## **Article**

# Artificial Intelligence: A Theological Perspective

Albert Erisman and Tripp Parker

Artificial intelligence is making inroads into many areas of our lives, and it will continue to do so. These changes will be substantial and disruptive. Surprising to some, the Bible has a great deal of wisdom to help us in this transition. First, we find that the development of such systems is a vital part of the creation mandate, and this work can be done to the glory of God. Second, because of the reality of the brokenness of our world, and ourselves, there are important cautions we should consider in the development and the deployment of these systems. In this article, we show why AI system development is different from ordinary programming, how it can be a part of the creative process God has called us to, and where we need to provide biblical cautions for developing and for using these systems.



Albert Erisman



**Tripp Parker** 

This time it looks real.<sup>\*</sup> The promises of Artificial Intelligence (AI) were first articulated over fifty years ago, but the excitement and hype died rather suddenly, in part from the book *Perceptrons* by Marvin Minsky and Seymore Papert.<sup>1</sup> They identified fundamental limits of the AI developed at that time.<sup>2</sup> That led to a pessimism about the exaggerated promises of AI and a significant decrease in research funding, ushering in the AI winter of the 1970s.

A second round started in the 1980s. Technology had become more powerful and the PC had come on the scene, distributing computing power to the masses. Expert systems (ES) were being touted as the replacement for many human decision-making challenges, including replacing much of what doctors (or pilots) did. Surely this was the time for AI systems to make a difference. In reality, complex decision making was much more challenging than AI enthusiasts had believed. Much of the work moved from expert systems to expert assistants – parts of the problem could be handled by the ES, but final judgment rested with a person. This was useful sometimes, but a long way from the promise.

In the 1990s, virtual reality (VR) became a focus – the creation of a virtual world in which humans could experience a different reality than they would be able to do in real life. Much of this was reserved for games. Albert Erisman remembers racing down a slalom course on a VR system, competing for time on ski slopes he would never attempt in real life. The vibrations in the skis and the visual cues were amazing and fun. The lack of pain from a crash was even better. In reality, this was far from reality.

At Boeing, Erisman's R&D team began to look at this technology for business use

**Tripp Parker** (*BSE*, *Duke University*) studied computer engineering, computer science, and philosophy of mind at Duke University. He currently works in R&D at Amazon in Seattle, WA, in the Alexa Health and Wellness division. His interests include the ethics of AI, the metaphysics of consciousness, and using AI to help humans flourish.

<sup>\*</sup> This article is a revised and expanded version of Al Erisman, "Artificial Intelligence," *ethix*, February 28, 2018, https://ethix.org/2018/02/28 /artificial-intelligence.

**Albert M. Erisman** (*PhD*, *Iowa State University*) is Executive in Residence emeritus at the School of Business, Government, and Economics (SBGE) at Seattle Pacific University. He is also the executive editor of Ethix magazine in which he has interviewed more than one hundred world leaders on the topics of business, ethics, and technology. In April 2001, Al completed a 32-year career at The Boeing Company, the last ten years as Director of *R&D* for computing and mathematics.

in their lab. Bob Abarbanel led a team in developing FlyThru, a VR system that allowed engineers, managers, and potential customers to "fly through" an assembly of electronic parts as if it were a real airplane. This became a key tool in the design of the 777 airplane. David Mizell headed the team to allow those in the factory to try assembly procedures in the virtual world. Tom Caudell had the idea of merging the virtual and real worlds together, projecting instructions for a repair procedure onto the physical part of the airplane, giving the mechanic hands-free access to information. He coined a new term for this, which he called "augmented reality," in 1990. It is interesting to see the current hype about virtual and augmented reality, as if they were something new. It provides an example of a popular phrase often used by technologists: "The future is already here, it is just unevenly distributed."

When IBM's Deep Blue defeated the then reigning world champion in chess, Garry Kasparov, in 1997, new hype began: "These systems are going to rule the world." Chess had been seen as the ultimate challenge demonstrating that any activity of the human brain is fair game, according to the promises of the 1960s. This thinking simply demonstrated that these researchers did not understand the human brain. In May 2017, a computer defeated Go champion Ke Jie. Since Go is considered the most complicated board game, the promise of AI seemed even more real.

Today, technology seems to have arrived at a point at which these systems will have a greater and greater impact on all of us. They will invade our lives, our workplaces, and society in ways that will produce much more substantial change than all that has happened in the past fifty years. Further, many of these systems will be invisible, not limited to a computer sitting on a desk, as Neil Gershenfeld predicted twenty years ago.3 These systems will make our lives safer and better; they will make products less expensive and better; and their promise is real enough to lead to substantial investment in the companies that build them. Yet, at the same time, there are significant questions about such systems that should engage us-not with an emotional resistance nor with fear, but with careful thought at the levels of design, personal use, organizational use, and societal policies and impact. To engage thoughtfully requires that we understand enough about these systems to inform our responses to them.

This suggests that we should give careful attention to two questions. The first is, what can go wrong with such systems? Albert Einstein once said, "We cannot solve problems by the same kind of thinking we used when we created them." The second question is, how might such systems impact society?<sup>4</sup> But first we want to briefly describe how AI systems are different from traditional computer programs, because this understanding informs our response to the two questions.

