

A Christian Perspective on Artificial Intelligence: How Should Christians Think about Thinking Machines?

Steve VanderLeest¹ and Derek Schuurman²

Abstract

Artificial Intelligence (AI) is a topic that deserves attention and careful thought from Christians working in technology. The technology of AI has shifted from the subject of science fiction literature to become the goal of serious engineering development. Recent developments have led to more machine-like humans with repairs and augmentation of our physical bodies through artificial limbs, artificial hearts, and replacement hips. More ambitious research plans are underway to study the brain with the goal of replicating it. Ray Kurzweil, director of engineering at Google, has publicly declared his goal to live long enough that technology will have developed far enough to allow him to download his brain into a computer and thereby achieve immortality. Accompanying recent technological developments in AI have been calls for caution, warning society that AI is different than earlier technology inventions in ways that could bring unforeseen consequences or irreversible harm. In order to understand whether it may be possible to build a human-like machine, we must first understand what it means to be human. This paper will not presume to fully answer this age-old question, but will pose some relevant questions and attempt to catalog some of the attributes that might be key to the definition. Having identified some candidate attributes of the human identity, we then turn to the scriptures, identifying some biblical principles that may be helpful in considering AI and what it means to be human, organized into the themes of creation, fall, and redemption. The paper concludes with a call to responsibility and humility.

1. Introduction

Artificial Intelligence (AI) is the idea of developing computers that can think, performing tasks that have normally required human intelligence. While AI arose as a possibility quite early in the history of computing, the concept of a machine that could mimic a human has a much longer history, going back hundreds of years to the idea of automatons. Science fiction literature and film has a rich tradition of exploring the ideas of AI and their implications. For example, the classic 1982 sci-fi film *Blade Runner* depicted human-looking androids called “replicants” which could think as well as humans and were physically stronger. One of the officers on the starship *Enterprise* in the 1990s television series *Star Trek: The Next Generation* was Lt. Cdr. Data, a human-looking android with great physical strength and great computing capability. Both the movie and television series depicted machines that were surprisingly human-like – exhibiting physical characteristics, but often poignantly displaying human emotions and values. However, these are simply stories played out by human actors and not possible in real life. Or is it

¹ Calvin College, Grand Rapids, MI, USA

² Redeemer University College, Ancaster, Ontario, Canada

possible? For another perspective, consider the converse. Science fiction has not only imagined human-like machines, but also machine-like humans. Consider the 1970s television series *The Six Million Dollar Man*, where fictional character Steve Austin suffers a terrible accident resulting in an extensive surgery to rebuild the injured parts of his body with artificial parts, resulting in augmented abilities for his legs, right arm, and one eye. What seemed visionary 40 years ago has partly slipped into reality, with substantial advances in the technology and science of artificial limbs, though an artificial eye with the capabilities imagined by the TV show has yet to be realized.

In this paper, we explore the topic of AI from a Christian perspective. We have divided our discussion into three sections. First, we make the argument that the topic of artificial intelligence is worthy of consideration by Christian scholars. Technological advances in AI are hinting at some incredible opportunities, but also warrant significant caution. Christians can and should be part of this debate. Second, we explore one of the key questions that underlie many of the issues that AI raises: what does it mean to be human? We dare not claim to answer this question that has perplexed philosophers, poets, and theologians throughout the ages. Rather, we'll mention a few interesting features of the human landscape, surveying and evaluating some of the potential litmus tests that have been suggested for our self-identification problem. Third, we explore a number of relevant scriptural themes that apply to AI as well as to our broader technological developments.

2. Artificial Intelligence is a Worthy Question

Artificial intelligence is a timely and important topic for Christians to consider as the gap between machine and human shrinks. Although we do not traditionally consider humans adding technology to themselves (machine-like humans) as a type of AI, we'll briefly consider it in this section to set the stage for exploring the more traditional concept of AI, machines with human-like attributes.

