ARTIFICIALLY INTELLIGENT

What is a Christian supposed to think about artificial intelligence (AI)? What does it mean to be intelligent? What does it mean to be human? Is AI dangerous? Will it be superhuman? There are lots of questions in response to the rise of AI. Let's consider briefly some of the good, the bad, and the ugly.

The Good

Let's begin with the good. The good news about AI arises from God's work of creation and providence.

AI, even in its earlier stages, was a marvel of human engineering. How it displayed the talents and ingenuity and perseverance of its human creators! And we who are Christians know who created the human creators—God. We know that this same Creator also created a world. This world has materials and resources that could be used to plan, construct, and empower the computers that run AI.

AI should stimulate a paean of praise to God. Certainly, we praise the talents and achievements of the human creators. But we find them admirable

precisely because they gloriously and beautifully reflect and emulate the talents and achievements of God. God was not only there at the beginning to create the world and the human race. He is here continually, giving ideas and energy to the human beings who conduct research into AI. The Bible teaches that God gives common-grace gifts to those who are evil as well as those who are good (Matt. 5:45; Acts 14:17; Ps. 94:10–11). Praising the human creators does not mean that we morally approve everything that they do, even within their specialty. It means, rather, that we recognize good gifts and their source in God.

We can also consider the details that go into making computers. Uncommon chemical elements are used in trace amounts to construct the integrated circuits that make up the computers. Chemistry itself, and a supply on earth of these special elements, must be there already. Consider also the complex technology of the miniaturization process, allowing the circuits to be small enough so that millions can be crammed on a single chip. Think of the history of the conceptualization of computers, which goes back as far as Blaise Pascal's calculating machine (1642), named *Pascaline*, and before it, the abacus. Think of the human work to understand electricity and semiconductors, which was a necessary foundation for the construction of increasingly powerful computers.

This element of praise impinges on my personal story. I can remember the time when a computer became

capable of beating most checkers players. And then later, a computer beat the leading grandmaster chess player.

God created human beings not only to think God's thoughts after him, but to work God's works after him.

I can remember when the IBM computer "Watson" played Jeopardy with some of the best human champions from the past. I watched the TV program, and I rooted for Watson to win. I'm a nerd, and I admit it. It is not that I do not admire the human champions. They are impressive, not only in the scope of their knowledge but also in their ability to retrieve it in a moment. They are clearly very intelligent people, as well as people with enormous, active memories. Praise the Lord for their talents, which are gifts from him. So, I do admire the human champions of Jeopardy. But I wanted Watson to win. Why? I admire the engineers who built it too. A win for Watson would be a win for all the people at IBM who worked for that moment of triumph. It was a human triumph that displayed, in a new way, the marvel of God's provision to human beings and the enormous potential that we have to build out our ideas.

Now, we are also seeing practical benefits from AI systems. God created human beings not only to think God's thoughts after him, but to work God's works after him—on the level of the creature, but empowered by the Creator. Technology is one of the results. We can use the internet to look up stuff that is far more extensive than the *Encyclopedia Britannica*. Printing with movable type was an earlier technology that empowered human beings. Now it gets better and more impressive. Think also of the aid to medical diagnosis. A computer can search thousands or even millions of medical records, looking for patterns that suggest what symptoms identify a particular disease, along with its most effective treatments.
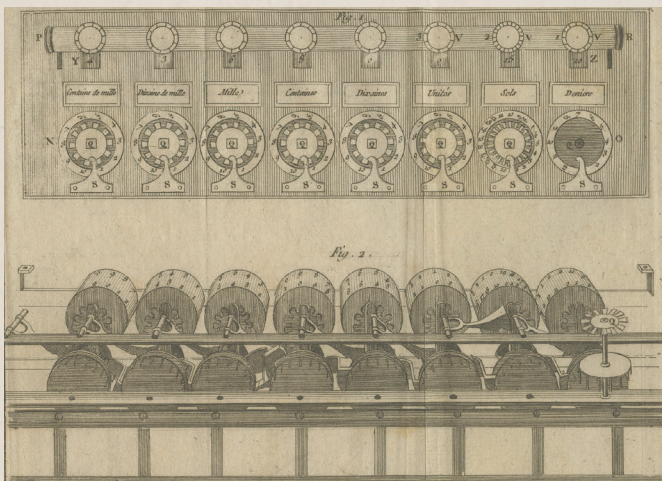
It seems clear that positive trends are only the beginning. We are going to see more and more benefits.

