ACSI

Leading Insights

Artificial Intelligence

ACSI
STRONGER TOGETHER

Printed in the United States of America
29 28 27 26 25 24 1 2 3 4 5 6

Edited by Swaner, Lynn E.
ACSI Leading Insights: Artificial Intelligence
ISBN 978-1-58331-568-2
eISBN 978-1-58331-569-9
Catalog#: LD01003
eLD01003
Designer: Patrick Flowers
Cover design: Patrick Flowers

Association of Christian Schools International
PO Box 62249 • Colorado Springs, CO 80962
Care Team: 800.367.0798 • www.acsi.org

CONTENTS

Introduction: Framing the AI Question[1]

Lynn E. Swaner, *Series Editor*

Artificial intelligence (AI) is far from new. It is already embedded in existing technologies like smartphone speech-to-text, digital assistants like Siri or Alexa, and GPS-enabled maps or routing apps. However, it was the launch of ChatGPT in November 2022 that brought widespread awareness of AI's potential for the educational sector. This was largely because ChatGPT was the first AI tool that successfully "positioned itself as a disruptive technology that is revolutionising the way students are taught, promoted, and supported in academic environments" (Montenegro-Rueda 2023, 2).

ChatGPT and similar tools are unlike previous AI applications in at least two significant ways. First, they are generative, meaning they use large language models (LLMs) to create new content (such as text or images) based on patterns present in the data they draw from. Second, chatbot tools are applications designed to simulate conversation with human users, typically through text-based interfaces. These tools use various techniques including rule-based systems, machine-learning algorithms, and natural language processing (NLP) to understand user input and generate appropriate responses.[2]

Put simply, these tools are conversational as they draw on nearly limitless data and easily perform academic tasks such as writing essays or providing answers to homework questions. These capabilities imbue the tools with the potential to affect education in significant if not radical ways. Educators around the world reacted to ChatGPT"'s launch with a range of responses, from implementing bans and installing detection software to embracing ChatGPT and similar tools in day-to-day teaching and learning (Gordon 2023).

Increasing educational effectiveness should be a perennial aim in schools, and AI holds promise for achieving this end. However, for

1 Portions of this chapter first appeared as part of a report co-published by ACSI and Cardus (Swaner and Djita 2024).

2 Throughout this chapter, "AI," "ChatGPT," and "chatbots" are used interchangeably to refer to this new class of AI tools that are generative, work off of LLMs, and are conversational (use NLP).

Christian schools, the question of whether to adopt AI in practice is not just one of utility. In *A Christian Field Guide to Technology for Engineers and Designers*, Brue, Schuurman, and Vanderleest (2022) describe "the ultimate and proper goal for technology—to help us be more fully human in relationship to each other and to God" (11). To be sure, the question of how to realize this goal—along with deciding whether and how to adopt new AI technologies in Christian schools—is a complex one.

Three lenses for thinking about AI in Christian schools can assist in catalyzing conversations. While these lenses are non-discrete and overlapping, they can be used together to frame discussions and planning around AI. These three lenses—the *use* lens, the *human* lens, and the *mission* lens—are described below, along with suggested reflection questions that leaders, teachers, students, and the school community together can use as they consider the challenges and opportunities posed by AI.

The Use Lens

The first lens to be considered is that of use. All industries, whether education, law, medicine, insurance, or others, are faced with new AI-driven or AI-supported technologies that are affecting their current work and will be shaping their fields for the future. This lens is often referred to in terms of developing "use cases" for these technologies, with an eye toward how they can help to improve performance in the industry (as mentioned earlier, the current use case discussion in education mostly centers around ChatGPT and related tools).

Employing the *use* lens, Christian schools can consider the following questions:

- As a staff, how can we increase our knowledge around AI (through readings, trainings, conferences, certificate programs, etc.)?
- Can we create small-scale experiments or pilots using AI in teaching and learning, from which we can learn without significant risk?
- How can we network with other schools to identify use cases or to collaborate on AI experiments or pilots?

- How can we effectively engage with stakeholders (leaders, teachers, parents, students, others) in discussions and decisions around the use of AI, whether through a task force or other method?

