our humanity:** spiritual, physical, emotional, mental—and at all levels: individual, community, society, world.
>
> **3. We recognize science and technology as tangible expressions of our God-given impulse to explore and discover** and as a natural outgrowth of being created in the image of God.

[24] "Frequently Asked Questions," *Christian Transhumanism Website*, https://www.christiantranshumanism.org/faq (accessed June 8, 2020).

[25] "Frequently Asked Questions".

> **4. We are guided by Jesus' greatest commands** to "Love the Lord your God with all your heart, soul, mind, and strength…and love your neighbor as yourself."
>
> **5. We believe that the intentional use of technology, coupled with following Christ, will empower us to become more human** across the scope of what it means to be creatures in the image of God.
>
> *In this way we are Christian Transhumanists.*[26]

At root, "Christians who embrace transhumanism tend to believe that God is not entirely done with the work of creation but is actively creating even now."[27] *Creatio continua* in Silicon Valley.

Christian Transhumanists are interested in gaining a place at the table with technologists and futurists. This is needed because already "Christianity has lost a propaganda war—no matter what we conclude in the dialogue with transhumanism, we currently do not have the power to create any substantial change" because Christians are primarily external to the conversation, much like the "bioethicist" operates separate from and outside the role of the doctor.[28] Instead, Christian Transhumanists hope for an evangelistic impact of sorts, an increased impact of Christian ethics on the development of transhumanism.

Other Christians are more critical of transhumanism because of its dependence on deficient ideas of enhancement. As Jeffrey Bishop puts it, enhancement technology "is the achievement of a rather dark view of the world. It is the achievement of a sinister metaphysics, originating from relatively recent Western cultural ideas about the ambiguity of the body."[29] Furthermore, "Enhancement technologies and the whole transhumanist lifeworld cannot be merely accepted by

---

[26]"The Christian Transhumanist Affirmation," *Christian Transhumanism Website*, https://www.christiantranshumanism.org/affirmation (accessed June 8, 2020).

[27]Ron Cole-Turner, "Introduction," in *Christian Perspectives on Transhumanism*, 9.

[28]Boaz Goss, "Christianity's Rigged Debate with Transhumanism," in *Christian Perspectives on Transhumanism*, 84.

[29]Jeffrey P. Bishop, "Nietzsche's Power Ontology and Transhumanism: Or Why Christians Cannot Be Transhumanists," in *Christian Perspectives on Transhumanism*, 118.

Christians because at the heart of these transhumanist lifeworlds is a metaphysics and an ontology that is alien to Christianity."[30] Many Christian Transhumanists identify the work of enhancement with the idea of being "co-creators" with God, who continues to create and work beyond the initial chapters of Genesis. But Bishop argues that the "co-creator" language is just a mask for an instrumental, utilitarian calculus that misrepresents the true nature of the world and is ultimately sub-Christian.[31] It sounds theological because it is rooted in Genesis and supposedly subordinated to God's work, but it in fact masks and defends a deficient and non-Christian approach to the world.[32]

Not only does transhumanism (and Christian Transhumanism) depend on a deficient metaphysics and ontology, it also promotes a paradoxical view of human nature. At the same time, "humanity is viewed as a formless work in progress, but also as fundamentally oriented toward desiring specific goods (namely, the goods of control and progress)."[33] Furthermore, there seem to be other paradoxes in play, such as the paradox between the language of artificial "intelligence," which operates based on some level of essentialist definition of "intelligence," but then the completely fluid approach to humanity as evidenced by transhumanism. I introduced this paradox at the start of this article, but hopefully now the substance of the paradox is clearer.

III. FAKE AND FUTURE "HUMANS": THE CHALLENGE TO THEOLOGICAL ANTHROPOLOGY

Now that we have introduced AI and transhumanism, we can explore some ways that these developments will challenge our understanding of what it means to be human, and how that relates to the pursuit of human flourishing in our communities and societies. In other words, we will look not only at traditional "doctrinal" issues but also to the ethical problems that are interwoven with our attempt

[30]Bishop, "Nietzsche's Power Ontology," 119.

[31]Bishop, "Nietzsche's Power Ontology," 131.

[32]Further, instead of buying into the promises of transhumanism, Christians should cling to the doctrines of creation and resurrection. At root, "The Christian message of resurrection is that bodies matter, they have significance, and they are not just clay to be molded to our wills." See Bishop, "Nietzsche's Power Ontology," 133.

[33]Ysabel Johnson, "Rivalry, Control, and Transhumanist Desire," in *Christian Perspectives on Transhumanism*, 230.

to follow Christ in the face of these particular opportunities and challenges.