# How Does AI Differ from "Normal" Programs?

AI systems differ from standard computer programs in an important way. A typical computer program follows an algorithm, a step-by-step procedure that starts with certain data and instructions and ends with a result in a repeatable, reliable way. A recipe for a cake follows this pattern. Given a set of ingredients, combine them in this way, cook them at this temperature for this period of time, and at the end we have our cake. This is also precisely what an accounting program does. Given this data, produce a cash flow or profit-and-loss statement. The human did the thinking, laid out the steps, and the computer carried out the calculations, producing the results.

AI systems work differently. A human may not understand the process, but the person feeds the system some rules of the game, some examples of good output derived from input, and the computer system (or supervised learning system, in this case) figures out how (through spotting statistical patterns in the data) to produce good output from given input. To emphasize, the person behind the system did not specify what those patterns were — indeed, he or she may not even understand what they are — but the learning system figures out a way to produce a result from the input. In a sense, this is how a child learns. Lots of trial and error, many false starts, some correction, and then she or he learns.

Here are three examples. It used to be that computer-based language translation was based on the programmer providing the instructions for translating a document based on vocabulary, rules of grammar, and so forth. The results were poor. Computer-based translations were barely readable and, at best, were aids to a human translator. More recently, work on computer-based translation has followed a different course. The learning system is provided with documents in one language and examples of good human translation, and then the system determines the procedure to change from one language to another. Such systems have made a significant improvement in language translation.<sup>5</sup>

A second, simpler example is about teaching a learning system to do what many children can do. How do you tell the difference between a wolf and a dog? The distinctions are challenging to describe in some sort of step-by-step procedure, though many children can tell the difference. In one famous example, researchers provided a series of pictures to the learning system (a neural network, in this case), properly identifying dogs and wolves in the sample learning environment.<sup>6</sup> Once the system had sufficient data, they subsequently fed a variety of new pictures to the system, which it began correctly labelling as either a dog or a wolf, demonstrating (it seemed) that it had learned to distinguish between them.

A third, harder example is self-driving cars. It would be impossible to lay out a step-by-step procedure for all of the decisions a person must make driving across town. But it would be enough to feed a variety of rules to the car's AI system and let it learn how to drive, much like a teenager learns to drive.7 Speeding is bad, going too slow is bad. Crashing into cars or pedestrians is bad. Anticipating and avoiding accidents involving other vehicles is good. Understanding the shortest way to get to the destination and following that path is good. With experience and testing, the car learns to drive, to navigate through traffic, to avoid accidents, to take the best route. The advantage of a car learning to drive is that the result can then be downloaded to other cars. As experience grows, cars can share their advanced learning with other cars. The result is safer, more reliable driving.

Auto accidents killed about 40,000 people on the highways of the US in 2016. For the first time in history, distracted or drowsy drivers killed more people than did drunk drivers. It goes without saying that computers do not get distracted or drowsy. Selfdriving cars, while new and frightening to many, may already be better drivers than humans, though the full complexity of this has not been proven. And they will continue to get better. Because self-driving cars are new, humans do a poor job of comparing risks. This explains why a single accident by a selfdriving car in California rates headlines around the nation, whereas thousands of other more serious accidents happen every day involving cars with drivers.

### Virtuous AI

The basic question is this: in a scenario in which any decision incurs a cost, how do we make sure the *least bad* decision is made? Can we even agree on what the least bad decision is? How do we build a virtuous AI that could reliably make such a decision?

In discussing this question, let's dispense for the moment with the prospect of an all-knowing and conscious AI. Let's ignore the science fiction version of AI in which it "wakes up" and does things we do not want it to do. Instead, let's focus on a simpler question: why is it difficult to build an AI that reliably does what we want? We call this the "AI alignment problem": How can we build an AI that acts in a way that is aligned with our values? What data might we give the AI that would teach it about our values and what we care about? A Christian might ask, "Can we just feed it the Bible as an input, and have it figure out what to do in a way that is wise and just?"

There are both theological and technical challenges with doing this, but let's start with the theological. This may seem obvious, but not even Christians can agree on what the Bible means when we read it. Even when we agree on the textual interpretation, we are not in full agreement on how to apply it in our everyday lives. How on Earth can we be confident to give scripture to an AI and have it reliably act in a way that is commensurate with our values? We do not agree on our values. We often misunderstand scripture. One can give ten people an algorithm for how much Tylenol® to take, and all ten people can interpret it correctly. If we give ten people the Bible, we will get ten different interpretations. In other words, the Bible is not some kind of holy algorithm for every answer to every problem we face in the modern world.

As Solomon wrote, "There is nothing new under the sun."<sup>8</sup> And like Solomon, we need to be wise in applying the deeper lessons of our faith, but wisdom comes from the Lord, through a relationship with him, not with an algorithm. A great description

of wisdom is something like this: Wisdom is not a rulebook, it is more like a dance.<sup>9</sup> We all have values that sometimes conflict with one another. Wisdom is knowing which value, in any given instance, should take the lead, and which should follow. The Bible is a book full of values, along with examples of wisdom and folly. Solomon, despite knowing what the scriptures said regarding right and wrong, despite being well versed in the law, asked God for wisdom. That did not come in the form of a rulebook. That did not come in the form of an algorithm. That came through a personal relationship.