The Human Lens

Many theologians, public intellectuals, technologists, and ethicists are concerned with the impact of AI on humanity. This second *human* lens centers on the practical yet profound question of how AI will shape human nature and human experiences writ large. Fundamental to answering this question is one's view of what it means to be human, from which flows one's ethical reasoning about technology. John C. Lennox, Oxford professor and author of *2084: Artificial Intelligence and the Future of Humanity*, makes the argument that while most technologies in and of themselves are values-neutral, the human question (of how humans use technology) is what imbues technological trajectories with an ethical dimension. As Lennox (2020) writes, "Of course, experience tells us that most technological advances are likely to have both an upside and a downside …. It is the same with AI. There are many valuable positive developments, and there are some very alarming negative aspects that demand close ethical attention" (54). Thus, while connected to the *use* lens, the *human* lens goes a step further to consider the "why" and not just the "how" or "what" of AI.

Using a *human* lens, Christian schools can consider the following questions:

- What is our theological framework at our school for understanding the nature of humans (for example, as created in God's image), and what are the implications of that framework for understanding human inventions and advances (like AI)?
- What is our educational purpose or philosophy? How are we trying to form our students as human persons? How might the use of AI in our schools enhance or detract from this purpose?
- How do we address ethical thinking at our school, especially when it comes to complex issues in society (including those around technology)?
- Do we offer students opportunities to wrestle with contemporary ethical issues (i.e., through reading, considering opposing viewpoints,

debating)? What training and support does our faculty need to help our students do this well?

The Mission Lens

People from any faith background or none can engage with questions around AI from the use or human lenses. But the question of how AI can be viewed through the lens of Christian mission is, of course, of importance to Christians inhabiting this moment in human history. Both the Great Commission (Matthew 28:16–20) and Great Commandment (Matthew 22:34–40) have implications for the use of any technology: They lead Christians to ask whether (and if so, how) AI can help spread the gospel, make disciples, and better love and serve their neighbors. This is the latest in a long line of similar questions regarding new technologies, similar questions that date at least as far back to the fifteenth century when Christians considered uses for the printing press (with the result that from the Gutenberg Bible to today, the Bible is the most widely printed book in human history). This third lens certainly overlaps with the first and second, and while Christians can and should consider both, they have a unique obligation to pay attention to this lens as well.

Using a *mission* lens, Christian schools can consider the following questions:

- What is our school's theological view of how Christians should engage the world? How does that view inform teaching, learning, and discipleship at our school?
- Given this theological view, along with our school's mission, how can we evaluate the potential of AI for our school—including whether and how AI can be used to nurture Christian beliefs and values, as well as serve others?
- How might our school winsomely engage faculty, parents, students, or other constituents who hold different theological views of AI?
- What resources can we draw on (books, speakers, webinars) that address AI from a distinctly Christian view?

To help Christian schools address this last question, we have developed the monograph that you now hold. This monograph is divided into three

sections: "Philosophy and Research," "Christian School Perspectives," and "Frameworks for Practice." In the first section, authors describe and define AI and its potential for education (chapter 1), present a biblical framework for understanding and responding to AI in the Christian school (chapter 2), and provide a snapshot of AI use across the Christian school sector (chapter 3). In the second section, "Christian School Perspectives," readers will gain insights from a head of school who is leading AI implementation at the middle and high school levels (chapter 4), a Christian school teacher and tech entrepreneur who discusses AI adoption in the classroom (chapter 5), and a focus group of high school students who share their experiences with AI (chapter 6). In the final part of the monograph, "Frameworks for Practice," authors provide a framework for adopting new technology in the Christian school (chapter 7) and suggest ways to consider digital well-being within the Christian school setting (chapter 8).

While AI may differ in significant and potentially profound ways from technologies that have preceded it, Christian school educators can take encouragement in remembering—as well as applying learnings from—previous waves of technological change they have navigated. When it comes to deciding whether or how to engage AI, they will need thoughtfulness and intentionality to chart a deliberate course into the future. We hope and pray this monograph will be a helpful resource along the journey.

Part I:
Philosophy and Research

AI in Education: Villain, Savior, or Something Else?[1]

Derek Schuurman, *Calvin University*

When I was a teenager, I purchased an early personal computer called a Timex Sinclair ZX-81 with money I had earned from my paper route. I was amazed at how computer programs enabled me to build "castles in the air ... creating by exertion of the imagination" (Brooks 1995). What started as a hobby later developed into a vocation as I worked in industry as an engineer and later pursued graduate work in the field of robotics and computer vision.