1. *Expansion: I define myself*. Transhumanism subtly tempts us to believe that our humanity is infinitely malleable by playing on our hopes for technology-empowered improvement. While we might think that we do not buy into this, we cannot deny that this ethos surrounds us and impacts the way we think about the world. As ethicist Jason Thacker puts it, "Because technology is woven into every aspect of our lives, it will naturally revolutionize how we see ourselves and those around us."[34] If we are not intentional about countering a transhumanist narrative, we will find ourselves and our churches slowly changed by it.[35] AI and transhumanism are poised to influence any self-definition of humanity we might be prepared to do, intentionally or unintentionally.

2. *Reduction: I am data; I am my work*. We will not only see ourselves expanding what it means to be human and thinking we can define it for ourselves. We will also find that as more and more of the world is turned into data (or, perhaps, recorded as data), we risk reducing ourselves and our neighbors more and more to sets of data.[36] As we find data about human behavior more and more interesting, and more and more useful (see comments on commercial interests below), we should see this as a helpful development that can illuminate for us some of the tendencies and consistencies of those around us. However, we must resist the idea that data can represent a person, full stop. A human person will always exceed what can be recorded as data, because humans are more than simply physical bodies with chemical reactions that can be recorded and stored. In short, the coming years are going to present us with a vast increase in the data we can know about ourselves and others. We are going to be sold on these things as though they reveal who we "really are." This data will be enlightening and could be used for great good. But we must not act like or buy into the idea that it fully represents a person.[37]

---

[34]Jason Thacker, *The Age of AI: Artificial Intelligence and the Future of Humanity* (Grand Rapids: Zondervan, 2020), 44.

[35]For more on this line of argument, see my *Transhumanism and the Image of God: Today's Technology and the Future of Christian Discipleship* (Downers Grove, IL: IVP Academic, 2019).

[36]Thacker, *Age of AI*, 66–67.

[37]The more we accept these ideas and interact with them uncritically, the more like machines we actually become. Some argue that mindless technology use actually turns people into simple machines, programmable and controllable by powerful interests. In *Re-Engineering Humanity*,

And while this difficulty is related primarily to the development of machine learning and AI, it also connects with transhumanism. Advocates of "mind uploading" believe that there may be technological pathways to "upgrade" a person from a biological body to a synthetic one. All you need to do is capture all of the data that make that person that person (which, according to many, is entirely housed within the brain, without remainder). One perplexing issue among transhumanists is the "reduplication problem": there can only be one you, so what happens when you make a downloaded copy?[38] In other words, even if you grant that a human can be reduced to a certain amount of data, and you can copy all of that data out of a biological brain, what do you have when you are done? Two persons? A clone?

The development of AI will also impact our sense of ourselves because it will challenge human beings' sense of work. As erstwhile presidential candidate Andrew Yang argued, "The lack of mobility and growth has created a breeding ground for political hostility and social ills. High rates of unemployment and underemployment are linked to an array of social problems, including substance abuse, domestic violence, child abuse, and depression… This is the most pressing economic and social issue of our time; our economy is evolving in ways that will make it more and more difficult for people with lower levels of education to find jobs and support themselves."[39] As he puts it later, "The challenge we must overcome is that humans need work more than work needs us."[40] These changes will not be isolated to jobs that we can immediately imagine robots doing—say, autonomous trucks replacing truck drivers—but may extend into jobs we had previously considered "safe" because we cannot yet

---

Brett Frischmann and Evan Selinger worry about "techno-social engineering," which "refers to processes where technologies and social forces align and impact how we think, perceive, and act. That's the 'techno' and 'social' components of the term. 'Engineer' is quite close in meaning to 'construct,' 'influence,' 'shape,' 'manipulate,' and 'make,' and we might have selected any of those terms" (4). They argue that we need the freedom to be "off" and freedom from an engineered determinism that many tech companies are after, whether in relation to AI or transhumanism. In other words, our resistance to the idea that we are merely lumps of data can help keep us from patterns of life that do in fact reduce us to almost that. See Frischmann and Selinger, *Re-Engineering Humanity* (New York: Cambridge University Press, 2018), 4.

[38]Schneider, *Artificial You*, 84.

[39]Andrew Yang, *The War on Normal People: The Truth about America's Disappearing Jobs and Why Universal Basic Income Is Our Future* (New York: Hachette, 2018), xiv.