Therefore, it is critical that Christians, people who have a personal relationship with the Lord of the universe, be involved in building, using, and guiding the future of AI. It will take wisdom, and wisdom cannot be easily prescribed.

Further, even if we completely agreed on what a wise decision looks like in any given scenario, there are technical problems with building such a system. Let's return to the example of telling the difference between wolves and dogs. The AI system seemed to have learned the difference, and was accurate based on the pictures provided. After more pictures were inputted, the researchers noticed that the system was giving a number of wrong answers. Why was it mixing up dogs and wolves? The decision criteria, the patterns between the input data and the correct answer, had not been prescribed by the researchers but had been developed by the learning system itself.

Eventually the researchers figured out, through many tests, that the system was not actually paying attention to the animal. It was looking at the animal's environment. If it saw snow in the picture, it declared that the animal must be a wolf, because the preponderance of the pictures fed to it had wolves standing in snow, whereas most of the dogs had been photographed on grass.

In other words, the researchers did not actually know what the system had "learned." AI solves problems in very different ways than we do, detecting patterns that would not occur to a human. When presented with a new, unexpected situation, AI's response can be unpredictable.

A metaphor that Tripp Parker often uses regarding this problem comes from the Disney movie Fantasia. Specifically, in a segment named the "Sorcerer's Apprentice,"<sup>10</sup> Mickey is an apprentice to a powerful sorcerer who gives him a task to do late at night: fill a cauldron with water. The Sorcerer retires for the evening, leaving his magic hat downstairs where Mickey is supposed to fill the cauldron.

Mickey, wanting to complete the chore with as little effort as possible, uses the magic hat to animate a broom to fill the cauldron for him. The goal that Mickey gives the broom is a completely full cauldron. The broom picks up a bucket, fills it with water, and carries it over to the cauldron. Mickey watches, and the system appears to be working. Mickey goes to sleep, leaving the broom to finish the task assigned to it.

You may remember what happens next: the broom overfills the cauldron, flooding the workshop. You can think of the broom as an AI that is trying to maximize the chance that it successfully fulfills the task given to it. What if there's a leak in the cauldron? What if someone took water out when it was not



Figure 1. Wolf!

Photo credit: Okssi/Shutterstock.com



Figure 2. Dogs!

Photo Credit: Alan Jeffery/shutterstock.com

looking? What if the broom does not have accurate vision, and while the cauldron appears full, it really is not? The way to maximize the chance that the cauldron is full is obvious to the broom: continuously pour water into the cauldron.

In other words, the broom's actual values did not fully align with Mickey's. Mickey told the broom about only *one* of his values, not all of them. He cares about a full cauldron, but he also cares about a flooded workshop. However, as we have discussed, the broom did not learn that, and therefore created a solution that was worse than the problem Mickey wanted solved. Such a solution would not occur to a human, who intuitively knows this to be a bad solution. But to an AI, it may make complete sense. Such knowledge is often referred to as *tacit* assumptions. Although they are obvious to all in the context of life, it is extremely difficult (or impossible) to document all of our tacit assumptions.

It is a difficult *technical* problem to ensure that the system you build reliably aligns with your values and learns the right output. There will be times when the system learns the wrong output, or it may lack maturity in its learning (not understanding the full range of values it should care about or pay attention to). In the case of human drivers, we have come to peace with making these judgments. We still license teenage drivers, even knowing that the risks of accidents are higher and that the maturity of judgment may be less. Many people are less comfortable with the prospect of an AI that lacks full maturity.<sup>11</sup>

And yet self-driving cars are the wave of the (near) future. Computers now defeat the best of the world's chess or Go champions. Smart algorithms sort through thousands of pages of material in minutes that once took weeks for well-trained (and well-paid) lawyers. Facial recognition software can identify a particular person running through an airport. Computers are changing the face of the factory with robotics and the medical research world with testing procedures. Data analytics frequently drive the decisions made, from the board room to the sports team.

Many excitedly embrace this new world with no thought of what could go wrong. Others fear a future without jobs, without a sense of control or understanding, and with scarcely a thought of the benefits. What do Christians have to offer to this conversation? What is it about this new technology that should offer hope and excitement, and what about it should give us pause? How do we navigate this new world we are entering?

## **Biblical Insight**

Starting with the scripture, it is not difficult to see where creative thoughts originate. Since people are made in the image of a Creator God, we see the roots of the passion and joy that come from the act of creating new things. The first two chapters of Genesis show God as a creator, God making humankind in his own image, and God bringing humankind into his work. Specific instructions include

- Oversight responsibility (Gen. 1:28-30)
- Care for the creation (Gen. 2:15)
- Classification responsibility (Gen. 2:19-20)

God stated that the creation was not complete, in that "there was no one to work the ground" (Gen. 2:5, New International Version [NIV]).<sup>12</sup>

God's purpose for humans in the design and discovery process is referred to beyond the creation account. In Proverbs, we are reminded of the delight in discovering the hidden things in God's creation:

It is the glory of God to conceal things, but the glory of kings is to search things out. (Prov. 25:2, English Standard Version [ESV])

Rosie Perera, a former Microsoft developer, put it this way:

As a software engineer, I have had the experience of creating something out of virtually nothing, which is pretty amazing. I love well-crafted, elegant computer code. And I love to see people's faces light up when they learn how to do something on the computer that previously mystified them.<sup>13</sup>

This sense of excitement is rooted in how we were made. It is developed by God and blessed by God. AI systems offer a special place within this creation process. The designer identifies key steps of the design, expecting the system to fill in the rest. These resulting systems can bring true and new insight.