The Bad

The bad news is that human beings, after the fall of Adam into sin, are fallen and sinful creatures. They marvelously, thoroughly, and pervasively distort the thinking of their Creator. They sin, not only in deed, but also in mind and in word.

For this reason, plenty can go wrong with AI—not that there is anything intrinsically evil about a computer. It is a thing, not a human being, who is ethically responsible. But the *people* who program AI and use AI and dream about AI are another matter.

There are at least two opposite tendencies, which paradoxically feed on each other. The first is *pride*. Human beings desire to use AI to become superhuman—virtually divine—in their powers. And, if they can, they will use those powers to control and manipulate other human beings—all the time telling themselves that they have only the best intentions. AI enhances the surveillance technology of totalitarian governments. It is already happening. Or, in the free world, it can be used more gently to remove unacceptable opinions from social media. Some people will use AI to produce pornography, to spread lies, or to help them build terrorist bombs. The underlying root problem is idolatry. Human pride is a form of idolatry in which a person worships himself.

The second tendency is *sloth*. The exertion necessary for normal human living seems too hard. So, AI will write the student's essay for him. He will no longer have the labor and sweat of research and of deep thought. He will farm out the work of analysis and the process of building an expressive written product. He will surrender to an AI chat system such as ChatGPT, which will produce the result for him in seconds.

But writing essays is only the beginning. The man or woman hungering for deep personal companionship will find an AI-simulated therapist who mimics understanding and sympathy and soothes wounded feelings. The simulation may even suggest strategies for solving one's emotional problems. But you may be sure that it will never confront sin.

It simulates, and it might stimulate, but it does not satisfy.

Or the boy or girl will find a way to make AI produce a linguistic simulation of a girlfriend or boyfriend that is less threatening, less selfish, and less prone to break a relationship than a real person. The person retreats from genuine and dangerous uncontrolled social interaction with real people. He retreats to the safety of a simulation. The real person surrenders his personality to what is subpersonal. It simulates, and it might stimulate, but it does not satisfy.

This, too, is a form of idolatry. Ancient paganism contained many instances of bowing down to creatures—to statues, to trees, to alleged sacred places (Rom. 1:18–25). It manufactured idols to simulate God. Below the surface, the person who is alienated from God feels the terrible weight of responsibility and the guilt of making bad decisions. He feels the frustration of dealing with people who are fellow sinners. It is easy for some people to surrender agency to AI, which represents a false source of wisdom. The person becomes a slave to the system.

The book of Revelation has visions of the beast and the prostitute. These visions summarize two main forms of idolatry. The beast is a *power* idol. People worship power, especially the power of the state, as a counterfeit of the power of God. The prostitute is a *pleasure* idol. People worship pleasure, especially the pleasures of sex and luxury, as a counterfeit for the joy of fellowship with God. AI can serve as an idol for both reasons—it is powerful, and it promises to serve our pleasures.

The Ugly

The principal ugly news is that AI, despite its name "artificial intelligence," is completely unintelligent. The doctrine of creation implies limitations. Among earthly creatures, human beings alone, as the image of God, have a personal relationship with their Creator. AI does not create persons. AI understands nothing—zero. It is the human programmers who understand and whom we should rightly acknowldge.

AI *simulates* intelligence. I can illustrate this in the following way. I have seen a video of a robotic dog. It is cute. It simulates a dog. But it is not a dog, and it is nowhere close to being a dog. Or I can use a computer illustration. I have tried playing the computer game "SimCity." It simulates managing a city. The player pretends to be the mayor or city planner, and he lays down roads and rails and business buildings and homes and power plants. But there is no city with bricks and mortar. It is all a simulation of a few dimensions of what it takes to make a city.

The field of AI research distinguishes two levels of AI. One is a specialized AI that addresses a limited field of knowledge—let us say, playing chess or searching records of diseases. Such systems exist even now. The other is "AGI"—artificial general intelligence. The hope is to produce a machine with the abilities of human intelligence across the board, and then to exceed those abilities with superhuman intelligence. AGI is a major goal of some researchers in the field. But is it going to be achievable? On the surface, the closest approach seems to be with AI systems that "chat" and appear to use language just like a native human speaker.

There is a famous article by philosopher John Searle (1932–) that illustrates the issue. The core of the argument is a thought experiment. Searle imagines himself locked in a room. He has a set of instructions equivalent to those of an AI system. The system is programmed to give answers in Chinese to questions posed in Chinese that are slipped under the door. By following the instructions carefully, Searle is able to produce for those outside the door the illusion that inside is a human native speaker of Chinese. However, Searle understands no Chinese at all and understands no meanings of neither the questions nor the answers that are being exchanged.