2.1. Machine-Like Humans

Humans have been augmenting themselves with technology since the beginning. Marshall McLuhan suggests that technology and media are “extensions” of ourselves.¹ Tools are, by definition, an augmentation of some ability, from hammers that help us pound harder to telescopes that help us see further. Sometimes we use technology to repair, such as using a splint to guide the healing of a broken bone. Despite this close relationship with our tools, until recently the line between ourselves and our technology was fairly clear. When the technology becomes an integral part of our bodies, the line is a bit more ambiguous. For a person with an artificial hip or heart, most would agree on which part of the person was human and which part was technology. It gets a bit fuzzier if we modify someone's cells using gene therapy. A machine of cold metal is easily identified as technology, but warm, supple artificial skin might fool us. Today's artificial limbs are surprisingly sophisticated, and current research is already experimenting with artificial limbs controlled directly by a person's nervous system.

The Christian philosopher Henk Geertsema writes that “the development and use of technical devices to heal or improve certain functions of the human body does not invalidate the difference between human beings and machines.”² But how much of ourselves can we replace and still remain human? What if I replace an amputated leg with an artificial limb? Surely, I am still human. What if I replace a faulty heart with an artificial organ? Although in ancient times the heart might be considered the seat of emotion and central to my humanity, in modern times my artificial heart would not disqualify me as being human. What if I start augmenting my brain? Would a “brain prosthesis” ever be possible? Nanotechnology enhancements might still be years away, but one could consider cochlear implants for the profoundly deaf to be a forerunner of brain augmentation. Suppose a microscopic computer is developed which can replace the function of a single neuron. Furthermore, suppose we begin using this to replace neurons in the brain. Replacing one neuron would have little or no effect, but replacing every neuron would essentially replace the brain with a computer. How much of our biological brain could I replace before I am no longer human?³

2.2. Human-Like Machines

While the previous section considered whether adding technology to ourselves might make us less than human, the question of this section is whether it is possible that technology might become more than machine, to be sentient. From the early days of computing, entertaining examples of software pretending to be a person have arisen, such as the ELIZA program written in the 1960s that responded in natural language using scripted pattern matching.⁴ Today’s expert systems can be surprisingly humanlike, such as computerized call center operators that understand a wide variety of spoken phrases, or the IBM Watson supercomputer that can defeat even the best humans on *Jeopardy!*, the popular TV game show. Even our smartphones are getting smart with us, with Siri providing helpful and sometimes entertaining answers to our queries. The Google autonomous car has demonstrated the ability for intelligent software to navigate a vehicle in complex environments. Web search engines routinely provide relevant results with both accuracy and speed.

A recent survey of current research in AI begins by noting how far the state of the art has advanced.⁵ The author notes that “[s]ome forms of computer vision and natural language processing can currently be done quite well.... The current generation of autonomous vehicles do quite well on the much more challenging problem of off-road driving. For certain kinds of technical image analysis, (e.g. in medical applications), computers do as well or better than human experts.” Despite these notable examples, AI still faces many fundamental challenges. The author concludes that even for simple narrowly defined tasks, AI still generally lags behind human abilities.⁶ It was also observed that AI capabilities often reach a plateau and that any incremental improvements typically require tremendous efforts and computing power. One area that presents particular challenges for AI is the area of common sense reasoning. A category of questions called “Winograd schemas” can be used to test such reasoning. An example is the following: “The man couldn’t lift his son because he was so heavy. Who was heavy?”⁷ Statements like these simply require identifying who the pronoun “he” refers to, yet they rely on

broader knowledge to infer that heavy items are more difficult to lift. Nevertheless, researchers remain busy trying to tackle problems such as these.

Ray Kurzweil, the inventor now working at Google on natural language processing, is famous for his publicly stated aim of transferring his consciousness into a computer and thereby becoming immortal. If Kurzweil tried it and the machine then said “I am Ray,” would we consider it to be human? Could we ethically turn it off or destroy it, saying it is simply a computer and therefore we can do what we wish with it? How certain would we be? What if it (he?) cried out for help and mercy? What role does humility play in considering such a technical marvel – or perhaps monstrosity? Does humility say humans should never dare to develop such devices? Kurzweil is not the only computer scientist imagining such a development. Danny Hillis, in a famous essay on AI, thought similarly: “Of course, I understand that this is just a dream. And I will admit that I am more propelled by hope than by the probability of success. But if, within this artificial mind, the seed of human knowledge begins to sustain itself and grow of its own accord, then for the first time human thought will live free of bones and flesh, giving this child of mind an earthly immortality denied to us.”⁸