If this were the end of the story, we could all share in this joy. But there is more to the story. In Genesis chapter three, we see the broad impact of sin in our world. In addition to the separation between God

and humankind that came from the Fall, there was a separation between people, and between the people and their work. This impact on work offers insight on how we respond to the changes in our world, including the changes caused by technology. God said,

Cursed is the ground because of you; through painful toil you will eat food from it all the days of your life. It will produce thorns and thistles for you, and you will eat the plants of the field. (Gen. 3:17b-18, NIV)

Thorns and thistles that grew in the crops were not part of the original plan. They crept into the farming work as impediments to the real task of growing food, sometimes choking out the intended growth. They interfered with the best of plans.

AI systems have their own unique "thorns and thistles." The conclusions that they can draw from incomplete information can be both dazzling and dangerous. The wisdom that is required to assess both what such systems do and how they do them, calls for creative insights and also creative awareness of how these systems can go off track.

Thorns and thistles have an application in the development and use of technology as well, in at least these four ways.

- 1. Our motives, both as developers and users, are not always pure. Both designers and users may approach technology with nefarious intent. These thorns and thistles can turn good work and good technology in a decidedly bad direction. For example, a talented designer might use his or her abilities to create phishing schemes that harm others. Or a user may employ a powerful system to sync together words and facial expressions in order to create a "deep fake" video of a person saying that which they did not say.<sup>14</sup>
- 2. Bugs, design flaws, and unanticipated results can show up in our work, producing unexpected results. For example, an AI system used in teacher performance ratings assumes that standardized test score performance is a valid measure for evaluating the performance of a teacher.<sup>15</sup> Sadly, some human performance evaluation is also carried out mechanically without the benefit of human wisdom.
- 3. The work product may be done well, meeting

all specifications, but it may have a surprising application that was not intended. For example, an automobile designed to provide safe, reliable transportation is used as a getaway vehicle in a bank robbery.

4. A great strength of AI systems is that they add insight that humans may not have. The great danger in such systems is that humans trust the system, turning off their own wisdom rather than applying it in new and unique ways. Some fear such systems, wondering if they are not more powerful than humans. Yet perhaps this question arises only because these systems are relatively new. Machines have always been more powerful than humans when muscle is needed. Computers have always been more powerful than humans in carrying out a long string of computation. A simple illustration here is the ability of computer cash registers to compute the proper amount of change in a transaction. Without human wisdom, the amount of change may be very wrong because a data entry error is made, but if no human estimate is made, the answer is believed. The challenge is to find the good and right role for such systems, and not to assume that they have insight and moral judgment as has been given to humans by God.

Interestingly, in spite of the thorns and thistles, the sense of joy and satisfaction that comes from good work remains a part of who we are. The reality of the brokenness of our world should not cause us to lose hope. Christ came to bring hope and healing to the brokenness in our world, and while this will not be complete until his return, we can be agents of reconciliation now. Paul said,

Therefore, if anyone is in Christ, the new creation has come: The old has gone, the new is here! All this is from God, who reconciled us to himself through Christ and gave us the ministry of reconciliation: that God was reconciling the world to himself in Christ, not counting people's sins against them. And he has committed to us the message of reconciliation. (2 Cor. 5:17–19, NIV)

He says this in a different way in another place:

The night is far gone; the day is at hand. So then let us cast off the works of darkness and put on the armor of light. Let us walk properly as in the daytime ... (Rom. 13:12–13, ESV) And Jesus said,

You are the salt of the earth ... You are the light of the world ... In the same way, let your light shine before others, that they may see your good deeds and glorify your Father in heaven. (Matt. 5:13–16, NIV)

Together, these passages remind us to bring the light of the truth of the gospel to bear on everything we do, including the building and use of technology. And while the final, complete healing will come later, we should look forward to that time when all will be set right.

These conclusions help us put AI systems in their proper place. AI can add to our ability to make decisions, but AI is not the autonomous decision maker. AI can bring insight to a problem, but not without supervision. As Edward Tenner wisely put it,

Pessimism about the effects of technology is a distraction from the real need for education and self-education on the best way to combine algorithms and intuition, digital and analog.<sup>16</sup>

Christians engaged in AI, either as builders or users, too often regard this work in a separate category from their faith. They are surprised by the bugs, misuse, and unanticipated consequences. They should not be. Sometimes they assume that this is just the difficult world in which we operate, and that somehow the light of the gospel has no connection to their work. In fact, biblical insight on why we can love our work, and why it can yet go so wrong, is valuable for everyone. We should raise questions others might not raise in the context of our work. As a result, people may even ask questions about the gospel!