As I pursued my graduate studies years ago, I found myself attracted to newer machine-learning methods that were being used in computer vision (Fortuna et al. 2002). I recall being astounded at the profound elegance of "training" a computer with a set of example images and then observing how well it could identify new images that were not part of the original training set. Even those early machine-learning techniques seemed magical.

Two things became apparent to me in the following years. First, the technology amplified opportunities to do good as well as to do harm. Already as a grad student, I had observed many research efforts being directed toward face recognition—an intriguing and challenging technical problem that had pitfalls for misuse and a myriad of privacy issues. I consciously chose a research direction that I felt was a more redeeming application of machine learning, such as automating the visual sorting of recyclable goods (House et al. 2011). I later recognized this approach as illustrating the theological notion of *structure* and *direction*: The possibility for machine learning is rooted in the *structure* of God's good creation, and *direction* refers to how we unfold this technology in either obedience or disobedience to God (Wolters 2005).

1 Portions of this chapter first appeared in the following articles by the same author: "ChatGPT and the Rise of AI", Christian Scholar's Review blog, January 20, 2023 (https://christianscholars.com/chatgpt-and-the-rise-of-ai/); "AI and Truth in a Post-Epistemic World", Christian Scholar's Review blog, February 27, 2024 (https://christianscholars.com/ai-and-truth-in-a-post-epistemic-world/); and "AI and Truth", Christian Courier, December 4, 2023, p. 20 (https://www.christiancourier.ca/ai-and-truth/).

The second thing that became apparent to me was that AI was developing faster than many of us would have predicted. As an engineering grad student some twenty years ago, I would have scoffed at the notion of an autonomous car; the computer vision challenges were simply too great in unstructured and unpredictable environments. However, within the decade, Google successfully demonstrated a self-driving car.[2] In the words often attributed to Yogi Berra, "It's tough to make predictions, especially about the future"—even for those who are developing technology.

One of the latest developments to catch widespread attention has been ChatGPT, a chatbot developed by OpenAI. ChatGPT relies on a large language model (LLM) that generates "statistically likely continuations of word sequences" to interact with a user by responding to questions and replying to prompts (Shanahan 2024). While some of the responses are amusing or simply wrong, the results are frequently astonishing, providing surprisingly coherent and cogent responses to a wide variety of prompts, including composing poems, stories, sermons, and essays. Unlike the modest amount of training data I used in my graduate work, ChatGPT-3 used 570 gigabytes of example documents (Tamkin et al. 2021). These LLMs demonstrate "that extraordinary and unexpected capabilities emerge when big enough models are trained on very large quantities of textual data" (Shanahan 2024, 73).

The results have been so remarkable it has led to speculations that essay writing is obsolete and has raised uncertainty about the future of many skilled jobs.[3] Indeed, computer programmers may be programming themselves out of a job. A tool called GitHub Copilot takes input prompts and generates computer code, leading some to speculate about the end of programming (Ito 2023). While rumors of the demise of writing and of programming are likely exaggerated, there will be definite impacts for education (Dakhel et al. 2022).

2 A short history of the Waymo project is available at https://waymo.com/about/#story.

3 For example, see Stephen Marche, "The College Essay is Dead," *The Atlantic*, December 6, 2022, and Paul Krugman, "Does ChatGPT Mean Robots are Coming For the Skilled Jobs?" *The New York Times*, December 6, 2022.

Which Way Is AI Headed?

Some are making dramatic claims about how AI will improve education. In an article titled "Why AI Will Save the World," web pioneer Marc Andreesen predicts that "Every child will have an AI tutor that is infinitely patient, infinitely compassionate, infinitely knowledgeable, infinitely helpful" (2023).

Such optimistic claims are not new. The rise of the World Wide Web came with predictions that it would be "a natural force drawing people into greater world harmony" (Negroponte 1996). As we know, while the web provided unprecedented access to information, it did not lead to greater world harmony. It has often become a medium for misinformation, echo chambers, bullying, and polarization.

Likewise, the development of AI will bring benefits in many areas like medicine, enhanced traffic safety, and environmental monitoring, to name a few. But there will be many pitfalls to address. Eric Horvitz, Microsoft's chief scientific officer, sounded a prophetic alarm about one of those pitfalls. He wrote that AI is moving us closer to a "post-epistemic" world "where fact cannot be distinguished from fiction" (2022). In short, AI will have an impact on how we perceive the truth, and educators will need to take heed.