2.3. Caution: Dangerous Curves Ahead

Futurists like Kurzweil look forward with nearly unbridled optimism towards the time when computers surpass us. “Before the next century is over, human beings will no longer be the most intelligent or capable type of entity on the planet.”⁹ But not all technology experts are so sanguine. Tom Dietterich and Eric Horvitz, the current president and a former president of the Association for the Advancement of Artificial Intelligence, highlight several risks and warn of the dangers of software errors in AI software.¹⁰ The growing complexity of AI software presents numerous challenges, especially when it is used to control automobiles, surgical robots, and weapon systems. It is a particular challenge to verify systems which rely on “machine learning” techniques. Another risk is the possibility of cyber-attacks, and AI systems which control safety-critical systems are also vulnerable to such attacks. Finally, the authors identify the risk illustrated in the tale of the “sorcerer's apprentice.” What happens if a super-intelligent AI program runs amok and out of control? Others have voiced similar fears. Bill Joy, Sun Microsystems chief scientist at the time, wrote a famous article for *Wired* magazine shortly after a thought-provoking encounter with Ray Kurzweil. He writes about AI that “Thus we have the possibility not just of weapons of mass destruction but of knowledge-enabled mass destruction (KMD), this destructiveness hugely amplified by the power of self-replication. I think it is no exaggeration to say we are on the cusp of the further perfection of extreme evil, an evil whose possibility spreads well beyond that which weapons of mass destruction bequeathed to the nation-states, on to a surprising and terrible empowerment of extreme individuals.”¹¹ Others are even more frank about these concerns. Stephen Hawking warns in a BBC interview that the “development of full artificial intelligence could spell the end of the human race.”¹² The entrepreneur and engineer Elon Musk says “I think we should be very careful about artificial intelligence. If I had to guess at what our biggest existential threat is, it's probably that.”¹³

3. What Does it Mean to be Human?

As we develop technology to provide some amazing capabilities, does there ever come a point where we have crossed a line, where machines have become human in some sense? Or is that impossible? To answer that question, we need to understand what it means to be human.

Humans have always had an identity crisis. For much of our recorded history, we have used rather specious definitions of humanness or personhood that granted power to some, while granting few or no rights to others. At times some have thought our gender or the color of our skin formed a key part of that definition. If we have erred too narrowly in the past, do we now risk erring too widely? In this section, we first consider the classical division of a person into mind and body. Second, we identify a few human attributes that might be considered essential to our self identity. Third, we describe the Turing test for intelligence and one example critique of the test.

3.1. Mind, Body, Soul

What makes up the human person? Different schools of thought in ontology (the philosophy that explores the nature of being or existence) have suggested anthropologies that affirm or deny the existence of at least three different parts: the body, mind, and soul. The body is composed of our physical self, including our neurons and brain. The mind consists of our thoughts and consciousness. The soul is “that part of us that might be said to be eternal or to transcend in some way the mortal body.”¹⁴

Most anthropological views can be categorized as either *monism* or *dualism*.¹⁵ Monism asserts that humans are made of one substance. Thomas Hobbes was an early supporter of monism by arguing that consciousness and souls arise from the functions of the body alone.¹⁶ In contrast, dualism holds that humans are somehow made up of two parts, often identified as the body and the soul. Dualism includes many theories about how the body and soul are separate but related. Platonic dualism saw the body as an earthly package for the spirit, something to be eventually discarded. René Descartes, an early modern philosopher who promoted a form of dualism, suggested the body was like a machine that interacted with the mind. Although the Bible is not a philosophical anthropology textbook, there are many verses indicating that we are more than our bodies. For example, Paul says that we “are confident, I say, and would prefer to be away from the body and at home with the Lord” (2 Corinthians 5:8, NIV). The view of “holistic dualism” acknowledges both the wholeness of body and soul as well as affirming the notion that one continues to exist after death without an earthly body.¹⁷