#### What Should We Watch Out For?

Even if AI is built *technically* in the right way, there are six main reasons why AI could affect us negatively:

- 1. Destabilization
- 2. Idolatry
- 3. Corruption
- 4. Unanticipated consequences
- 5. Contextual misfit
- 6. Isolation

Let us look at each one briefly.

#### 1. Destabilization

It is no secret that robotics, for instance, has taken many manufacturing jobs away from those who have historically performed them in the US. AI will do the same, but to a larger degree, and will do so much faster than the technology we have created in the past.

Today, three million people earn their living from driving buses, cars, trucks, and other vehicles. If the switch to driverless vehicles happens quickly, as many are predicting, this will be a huge disruption in the labor force. Add to this the medical jobs that involve reading X-rays, supporting diagnoses in general, and testing pharmaceutical drugs, and we see a significant number of jobs that are vulnerable. Sorting legal documents, working though accounting categories, and other diverse jobs are at risk. In fact, *any* job that is repetitive and predictable is at risk of rapid automation.

It is easy to argue that we have been through this before. The Industrial Revolution is one example, as is the larger migration from farm to city jobs. Many of these, however, had longer implementation times, enabling the retraining of workers for other semiskilled positions. Further, the new positions were similar to the old ones. It is one thing to train someone who had repaired old tools to work in a factory that builds new ones. It is quite another problem to train a tractor-trailer driver to be a physical therapist.

This time, changes will probably happen much more quickly, and the retraining may involve much more complex skills that take longer to learn and do not match everyone's abilities. On the other hand, there are many jobs that need to be filled but do not pay very well—service jobs that support an aging population is but one example.

Society and the church will need to wrestle with how to best address these issues. How will we care for those affected? How do we best help them and their children for the long term? If we are subcreators, made in the image of God, how do we help people find their place in this new world without robbing them of the ability to contribute to it?

#### 2. Idolatry

AI can be a temptation toward idolatry. One could be forgiven if it were suggested that we already wor-

ship our technology. How many times have you seen a family in a restaurant, sitting together but paying attention solely to their smart phones? Is this not a form of worship? How many of us would immediately return home if we realized we had forgotten our phone, whether or not it was *really* needed? Could we not consider this an idol?

AI will make the problem worse. Movies such as *Her* and *Ex Machina* play with this idea, that as the technology starts to better imitate a person, as it caters more and more to our every whim, we may see it less as a tool to use for God's designs and more as an idol that gives us what we want. Why bother with real relationships, when an artificial one will give me what I want (but not what I need) without the messiness of involving another sinful person? Why be present in the real world with all its messiness, when we can interact with an artificial one that's much cleaner and more to our liking? Luke Dormehl develops this case.<sup>17</sup>

Further, it is easy to see how we might use AI to make gods of ourselves. You can see this in, for instance, election advertising. Using one type of AI system, people are classified and put into categories. Another classifies the content needed to target them, and yet another targets people to get them to act in the way the creators wanted. Maybe they can influence you to change your vote? Maybe you were planning to vote, but they make you a little less likely to do so? By selectively providing people with information (regardless of whether the information is *true*), over large populations one might be able to swing a few thousand votes in Michigan, or Virginia, or Georgia. And a few thousand votes in the right place can swing an election.

In other words, these tools allow their creators to influence people *en masse*. I do not need to know you as an individual. I never sit across from you at a table and hear your story. I reduce you to a vote, as clay to be molded so that I can get what I want in the end. I treat myself as God. I commit the original sin. I idolize not the technology, but, through the technology, I idolize myself.

#### 3. Corruption

Here we disagree with a common refrain that you may have heard: "X (*whatever one might be referring* 

to) is just a tool. And any tool can be used for good or for evil."

Despite what these common sayings suggest, tools *are not* neutral. One cannot approach a chair and do just *anything* with it. It begs to be sat on, and sat on in a certain way. Your iPhone cannot be used in just any way, and it does not sit idly by as if it were indifferent to how you use it. Try it for yourself. Set it down next to you. Sooner rather than later, it will light up, whether or not you actually received a call or a text message. An app may give you a notification. A news alert will pop up. It is almost as if it were saying, "*Pay attention to me.*"

All these tools were made with a particular purpose in mind. Your phone and the apps on it have *success metrics*. Their creators have defined how their product will serve its purpose; that is, how each part of the product will encourage you to use it for each of its purposes. Therefore, depending on how a tool is made and the purposes for which it was made, you will find specific incentives to use that technology in certain ways. AI is no different. Here are a few examples of the perverse incentives that AI could create.

As we have discussed, AI needs to predict outcomes based on input data. That is how it is trained. AI systems spot patterns between inputs and outputs, make predictions, and perform tasks in order to produce the desired output. The more data the AI has, and the more diverse the dataset, the better it will be able to give the desired output. Therefore, the creator of an AI system has an incentive to acquire as much data as possible from you in order to better train the AI system to make more-accurate predictions. However, your privacy is an impediment to this goal. Next time you click "Yes" on a software user agreement, just ask yourself, "Why is this agreement so long and complicated?"