Like the robotic dog, the Chinese room is a simulation of a Chinese speaker. Due to the intelligence of the programmers, it contains a program that cleverly simulates human intelligence and understanding. But it is only a simulation. Searle understands nothing of Chinese. Likewise, the program understands nothing. The hardware on which the program runs understands nothing. A simulation is a kind of image. Human beings create images. Artists do it all the time. But this kind of image is not the real thing.

Real understanding of meanings, such as we regularly experience with fellow humans, is unique to human beings. AI is a simulation but not the real thing. Not at all. Not close. Adding more lines of programming and more silicon and faster circuits will not change the fundamental difference between a simulation and the real thing.

My wife and I sometimes use Google Assistant in our home. Its most common non-answer is, "I don't understand." True.
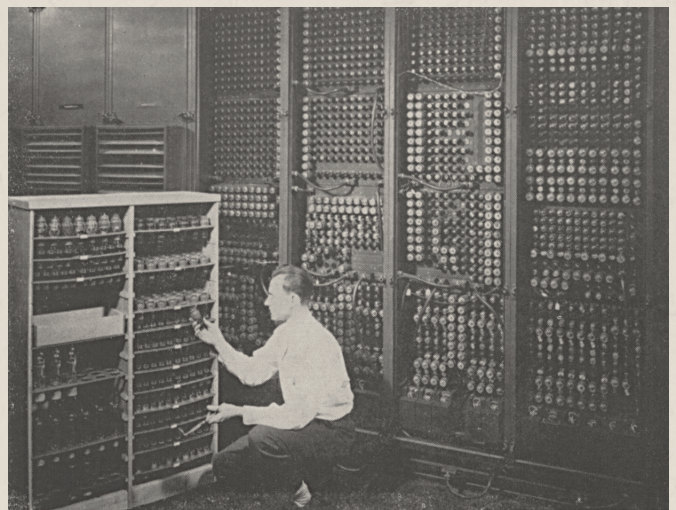
Value of the Products

And what of the products of AI in the intellectual arena? AI can produce impressive results when paired with a well-defined goal, such as winning a game of chess or increasing the accuracy of medical diagnoses. We should be grateful to God and should praise the programmers. But even in such areas, success depends heavily on good input. If the medical records are cluttered with too many inaccurate diagnoses by human physicians, it will do no good to set an AI system to "learn" from the mass of data. Computer programmers have a motto, GIGO: "garbage in, garbage out." In the case of AI, there is always an element of garbage in (GI), namely, the effects of human sin. AI cannot rescue us from sin. Consequently, in any area affecting human living, AI produces garbage out (GO), in unpredictable ways.

There is also the question of ethical evaluation. Suppose an AI system is asked to deal with questions that depend on ethical values. For example, consider a question about the next COVID-like epidemic. Can an AI system deal responsibly with the trade-off between suppressing the transmission of disease on the one hand and giving students the benefits of learning in person on the other hand? Can AI give a responsible answer to the questions of the moral permissibility of abortion or transgenderism? The answers will depend on the ethical inputs. Who decides what these are?

Think of it this way. AI systems are programmed to mash together enormous amounts of data or enormous amounts of language (so-called LLM, "large language model"). But what is the quality of this input, especially with regard to ethics?

What happens in detail cannot be predicted by the user, or even by the programmers, because the system continues to evolve using more and more pieces of language. That seems exciting. The results can, at times, be surprising. But when the base is a large amount of language, what will come out is a mash—a mash that averages out patterns in the language sample. The mashing takes place without actually understanding a single meaning. The result is not going to be creative in the way a human being can be because he actually understands the subject matter and may generate new insights. Rather, AI gives the user a dull average, combining all the creative language of a multitude of human beings into a center sum. It is as if one started with an image representing the star of Bethlehem, a multifaceted, shining, lovely thing, and converted it into a bright yellow blob. AI systems in art can be programmed to look at many features, and not just average a single feature like color or location. But they do so without understanding a thing.

The Good News

The misuse of AI is already around us. AI is new, but the underlying problem is not—it is the problem of sin. And the Bible gives us the one remedy—Jesus Christ our Lord, who died and rose and reigns. Personal fellowship with him, listening to the Bible, addressing him in prayer, and worshiping in person in church are hard and dangerous in comparison with the simulation of an idol. But it is the way in which Jesus Christ works for us and in us, and then also with us. And then we serve him with genuine agency, free from pride and from sloth, and free from the enslavement to our technology.