A third view which is far less common is *trichotomy*, the notion that humans are comprised of three components: body, soul, and spirit. The spirit is the human self, where the soul mediates the spirit and the body.¹⁸ Trichotomists could appeal to literal interpretations of New Testament verses such as 1 Thessalonians 5:23 “May your whole spirit, soul and body be kept blameless at the coming of our Lord Jesus Christ.” However, care should be used when using the scriptures in

this way. The point of this verse does not appear to be a lesson on anthropology, but rather that our whole person be preserved and kept.¹⁹

More recently, many modern western philosophers have embraced *materialism*, which is a form of monism that denies the presence of a soul and holds that reality is made up of only the physical stuff around us. In his book, *The Concept of Mind*, Gilbert Ryle rejects dualism and ridicules it as “the myth of the ghost in the machine.”²⁰ More recently, a materialist view has been promoted by Ray Kurzweil in a series of books such as *The Age of Spiritual Machines* and *How to Create a Mind*. This perspective dismisses the notion of a soul, concluding that our mind and consciousness arise entirely from the physical brain. Some materialists account for the complexity of the mind by attributing it to the interactions of many simple entities, like an ant colony. Although each ant appears to act at random, more complex behavior emerges at the level of the colony.²¹ “The notion of emergence would suggest that such a network, once it reached some critical mass, would spontaneously begin to think.”²²

Our view of human personhood has profound implications. For instance, a materialist view applied to the mind (sometimes referred to as *physicalism*) will conclude that all illnesses of the mind or spirit are reducible to an illness of the body which can be treated by pharmaceuticals.²³ Matthew Dickerson argues that physicalism has significant implications for areas like creativity, heroism, ecology, as well as for reason and science.²⁴ Despite its rejection of the spiritual aspect, materialism itself has religious aspects. The belief that there is nothing but the physical is itself a religious belief since it identifies the physical as an unconditional reality that is the ultimate explainer of our existence. For a physicalist, being human simply reduces to the interactions of basic particles.²⁵ However, if we consider ourselves as more than simply a physical body, how does that shape our view of what it means to be human?

3.2. Which Attributes are Essentially Human?

In this section we inventory a number of the characteristics that some have suggested are essential human qualities, i.e., attributes that define us as human. We do not claim the list is exhaustive. Rather, it demonstrates the range of ideas thought to be core to our identity.

3.2.1. Intelligence

Some technologists working in AI have not aimed for replication of humans, but rather for intelligence. Is intelligence, the ability to learn and to apply that learning, an essential quality of humanness? Is it a unique talent of humans alone, unattainable by any other natural or artificial creature? Could a machine have excellent logic and rationality, surpassing humans at deductive reasoning? The IBM supercomputer Deep Blue was able to beat the world chess champion Garry Kasparov in 1997. If a machine can play chess better than the very best human player, does that make it intelligent?

3.2.2. Sentience

Closely related to intelligence, sentience is the ability to perceive. Is perception an essential human quality? Is consciousness, self-awareness, the key ingredient to personhood? The AI

community has long struggled with what self-awareness means. Douglas Hofstadter, in his famous 1979 book *Gödel, Escher, Bach*, explored ideas of recursion, self-reference, and the idea of the “strange loop” as possible layers that might allow the whole to be greater than the sum of its parts. This is the idea of emergence, that simple components can interact so that a more sophisticated, perhaps intelligent, behavior emerges.

3.2.3. Emotion

Emotion is often considered a part of our intellect, but a peculiar component that is not logical or calculating, even though it can often be predicted. Emotion seems to be connected both to our state of mind and to our bodies. Emotion makes our hearts race and our hands sweat. It puts the bounce in our step or the frown on our face. In order to feel emotions, one must have both intelligence and self-awareness. For example, someone becomes angry after a barrage of insults only because they both understand the meaning of the insults and they perceive the insult as directed at them personally.

3.2.4. Soul

As Christians, we consider the soul an essential part of our being, in fact, the one part that survives our death. In addition, this is often the attribute that many believe uniquely defines us as human, particularly when other attributes do not seem sufficiently unique because we find them at least partially in other creatures. For instance, we observe emotion in chimpanzees or dolphins and might perceive intelligence in a chess-playing computer. By contrast, the soul seems to be confined to humans. However, one cannot measure for the presence of a soul as a test of humanity. The computer scientist Matthew Dickerson makes the astute point that assuming we can scientifically test for the spiritual assumes that the spiritual is reducible to the material, which is equivalent to saying that the spiritual does not exist.²⁶

3.2.5. Living being

If our body is an essential part of our humanity, then it seems biology is a necessary component of humanity. This is the part of our humanity that we share with other living creatures. However, having a natural biological birth does not appear to be part of that qualification, since we do not see any distinction for test-tube babies. It remains to be seen how we would treat human clones, or humans with significant alterations to their genetic makeup.