In what other contexts might this corruption happen? We have already discussed elections, and the selective and targeted spreading of information. Often, people *want* to be told that they are right. Often, people *want* to hear what they want to believe. AI systems that classify you can give you *what you want, even if it is bad for you.* AI systems can give you the bubble you are looking for, efficiently and without complaint, without using information you might find uncomfortable. We ought also to be concerned about dehumanization. As AI systems behave more and more like people, there is a question about how we ought to treat them. Are they people? Ought we treat them as if they were? Will we?

Immanuel Kant once said, "He who is cruel to animals becomes hard also in his dealings with men."<sup>18</sup> If we can treat with callousness a living being that is not human, often that means we will end up doing the same to people. One might be able to extend this argument to AI systems: if I can treat an AI that *acts* like a person as a tool to be used by me, might I treat other humans that way as well?

As with any human endeavor, we ought to be constantly asking ourselves, "Who is God? Who are we?" and "What does that mean about our current endeavors?" Are we fulfilling our calling, that of subcreators in attempting to redeem the earth and be fruitful? Are we trying to build the Tower of Babel, thinking that we can build heaven on Earth without the blessing of the Creator of the universe? The answer is not simple. However, if we are trying to build and use this technology in a way that is in keeping with our faith, what might that look like?

#### 4. Unanticipated consequences

The more complex the technological development, the more likely we are to encounter unanticipated outcomes. This hearkens back to the earlier example of Mickey in the "Sorcerer's Apprentice"; Edward Tenner develops this case in general.<sup>19</sup> We need to be vigilant and forward looking as we roll out the technologies, but often our culture of short-term thinking and immediate gratification overrides our best intentions.

#### 5. Contextual misfit

We often find that an AI system works well in the lab, but it has difficulty when placed in a bigger context. For example, consider self-driving cars. Driving laws were created for human drivers and focus on the human tendency to create unsafe conditions. Thus, we have laws against speeding (or driving too slowly), failure to obey traffic lights, and so forth. A good AI system can obey all of these laws, but where might the difficulties lie? What laws are needed to create a safe overall system? Insurance is for the driver of the vehicle. Who is liable for an accident with a self-driving car?<sup>20</sup>

In the transition between horses and cars in New York City at the early part of the twentieth century, the most dangerous time was found to be when both horses and cars were on the road together. The transition involved people holding onto their horses because they liked them, governments trying to deal with changing laws, and difficult interactions between two types of transportation. How will the transition to driverless cars be managed?

#### 6. Isolation

There is the important question of how AI interactions may affect the relationships we have with humans. God made us to be in relationship with people, but might our interaction with AI systems be more comfortable and reliable, and undermine our willingness to engage in the hard conversations we should have with others?<sup>21</sup> Sherry Turkle raises this and related issues in her book Reclaiming Conversation.<sup>22</sup> It is possible that reducing the number of human "transactions" (impersonal tasks that we carry out with other people) may cause us to step away from thinking of the other person in terms of the transactions we have with them.23 This could allow us to focus our true human relationships on a smaller number of people (family, friends, neighbors, some coworkers, church members) and to take these relationships more seriously. Perhaps we would see other people more in the way it was intended, rather than simply as persons who can meet our needs. This will require us to be intentional about developing and fostering relationships with others in spite of their messiness.

The other side of this question, discussed by Turkle as well as by Beavers and colleagues,<sup>24</sup> is that our personal relationships need to go deeper than "technology mediation." Texting, emails, video conferencing, and phone calls are helpful, but they are not enough for properly relating to another person. We need to relate at a deeper level, sharing reality well beyond human transactions. Perhaps AIs will free us to do this.

#### **Frequent Responses**

Research is leading to new AI tools and systems that will change our lives. There are several possible responses to these changes.

• The Blind Enthusiast: Some will embrace the changes with little thought to a potential challeng-

ing downside. We need to listen to these people because they suggest new possibilities that we might not have considered.

- The Luddite: Some will push back against all change, resisting almost all new technology. We need to listen to these people as well, because they might remind us of a downside we would not have considered in our own enthusiasm.
- The Disaffected: Some will respond negatively, their objections based solely on how the technological changes will affect them personally. We can easily note potential problems with the technology, but we need to look at the questions from a broader viewpoint. For example, some resist driverless cars because they personally like to drive. Yet with such significant loss of life associated with person-driven vehicles, we cannot afford just one narrow viewpoint.
- The Observer: Some will sit on the sidelines for a long time, waiting to see whether the new technology offers a good or a bad outcome. From these people we can learn not only to avoid a rush to judgment but also to make wise calls.
- The Ambivalent: Some have little interest in technology and simply want to avoid the questions. These people do not need to be experts, but they need to go beyond naysaying and be open to constructive conversation.
- The Wise: Some will immediately try to understand and seek to steer development of AI systems in a way that keeps the big picture in mind. They need to be open to new possibilities, be aware of potential downsides, and be careful to avoid premature judgments. They also need to listen to questions from those who do not understand the technology.

We are delusional if we believe that we can stop this development. The technology is rapidly advancing, and will continue to do so. We were made to do this: to relentlessly create as one made in the image of the Creator. We have an opportunity to be a part of shaping it.