Horvitz identifies one area in which AI will obfuscate the truth: "deepfakes." Deepfakes use AI to create synthetic videos that can impersonate people. In a famous demonstration, researchers at the University of Washington posted a deepfake video of President Obama, making him say whatever they wanted (Langston 2017). Seeing is no longer believing; "truth" can now be manipulated and fabricated.

Another recent development is "astroturfing"—using AI to generate a fake campaign that gives the illusion of a grassroots movement. AI chatbots can be harnessed to post massive amounts of tailored content on social media with the purpose of capturing attention and manipulating people's opinions. Some predict that astroturfing will increasingly distort truth and reality, a prediction that if realized, would pose a direct threat to democratic societies (Schneier and Sanders 2023).

Another impediment to discerning truth occurs when LLMs generate false information referred to as "hallucinations." This should come as

no surprise: By their very architecture, LLMs are "simply a system for haphazardly stitching together sequences of linguistic forms … without any reference to meaning: a stochastic parrot" (Bender et al. 2021). According to Grady Booch, an IEEE (Institute of Electrical and Electronics Engineers) Fellow and chief scientist for engineering at IBM, "Generative modes are unreliable narrators" that can "generate misinformation at scale," and they are now being "unleashed into the wild by corporations who offer no transparency as to their corpus" (Anderson 2023). According to the research of one company, "A.I. chatbots invent information at least 3 percent of the time, and some as much as 27 percent of the time" (Metz 2023).

Other researchers have begun to recognize that AI chatbots are trained with a particular worldview and users are subject to something called "latent persuasion" (Jakesch et al. 2023). Regular usage of AI chatbots can be like having a Jiminy Cricket on your shoulder, autocompleting your thoughts. Over time, such nudging can shape your opinions without your realizing it. In fact, one startling study demonstrated that increasing use of AI is correlated with a decline in religiosity (Jackson et al. 2023). All of this can lead to a kind of syncretism, in which Christians amalgamate secular ideologies promoted by AI alongside Christian thought.

Discerning a Response to AI

We shape our tools, but our tools can also shape us—including shaping our perception of truth. What follows are three general guidelines for Christian educators as we discern a Christian response to AI.

First, we need to avoid the pitfalls of viewing technology with either too much optimism or with undue pessimism. We must reject a reductionistic worldview that sees everything (including education) as amenable to technical solutions. A trust in technology and progress, sometimes referred to as *technicism*, is essentially a form of idolatry (Schuurman 2013, 60). On the other hand, we should not view technological developments with a despair that they will inevitably threaten humanity. AI is part of the latent potential in creation, and we are called to responsibly unfold its possibilities (Schuurman 2022). Theologian Al Wolters (2005) writes that "the Bible is unique in its uncompromising rejection of all attempts … to identify part of creation as either the villain or the savior" (61).

Second, rather than focusing on what AI can *do*, we need to start with an ontological question: *How are people distinct from machines?* A common tendency is to anthropomorphize our machines, thereby elevating the status of our machines and, in doing so, reducing the distinctiveness of human beings. "Exchanges with state-of-the-art LLM-based conversational agents, such as ChatGPT, are so convincing, it is hard to not to anthropomorphize them." But we should "resist the siren call of anthropomorphism" (Shanahan 2024, 73, 79).

In *Humans Are Underrated*, Geoff Colvin suggests asking the following question: "What are the activities that we humans, driven by our deepest nature or by the realities of daily life, will simply insist be performed by other humans, regardless of what computers can do?" (2016, 42). Already in the 1960s, the early AI pioneer Joseph Weizenbaum explored the notion of automating psychotherapy with a chatbot named ELIZA. He concluded, "There are limits to what computers ought to be put to do" (1976, 5–6). An AI chatbot or robot should never substitute for the human wisdom and empathy of a caring teacher. Without a biblically informed ontological grounding, we will be susceptible to various reductionistic philosophies like physicalism and Gnosticism (Schuurman 2019, 79). The biblical story is clear that while humans are also creatures, we are uniquely created in the image of God and distinct from machines. The notion of the *imago Dei* endures even as our machines become more capable. Theologian Herman Bavinck argued that "a human being does not *bear* or *have* the image of God, but … he or she *is* the image of God" (2003, 554, emphasis in original).