3.2.6. Creativity, Use of tools

Rather than distinguishing humans by how they think, as *homo sapiens*, many point to our ability to make tools as what distinguishes us from other creatures. Thus we are *homo faber*, humans as makers. Inventing novel devices, composing new music, and innovating in business are all examples of creativity that may also be hints of an essential quality of our humanity. Humor is a type of creativity required to banter about with one another, and laughter is sometimes considered uniquely human.

3.2.7. Free will

Do humans have free will, the ability to choose? Or do our circumstances, genetics, and state of mind determine our course of action? Vaclav Havel pointed to this attribute as essential: “The

secret of man is the secret of his responsibility.” We cannot be held morally responsible for an act unless we have a choice (to act or not to act). Moral agency, the ability to choose, and to be held morally accountable for our choice, is perhaps uniquely human. While some would argue that a computer can never be human because it cannot truly make a free will choice, others counter that humans cannot make a free choice either, thus subscribing to some version of determinism.

3.3. Implications of Materialism

Taking a materialist view to its logical conclusions would deny the very possibility of many of the attributes listed above. Physicalism applies the principles of materialism to the brain such that “our brains are natural phenomena and this must follow the cause-and-effect laws manifest in machines.”²⁷ A strict physicalist view would deny the presence of a soul, suggesting we are just bodies operating under physical laws. In addition, it would reject the notion of free will. Kurzweil notes that “if human decision making is based on such predictable interactions of basic particles, our decisions must also be predetermined. That would contradict human freedom to choose.”²⁸ Furthermore, if our thoughts are merely the “interactions of basic particles” then true creativity is also an illusion. Matthew Dickerson argues that “to the extent that creativity is defined in terms of originality ... physical automata, whether the digital computer variety or the biochemical human variety, are not capable of originating anything.”²⁹ According to this view, emotions would also be simply reduced to the interactions of basic particles.

Materialism and physicalism are highly reductionistic with profound implications on how we view our humanity. Human attributes such as souls, free will, creativity, and emotions are essentially an illusion reducible to the laws of physics. Furthermore, it has implications for our understanding of knowing and truth. Materialism is a type of naturalism, and C.S. Lewis observes that naturalism “offers what professes to be a full account of our mental behaviour, but this account, on inspection, leaves no room for the acts of knowing or insight on which the whole value of our thinking, as a means to truth, depends.”³⁰ Ironically, a physicalist view of reality even leads to devaluing of our physical bodies, potentially leading to a new kind of gnosticism. Kurzweil suggests that “We don’t always need real bodies. If we happen to be in a virtual environment, then a virtual body will do just fine.”³¹

What are the attributes that define our humanity? Even if one rejects a physicalist view, it is not clear that any of these attributes, or any other proposed characteristics can definitively categorize humans and nonhumans. Not only is it difficult to conclusively identify those attributes that are sufficient, it is also difficult to simply list which ones are necessary. Perhaps the value of such a list is not as a tool to determine who is in the human “club,” but rather to encourage and challenge each other to flourish and grow towards the best humans we can be. In order to determine what is best, we must identify our purpose. However, before we turn to that subject through the lens of Scripture, let’s look at one other proposed test, a rather well-known one within the AI community.