## The Role of Christians in This Discussion

Christians do not have a corner on wisdom in any of these areas. We have found both Christians and

those with no religious beliefs in all categories of wisdom and foolishness. But followers of Christ who take the Bible seriously need to consider some other factors as they engage in this discussion. There are four principles that Christians should adhere to:

First, God has called his people to not make themselves the center of the issue.

Do nothing out of selfish ambition or vain conceit. Rather, in humility value others above yourselves, not looking to your own interests but each of you to the interests of the others. (Phil. 2:3–4, NIV)

Second, we should be people who do more than say "no." David Gill developed this case from Titus.<sup>25</sup>

For the grace of God has appeared that offers salvation to all people. It teaches us to say "No" to ungodliness and worldly passions, and to live self-controlled, upright and godly lives in this present age ..." (Titus 2:11-12, NIV)

We are to say no to ungodliness and worldly passions, but we are to "live ... in this present age." How do we work together to properly discern the times?

AI tools can be a part of healing the sick, producing safer cars, understanding potential implications that lead to new public policy. If we sit on the sidelines, others will shape the future without the insight we can bring to the issues.

Third, in our instantaneous technological society, there is a tendency to look only short term. We may look at short-term gains or short-term losses. We can be caught in excitement or fear. As the people of God, we should broaden our thinking, living as the people of God with the end in mind. Romans 13:12 says we are to live as the people of light even in the present darkness.

Fourth, God's command to his people in exile can as well be meant for us today.

This is what the LORD Almighty, the God of Israel, says to all those I carried into exile from Jerusalem to Babylon: "Build houses and settle down; plant gardens and eat what they produce. Marry and have sons and daughters; find wives for your sons and give your daughters in marriage, so that they too may have sons and daughters. Increase in number there; do not decrease. Also, seek the peace and prosperity of the city to which I have carried you into exile. Pray to the LORD for it, because if it prospers, you too will prosper." (Jer. 29:4–7, NIV)

#### Conclusions

Both the opportunities and the problems of our technological society are real. The effects of AI systems both now and soon to come will challenge our suppositions and draw us into places where we have not been. Even here, we need to live fully for God.

It is no accident that God has placed us in the twentyfirst century. Some would sound a call to retreat, but God commands us to be salt and light in *our* world. This means we do not hide from the changes, or simply embrace them as inevitable, but we seek to understand them from the light of the scripture. Like the body of Christ that Paul talks about in 1 Corinthians 12, we do not all have the same role, but different roles. Let us challenge and encourage each other in this world where God has placed us.  $\triangle$ 

#### Acknowledgments

The authors are grateful to the referees who reviewed an earlier draft of this article. Their helpful comments encouraged us to add references and to make changes to the original draft.

#### Notes

<sup>1</sup>Marvin Minsky and Seymour Papert, *Perceptrons* (Cambridge, MA: MIT Press, 1969).

- <sup>2</sup>We have chosen not to develop the distinctions between strong and weak AI, or to mark the distinctions between machine learning and AI. While this would be straightforward, it would unnecessarily complicate this article, and we believed that it would obscure the major conclusions. <sup>3</sup>Neil Gershenfeld, *When Things Start to Think* (New York: Henry Holt, 1999). Gershenfeld argues that computers will disappear as things begin to think.
- <sup>4</sup>Many authors have addressed pieces of this problem. The specific issue of job displacement through technology has been addressed, for example, by Carl Benedikt Frey and Michael A. Osborne, "The Future of Employment," a working paper from Oxford Martin School, Oxford University, 2013, retrieved from https://www.oxfordmartin .ox.ac.uk/downloads/academic/The\_Future\_of\_Employment .pdf. Other references related to the societal impacts of AI can be found here, https://medium.com/@eirinimalliaraki /toward-ethical-transparent-and-fair-ai-ml-a-critical -reading-list-d950e70a70ea. Many cultural critics have addressed aspects of the issues, and we have been influenced by many sources beyond our own experiences.
- <sup>5</sup>An excellent summary of this can be found in Gideon Lewis-Kraus, "The Great AI Awakening," *New York Times Magazine* (December 14, 2016).
- <sup>6</sup>Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin, "Why Should I Trust You?': Explaining the Predictions of Any Classifier," *Cornell University* (last revised August 9, 2016), https://arxiv.org/abs/1602.04938.