[40]Yang, *The War on Normal People*, 68.

imagine an AI doing them.[41] As one scholar puts it, "The threat to jobs is coming far faster than most experts anticipated, and it will not discriminate by the color of one's collar, instead striking the highly trained and poorly educated alike."[42]

Others argue that this line of thinking falls prey to three myths. These myths assume AI will follow a clear line of "progress" away from human involvement, eventually replacing all human jobs, and leading to a fully autonomous intelligence that can operate on its own. Instead, others believe there will be more creative ways of interacting with and utilizing AI, maintaining human control, jobs, and so on. The future, for these thinkers, is collaboration, not replacement.[43] In fact, "For the vast majority of professions, the new machine will actually enhance and protect employment. We don't think, for example, that a single teacher or nurse will lose their job due to artificial intelligence. Instead, these professions will become more productive, more effective… and more enjoyable. Workers in such professions will come to view the new machine as their trusted colleague."[44] Such collaboration will raise a different set of questions for the meaning of human work, and we must be better prepared not to reduce our sense of humanity or our primary identities to our work.

3. *Big business: aligning commercial interests and the common good*. Another economic challenge presented by these developments emerges when we look beyond the impact on jobs to the way economic incentives drive the growth and implementation of AI and the implications these decisions have for society at large. We also must consider the impact that AI will have on the development of human economies and societies. In *The Big Nine*, Amy Webb draws out how nine major corporations have a large impact on the direction

[41]See, for instance, the work of West, *The Future of Work*.

[42]Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (New York: Houghton Mifflin Harcourt, 2018), 5.

[43]See, for instance, David Mindell, *Our Robots, Ourselves: Robotics and the Myths of Autonomy* (New York: Viking, 2015), 8–9.

[44]Malcolm Frank, Paul Roehrig, and Ben Pring, *What to Do When Machines Do Everything: How to Get Ahead in a World of AI, Algorithms, Bots, and Big Data* (Hoboken, NJ: Wiley, 2017), 8–9. See also Paul R. Daugherty and H. James Wilson, *Human + Machine: Reimagining Work in the Age of AI* (Boston: Harvard Business Review Press, 2018); Andrew McAfee and Erik Brynjolfsson, *Machine, Platform, Crowd: Harnessing Our Digital Future* (New York: Norton, 2017); Thomas H. Davenport and Julia Kirby, *Only Humans Need Apply: Winners & Losers in the Age of Smart Machines* (New York: Harper, 2016); Nick Polson and James Scott, *AIQ: How People and Machines Are Smarter Together* (New York: St. Martin's, 2018).

of this field, and there are a variety of ways that it could turn out.[45] Christians must consider these elements not only to hope for the ideal direction, but also to consider how best to minister to people in the midst of some of the less-optimistic future scenarios. Webb's basic argument is that the development of AI is currently controlled by nine main companies that could take it in three different directions depending on a variety of factors. She wants especially Western countries to invest more in AI so that it does not have to simply be about being quick to market and therefore making a profit for investors and shareholders.

As Webb lays out the nine companies, they fall into two main groups or tribes. G-MAFIA is the Western group (Google, Microsoft, Amazon, Facebook, IBM, and Apple), and it is primarily dependent upon the profit motive. They are well intended, but they have to focus on products that are quick to market and fit the consumeristic desires that would make them attractive. Meanwhile, the coding, etc., that is going on right now will be incredibly important for the way AI continues to develop. Webb hopes that Western countries can help the G-MAFIA collaborate and be motivated and guided by the common good, not just profit.

The other group, BAT (Baidu, Alibaba, Tencent) are the Chinese companies controlled by that government. According to Webb, China is considering the long-term in a way the West is not. But their long-term goals are bent on world domination. They have more data to build on, etc., so they are ahead in many ways.

Webb's three futures are interesting and well developed. There is an ideal scenario, in which we learn to collaborate and align the development to a common good future. There is a pragmatist scenario, in which Webb describes many "paper cuts" that lead to an adequate but still difficult future. The worst-case scenario is one in which China comes to dominate and ultimately eliminate the West. While only time will tell the outcome, this angle should encourage Christians to consider how to align technology with neighbor love, not only on the individual level, but also in how we hope and work to see technology deployed in our societies. The common good must be a human good and one rooted in a true sense of human flourishing.

[45] See Amy Webb, *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity* (New York: Public Affairs, 2019).