3.4. The Turing Test

Rather than attempting to catalog the attributes necessary to demonstrate intelligence, in a famous 1950 paper titled “Computing Machinery and Intelligence,” the computing pioneer Alan Turing suggested a test he called the “imitation game” that could be used to answer the question “Can machines think?”³² The proposed test would have a human interrogator send messages remotely to both a human and a computer (not knowing which was which), and get responses back. If the interrogator could not tell the difference between the human and the machine, Turing suggested that the computer could be said to be thinking. This test has since become known as the “Turing test.” Turing’s prediction was that the test he proposed would be passed by about the year 2000. A recent computer program which simulated a 13-year-old Ukrainian boy purportedly passed the the Turing test in 2014, but it passed by only a slim margin.³³

There have been many challenges to the presuppositions inherent in the Turing test. The validity of the Turing test rests on a philosophical notion called *functionalism* which suggests that mental states can be reduced to mapping sensory inputs to behavioral outputs. The philosopher John Searle challenges this idea by describing a thought experiment called the Chinese room experiment.³⁴ Searle first considers a computer that passes the Turing test, but he adds the stipulation that the test is conducted in Chinese. He now conducts a second experiment, in which a person who only understands English is placed in a room. Messages in Chinese are passed into the room. Following an elaborate set of English instructions which are based on the program of the computer which passed the Chinese Turing test, the English-speaking person matches the Chinese symbols with instructions on how to return a message by arranging other symbols. To a person fluent in Chinese on the outside, it would appear that the English-speaking person in the room understands Chinese. However, as this experiment demonstrates, the English-speaking person in the room does not have any understanding of Chinese. Searle argues that the computer is comparable in that it is essentially a symbol processing machine and it cannot be said to think. He distinguishes between “weak AI” and “strong AI.” Strong AI holds that machines running the right software could become a mind, whereas weak AI holds that machines can only simulate thinking.

If a machine can play chess better than the very best human player, does that make it intelligent? Searle would claim no, that it is merely the human programmers that are intelligent, who have programmed the machine to implement their ideas. However, how is this different than a human mentor that teaches a student to play chess? Do we say the mentor is intelligent and the student is merely deterministically following the rules she was taught? We now turn to the Bible for further insight into that question, looking beyond our self-identified attributes and focusing rather on our purpose, as determined by our Creator.

4. Human Identity in Scripture

The previous section explored a few approaches that philosophers and sometimes scientists have taken in distinguishing what makes us human, or what constitutes thinking. Christians should find these questions particularly important, since our identity as humans should be closely tied to

our created purpose. In this section, we thus explore some scriptural concepts and references that help us understand our identity as humans and the role of technology. We will explore these concepts in categories corresponding to the three part narrative of creation, fall, and redemption.

4.1. Created Human

Genesis 1:27 recounts the creation of man, and in comparison to all the prior acts of creation, we are unique: the only creatures God created that he also endowed with his image. As the *imago Dei*, we reflect the creator. The Genesis passage offers scant explanation for what it means to be created in God's likeness, though the immediately succeeding sentence indicates God created humans male and female. Many have taken this adjacent statement to imply relationship is a part of that reflection. Richard Mouw suggests as much, observing that "we are social beings because God created us with deep communal longings and needs."³⁵ He goes further to suggest that the image does not find its fullness in any single individual human, but only in "rich diversity of humankind spread over many places and times."³⁶ Like Mouw, Gunton also argues that part of our essential being is our placement in relation: "[that] the Platonic view is with us still in deep-seated assumptions of our culture is shown, for example, by the widespread belief that if a computer could be made to think, it would be a kind of person, as if relationality and especially love were not also essentials of our being."³⁷ These twin ideas of relationship and love are concisely seen in our Lord's instructions to his disciples as he is preparing them for his impending death, commanding them to "love one another" (John 15:17). Perhaps we are the most human when we reflect God's likeness in these twin ideas.

Another reflection of God's likeness is our ability to create. Although all humans have creativity to some extent, the Bible tells of an extra dose of creativity that God gives Bezalel, the craftsman tasked with building the tabernacle.

Then the Lord spoke to Moses. He said, "I have chosen Bezalel, the son of Uri. Uri is the son of Hur. Bezalel is from the tribe of Judah. I have filled him with the Spirit of God. I have filled Bezalel with wisdom, with understanding, with knowledge and with all kinds of skill. He can make beautiful patterns in gold, silver and bronze. He can cut and set stones. He can work with wood. In fact, he can work in all kinds of crafts.

Exodus 31:1-5 (NIRV)

Matthew Dickerson argues that creativity is the ability to make something original and "to bring into being something new, which does not proceed entirely from what has gone before or what already exists."³⁸ He argues that machines are controlled by physical causes (predictable or unpredictable) and therefore, by this definition, cannot be creative.³⁹ Aesthetic ways of knowing cannot be simply reduced to physical processes. In contrast, a materialist view of humans would suggest that creativity is just an illusion.