- <sup>7</sup>A self-driving car involves much more than the AI system to make it work. It depends on visual systems, including the recognition of objects, reading signs, links to navigation tools, and so forth. For this discussion, we bring these all into one system.
- <sup>8</sup>Ecclesiastes 1:9.
- <sup>9</sup>This insight is attributed to Malcolm Guite, an Anglican priest, poet, lecturer at the Cambridge Theological Foundation, musician and current chaplain of Girton College Cambridge, from a lecture he gave at Cambridge University. He said he "was riffing on the way Dante has the philosophers whirling round in joyful circles in the Paradiso."
- <sup>10</sup>https://en.wikipedia.org/wiki/Fantasia\_(1940\_film) #The\_Sorcerer's\_Apprentice.
- <sup>11</sup>When roughly 40,000 people per year die in car accidents in the US, and 90% of crashes are due to human error (distracted driving, alcohol, sleep), a strong case is sometimes made that the current early version of self-driving cars is already safer than cars with human drivers. This has not been established, and the transition period will perhaps identify new things: S. Singh, "Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey," Traffic Safety Facts Crash•Stats. Report No. DOT HS 812 115 (Washington, DC: National Highway Traffic Safety Administration, 2015), https://crashstats .nhtsa.dot.gov/Api/Public/ViewPublication/812115.
- <sup>12</sup>Note that while Genesis 1 gives God's orderly account of creation, Genesis 2 looks at the same account from a different point of view. The good creation God had made as stated in Genesis 1 means that the world was perfectly provisioned. Genesis 2:5 says that God invited humans to develop this perfectly provisioned world. God could have created a computer bush allowing us to pick what we need, but he provisioned his world and allowed humans to develop a computer.
- <sup>13</sup>Rosie Perera, "Technology: Love It or Hate It?," *Ethix* (February 1, 2008), https://ethix.org/2008/02/01/technology -love-it-or-hate-it.
- <sup>14</sup>Deb Riechmann, "I Never Said That! High-Tech Deception of 'Deepfake' Videos," *The Seattle Times*, originally published July 1, 2018, updated July 4, 2018, https:// www.seattletimes.com/nation-world/nation-politics /apxi-never-said-that-high-tech-deception-of-deepfake -videos/.
- <sup>15</sup>Cathy O'Neil, Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy (New York: Crown, 2016).
- <sup>16</sup>Edward Tenner, *The Efficiency Paradox: What Big Data Can't Do* (New York: Alfred A. Knopf, 2018), xxii.
- <sup>17</sup>Luke Dormehl, *The Formula: How Algorithms Solve All Our Problems ... and Create More* (New York: Perigee, 2014), chapter 2.
- <sup>18</sup>Immanuel Kant, "Moral Philosophy: Collins's Lecture Notes," in *Lectures on Ethics*, ed. Peter Heath and J. B. Schneewind, trans. Peter Heath (Cambridge, UK: Cambridge University Press, 1997).
- <sup>19</sup>Edward Tenner, Why Things Bite Back: Technology and the Revenge of Unintended Consequences (New York: Alfred A. Knopf, 1996).
- <sup>20</sup>Early indications are that self-driving cars will be insured by the auto manufacturer. What unintended consequences might come from this arrangement?
- <sup>21</sup>Justin Anderson, Parker's pastor, brought this up in discussing the cultural consequences of AI, and we found

it insightful. Justin is launching Icon Church (https://iconchurch.org) in the heart of tech country in Seattle in Spring 2019.

<sup>22</sup>Sherry Turkle, *Reclaiming Conversation: The Power of Talk in a Digital Age* (New York: Penguin, 2015).

<sup>23</sup>For many of our routine tasks (e.g., ordering in our local coffee shop), we often mentally reduce the person we are interacting with to a functional level. That person serves to allow me to perform a transaction, and we rarely deviate from thinking of them that way. In other words, we dehumanize them. In fact, many of us might even get a little annoyed if the person taking our order is a little too chatty, a little too human, when all we really want is our coffee!

<sup>24</sup>Randy Beavers, Denise Daniels, Albert Erisman, and Don Lee, "Communication Technology Mediated Relationships: Some Considerations from Theology," *Christian Business Review* 7 (Fall 2018): 12–21.

Business Review 7 (Fall 2018): 12–21. <sup>25</sup>David W. Gill, "Light a Candle," Baccalaureate speech at Gordon Conwell Seminary, *YouTube* (June 2, 2016), https://www.youtube.com/watch?v=AJMWIFwcY0Y.

**ASA Members:** Submit comments and questions on this article at www.asa3.org $\rightarrow$ RESOURCES $\rightarrow$ Forums $\rightarrow$ PSCF Discussion.

## Call for Papers

### TRANSHUMANISM: CHRISTIAN DESTINY OR DISTRACTION?

David C. Winyard Sr. (PhD, Virginia Tech) is an engineer and Science and Technology in Society scholar. From 1975–2012, he worked in R&D for the US Navy and Defense Logistics Agency, receiving multiple awards and patents along the way. After retiring from federal service, he taught at Mount Vernon Nazarene University from 2014 to 2018, building its new engineering program.

Since 2010, Winyard's focus has shifted from technology development to its meaning. His 2016 dissertation analyzed connections between Christianity and transhumanism which seek fundamental enhancements of life by science and technology.

In an essay on the ASA and CSCA websites, Winyard describes for us the latest developments and challenges from transhumanism for our society and Christian faith. The essay is intended as an invitation. Readers are encouraged to take up one of the insights or questions, or maybe a related one that was not mentioned, and draft an article (typically about 5,000–8,000 words) that contributes to the conversation. These can be sent to Winyard at winyard.david@gmail.com. He will send the best essays on to peer review and then we will select from those for publication in a transhumanism theme issue of *Perspectives on Science and Christian Faith*.

The lead editorial in the December 2013 issue of *PSCF* outlines what the journal looks for in article contributions.

For best consideration for inclusion in the theme issue, manuscripts should be received electronically before **October 31, 2019**.

Looking forward to your contributions, **James C. Peterson**, *editor-in-chief*