Third, we need to discern *norms* for the responsible use of AI. The creators of ChatGPT bumped up against the "AI Alignment" problem—the challenge of aligning an AI system with the goals and values of the designers. The developers had to grapple with bias (including racism) in their training set. Technology is not neutral, and neither are the algorithms and the training data used in AI. Consequently, AI systems can perpetuate injustice, a real threat as big data is employed in a wide variety of fields including insurance, policing, marketing, loans, and politics (O'Neil 2016). Furthermore, AI tends to favor efficiency over other normative considerations. The Christian philosopher of technology, Jacques Ellul,

warned against the ideology of *technique*, which he defined as the drive for "absolute efficiency" applied to "every field of human activity," including education (1989, xxv). We must resist absolutizing efficiency and instead discern creational norms for AI that include considerations like justice, cultural appropriateness, caring, social norms, stewardship, transparency, and trust (Schuurman 2019, 71–108).

Beginning with Ourselves

How do we discern normative uses for AI in education? It begins with ourselves. Ellul provides general advice in his book, *The Presence of the Kingdom.* He points to Romans 12:2, "Do not be conformed to this world, but be transformed by the renewal of your mind, that by testing you may discern what is the will of God, what is good and acceptable and perfect." Ellul argues that "faith produces a renewal of intelligence" and that it "takes place in Jesus Christ, through the action of the Holy Spirit" (1989, 80–81). Ellul argues that this requires a "new style of life" that includes the "whole of life," from "the way we dress and the food we eat" to how we treat our neighbors (119–122).

In our context, a new style of life might include habits of mind and the cultivation of various intellectual virtues. In his book *Epistemology: Becoming Intellectually Virtuous*, author W. Jay Wood reminds us that "wise persons not only possess knowledge of eternal or ultimate significance but have undertaken to become the kinds of persons who naturally desire and pursue this knowledge" (Wood and Wood 1998, 69). Some countercultural habits that might help us cultivate wisdom include observing Sabbath and limiting our exposure to the constant stream of AI-driven media.

In the mid-twentieth century, C.S. Lewis described the pitfall of developing a "blindness" to certain truths by reading "only modern books." His advice is "to keep the clean sea breeze of the centuries blowing through our minds" by "reading old books" (2014). In our twenty-first-century era, "modern books" are no longer the issue, but rather "modern media." None of us are immune to the "blindness" that may be caused by misinformation, latent persuasion, and astroturfing. We ought to renew our minds with the clean sea breeze of older books including, of course, the Scriptures.

Charting the Course

Ultimately, appropriate norms should point us toward using AI to open up new possibilities for showing love to our neighbor and caring for the earth and its creatures. Already, AI has shown amazing redemptive applications in medicine, drug discovery, environmental monitoring, wildlife preservation, assisting people with disabilities, and enhancing traffic safety. These are fruitful directions for computer scientists to explore. However, computer scientists will need the help of philosophers, theologians, educators, social scientists, and others in the humanities to help direct technologies like LLMs in normative ways. Christian educators must join the wider conversations around AI and discern its impact on teaching, learning, and advancing truth. We ought to heed the apostle Paul's injunction to "test everything; hold fast what is good" (1 Thessalonians 5:21, RSV).

Frederick Brooks, a respected Christian computer scientist, wrote, "It is time to recognize that the original goals of AI were not merely extremely difficult, they were goals that, although glamorous and motivating, sent the discipline off in the wrong direction" (1995). He advocates for IA (intelligence amplifying) systems over AI, suggesting people and machines will be able to do far more than AI alone. As an example, one of my colleagues at Calvin University has been exploring the use of AI for helping people write better as opposed to writing for them (Arnold 2021). It is my strong sense that such an approach will be the most fruitful in education.

We shape our tools, but our tools can also shape us—including shaping our students and their perception of truth. Despite the possibilities for sinful distortions, AI is part of the exciting possibilities in creation that Christians can help unfold in God-honoring ways. Christians will need to join the wider dialogue surrounding AI, recognizing that AI is neither the villain nor the savior, bringing biblical insights into what it means to be human, and discerning norms for its appropriate use in education.