4. *Surveillance and privacy; policing and justice*. Another way that the technology of algorithms becomes more problematic in societies is when combined with machine learning. As noted above, machine learning takes a known data set and then teaches itself how to create an algorithm that can work with future data points for accurate predictions. So, for instance, you could "give the computer" a dataset on criminal statistics that pull in all sorts of factors, including verdicts. Once it teaches itself by interpreting patterns, you can plug in other data, let it work, and it'll give you results that fit the pattern of the original data set. Such systems are used in policing (to determine which areas of a city to patrol more carefully) and in sentencing (to determine how likely a particular person is to re-offend). The problem is that no one knows how it works. For instance, an algorithm built via machine learning for criminal justice could be racist, relying overtly on race or racial signifiers in sentencing. If no one knows how it works because it is too complex, there is no way to evaluate the ethics of the way it is making decisions.

One of the most profound questions for AI, I think, is how to make machine learning ethical, if we can. One of Hannah Fry's most helpful ideas is the notion of "algorithmic regulation": "Should we insist on only accepting algorithms that we can understand or look inside, knowing that taking them out of the hands of their proprietors might mean they're less effective (and crime rates rise)?... In part, this comes down to deciding, as a society, what we think success looks like. What is our priority? Is it keeping crime as low as possible? Or preserving the freedom of the innocent above all else? How much one would you sacrifice for the sake of the other?"[46] Or would such regulation grind development and profit to a halt?

These issues weigh heavily in the actual pursuit and prosecution of justice, but they also impact our overall understanding and expectations of privacy. In her book *The Age of Surveillance Capitalism*, Shoshana Zuboff reveals how companies are built on collecting, analyzing, selling, and utilizing data.[47] This issue of surveillance ties in with the issues related to how AI can be turned more toward the

[46] Fry, *Hello World*, 173.

[47] Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (New York: Public Affairs, 2019). See my review, "On Being Watched, and Remembered," *Front Porch Republic*, May 15, 2019, https://www.frontporchrepublic.com/2019/05/on-being-watched-and-remembered/.

common good rather than merely short-term financial interests. Easy answers aren't options here, but we must be ready to consider how viewing humans primarily as data, building companies to turn that data into profit, and the ubiquitous surveillance it relies on impacts our understanding of what it means to be human.

5. *Warfare and world domination*. Of course, ad companies using surveillance and AI might end up being the least of our worries. As Vladimir Putin said in 2017, the country that takes the lead in AI will rule the world.[48] But why?

To understand the current developments in this, we need to rewind to the ways AI has developed. Kai-Fu Lee explains the new era in AI by tying it to AI's history and then putting all of that in a political context. Basically, there were two camps: rule-based approaches (which sought to program algorithms) and neural-network approaches (machine learning and ultimately deep learning). AI research has gone through "winters," when development is slow. Deep learning is narrow AI, which draws from data in one field for use to achieve a specific outcome. Basically, in the mid-2000s neural networks research made a leap forward and then proved better in competition in 2012.[49] This leap puts us into the age of implementation.

Neural networks need three things: data, computer power, and the work of strong engineers.[50] Computing power and engineers are easier to get. What is going to make the difference going forward is access to data. China is way ahead on this front because their Internet has developed differently and has gobbled up so much more data on so many more people. All of this data can be fed into innovative algorithms for implementation. We have shifted from the discovery phase (figuring out how it works) to the implementation phase (applying it in a variety of ways); from the age of expertise (when we need experts to develop the theory) to the age of data (the neural networks work; they just need more data). While the West had advantages in the early stages of development, now China has the clear edge.[51]

But how might this tie into not only economic advantage but

[48]Radina Gigova, "Who Vladimir Putin Thinks Will Rule the World," https://www.cnn.com/2017/09/01/world/putin-artificial-intelligence-will-rule-world/index.html (accessed June 10, 2020).

[49]Lee, *AI Superpowers*, 9.

[50]Lee, *AI Superpowers*, 14.

[51]Lee, *AI Superpowers*, 15.

to ruling the world? Paul Scharre served in the military and has been involved with policymaking regarding autonomous weapons. His book wrestles with "lethal autonomy" and how nations should approach that, given that AI is getting faster and faster but warfare requires an understanding of context that seems to require a human "in the loop." He is not against using AI, but he warns against a rush to autonomous robot killers. These are questions we must face now, because the technology is already available to make many of these things happen. What policy can limit this on an international scale? Over ninety nations have drones patrolling the skies, and more than thirty already have defensive supervised autonomous weapons.[52] The Israeli Harpy drone has already crossed the line to full autonomy: it "can search a wide area for enemy radars and, once it finds one, destroy it without asking permission. It's been sold to a handful of countries and China has reverse engineered its own variant."[53] As Scharre puts it, "AI is emerging as a powerful technology. Used the right way, intelligent machines could save lives by making war more precise and humane. Used the wrong way, autonomous weapons could lead to more killing and even greater civilian casualties."[54] We should not underestimate the significance AI will play in future global conflicts and balances of power.