Vern Sheridan Poythress

Vern S. Poythress

Distinguished Professor of New Testament, Biblical Interpretation, and Systematic Theology
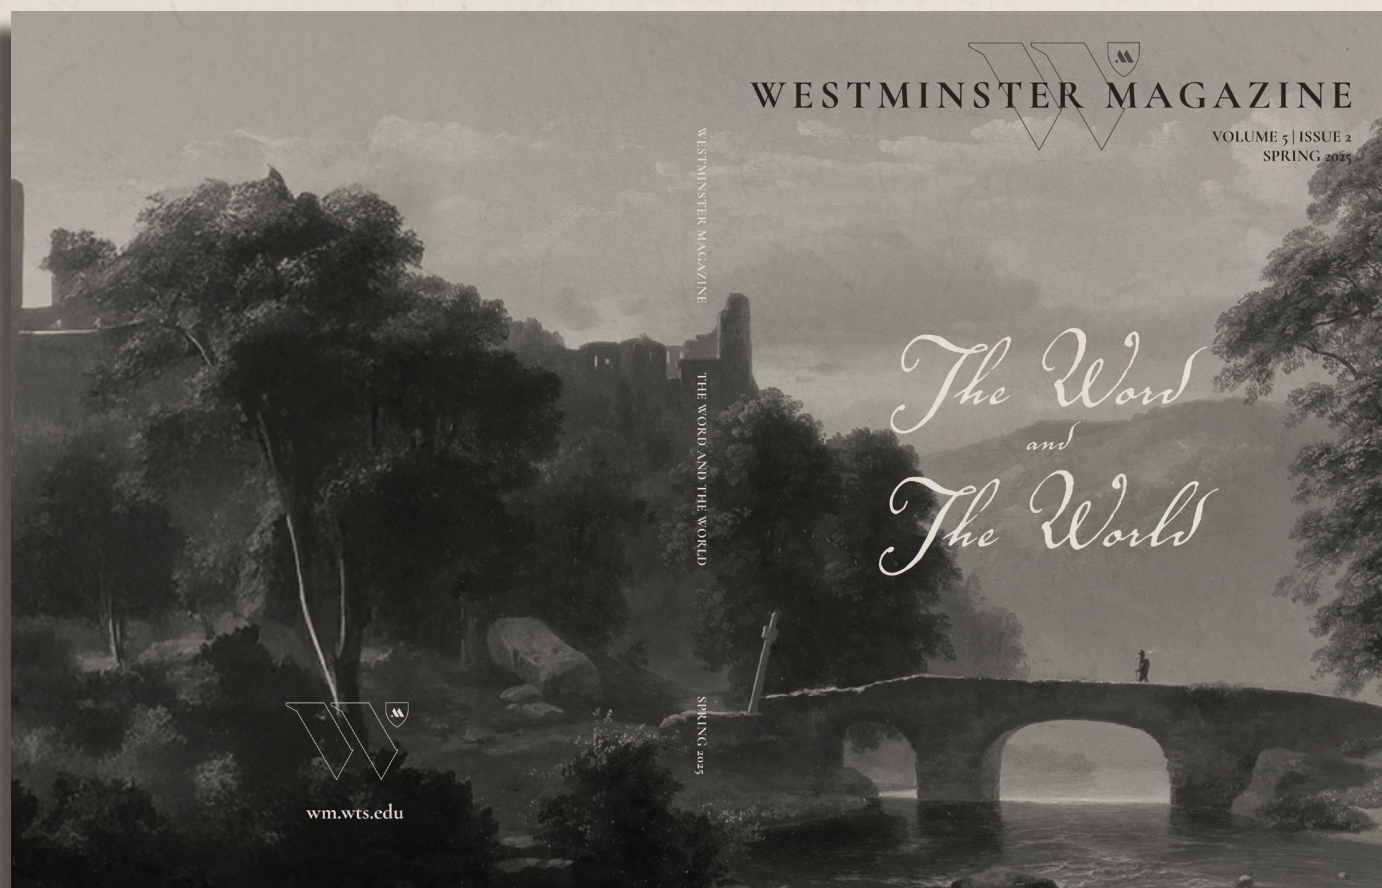

VIA ET VERITAS
ET VITA

Will you join us in proclaiming the glories of Christ in today's world?

WTS.EDU/DONATE

Westminster Theological Seminary
2960 Church Road
Glenside, PA 19038, USA

Stewardship Office: 215-572-3830

POINT *of* CONTACT