6. *The limitations*. We should certainly be wary of the many ways that technology could go wrong. At the same time, we should be wary of too much hype. A robust doctrine of humanity reminds us that humans are the crown of God's creation. This does indeed mean we can do great things, but we should not expect our own creations to do everything. One realm to consider, even from a secular perspective, comes down to meaning and value. As Scharre explains, "Machines can do many things, but they cannot create meaning. They cannot answer these questions for us. Machines cannot tell us what we value, what choices we should make. The world we are creating is one that will have intelligent machines in it, but it is not for them. It is a world for us."[55] We might want to situate that sentiment a little more theologically, but at its root we are reminded of the limitations

[52]Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: Norton, 2018), 4.

[53]Scharre, *Army of None*, 5.

[54]Scharre, *Army of None*, 8.

[55]Scharre, *Army of None*, 362.

of our technology and our responsibility to orient not only the tools but the culture around the tools in a way that honors the kingdom of God rather than building the idolatrous kingdom of man.

IV. CONCLUDING WITH PERSONHOOD

Where do we go from here? There are many possible routes to address AI, transhumanism, and the challenges and opportunities they raise from a Christian perspective. We could talk about the *imago Dei* in Genesis, the prohibition of idolatry throughout the Bible, the prophetic call for justice, Jesus's teachings on caring for the marginalized, or the Great Commission's charge to make disciples. In *2084*, John Lennox turns to the book of Revelation for insight.

But what about considering personhood, seeking a better understanding of how we can know a person when we "see" one? This idea can help us notice the difference between humans and artificial intelligences, as well as the false promises of transhumanism. While we're used to the language of personhood in a theological context, its use in a secular context is already growing in significance in relation to these challenges. Susan Schneider asks the question "What is a person?" in her book *Artificial You: AI and the Future of Your Mind*. She goes on to highlight four main theories, before roughly combining two of them to argue for ways that personal existence could persist outside the physical brain. Going into her argument would take us too far afield at this point,[56] but this shows that the question of personal existence is tied into these questions of what exactly a person is and how that relates to the material world and the "digitizable" world. Here we are, back at the doctrine of humanity.

Secular approaches to AI and transhumanism have to make a call on what it means to be a person, because they must explain whether AIs should be considered persons, and they must also explain how some of these radical extensions of "life" would still be the same "person." But they actually lack the ability to provide a solid definition. They lack this because they refuse to allow God to speak, and they also lack it because they are pulled in opposite directions. Techno-utopians insist on essential definitions of things like "intelligence," but they resist any essential definition of "human" or "person," because the whole transhumanist project is built on

[56]See Schneider, *Artifical You*, 74–81.

exceeding and improving everything, which resists the idea of preserving any "essence."

As Christians, we must develop a strong doctrine of humanity not only to guide our use of particular technologies for ourselves (the temptations associated with transhumanism) but also in how we consider, evaluate, and "treat" emerging technologies (AI).

One article cannot provide a robust enough treatment of the doctrine of humanity, nor can a single issue of a journal. But we can start, and we can point in directions of further development. I would like to propose one quick litmus test question for evaluating whether or not something is or is not a person. Can it make or break a covenant with God? To be a person is to be one who can enter covenant with God. Or, perhaps we might say, to be a person means to be able to exist in obedience or disobedience to the Triune God. (Angelic persons, then, fit into this, without our having to determine some sort of "angelic covenant.") As Michael Horton puts it, "Can there be any doubt that human beings are uniquely suited among the creation to be covenant partners with God?"[57] If we develop our understanding of the image of God into a series of capabilities, we might very well see that AI can replicate many of them. Some sort of transhumanist intelligence built off a copy of a biological brain might also be able to replicate some. But does that make either of those *things* into *persons*? I do not think so, because personhood is ultimately given by God, the Creator, to those he calls into relationship with himself for his glory. We can only acknowledge that we have received this gift; we cannot create it ourselves.

We could also recast this litmus test with the question the Gospels writers put before us, reminding us that Jesus asked, "But who do you say that I am?" While an AI might be able to answer with facts, or even repeat statements that sound like praise, only a person can give and live by Thomas's later exclamation: "My Lord and My God."

[57]Michael Horton, "Image and Office: Human Personhood and the Covenant," in *Personal Identity in Theological Perspective*, ed. Richard Lints, Michael S. Horton, Mark R. Talbot (Grand Rapids: Eerdmans, 2006), 184.