Regardless of how we understand our status as created beings in the likeness of our Creator, there is a danger to defining a line too tightly around our humanity. If we define ourselves as

intelligent beings, then what does that say of people of below average intelligence? If we define ourselves as having a natural, biological birth, then what does that say of test tube babies or human clones (should they ever occur)? If we require emotion or creativity, what does that say of the person lying in a coma? Thus the danger in circumscribing our humanity too tightly is that we consider certain people as somehow less human, whether they are a fetus, a senile elderly person, a man in a coma, or a child with a severe brain injury. It would be a mistake to define some kind of litmus test for what it means to be human based on attributes such as intelligence or creativity. Instead we should affirm the clear biblical teaching that God has made humankind in his image. Our ontological status as humans seems to be distinct from the rest of the creation (including machines). This implies that human personhood needs to be attributed even when certain human attributes are less evident due to age, capacity or infirmities.



Illustration 1: "In Body and Soul" by Carina Schuurman

A materialist would reject the notion of a soul, free will, and creativity since this viewpoint sees everything as determined by natural laws. In contrast, a Christian view acknowledges the diversity as well as the irreducible complexity of reality. Geertsema concludes that “[f]or a responsible implementation of all kinds of technology it is of crucial importance that the distinct nature of human personhood be taken into account.”⁴⁰ A Christian perspective is shaped by the understanding that God created us with the ability to respond to him, and with that ability comes freedom and responsibility. Part of that responsibility is given in the cultural mandate (Genesis 1:28) where humans are called to take care of the earth and unfold all the latent possibilities in creation. Freedom and responsibility imply we have a choice. The ability to choose (and especially to make a moral decision) is perhaps the most difficult attribute to understand about ourselves. How can the creator give the thing he creates the ability to do something other than what the creator intends? And yet that is what God gave us. This ability touches on the paradox between election and free will. This ability makes us human, and perhaps more than any other ability makes us distinct from machines, but this ability also allowed us to fall.

4.2. Fallen Human

Views such as materialism elevate “one aspect of human being to be the ultimate, apart from any dependence or responsibility to God the creator.”⁴¹ Anthony Hoekema identifies this as a type of idolatry: worshiping an aspect of creation rather than the creator. Others (such as Kurzweil) place their trust in the hope that one day we will be able to download our brains into a computer and thereby achieve a kind of immortality. This is an example of *technicism*, placing our trust in technology as savior of the human condition. Fred Brooks observes that the rhetoric in the field of AI has “echoed the builders of the Tower of Babel: ‘We will make machines that think; we will make Giant Brains.’”⁴² Brooks suggests that these goals, “although glamorous and motivating, sent the discipline off in the wrong direction.”⁴³ We are responsible for the direction of our technology, also in the area of AI.

As Brooks observes, this posture is not new. The story of the Tower of Babel describes people who use technology to build a name for themselves (Genesis 11:4). Other examples in Scripture include King Uzziah. He built towers and invented devices for military use, which made him powerful. “But after Uzziah became powerful, his pride led to his downfall” (2 Chron. 26:16). The technology of AI, like many other technologies, makes the user more powerful. That power can make us proud and lead to our own downfall. In response to technicism, we affirm with the psalmist: “Some trust in chariots. Some trust in horses. But we trust in the Lord our God” (Psalm 20:7 NIRV).

Our intent is not to suggest that technology itself is inherently evil. Rather, it is corrupted by sin just like all of creation. Following ideas from Wolters⁴⁴, we believe that God created inherent structures as originally good, including marriage, government, music, math, technology, and more. However, sin affects the direction of these structures, turning them away from God. Thus, just as examples of bad marriages or governments should not lead us to conclude that marriage or government themselves are evil, neither should examples of bad technology lead us to avoid all technology.

All humans are fallen by virtue of the choice made by our first parents. Romans 8:22 tells us that not only we, but the entire creation now groans under the weight of sin. So do we pass on that sinful nature not only to our biological children, but also to our technical creations? If all creation is tainted, then yes, our technical inventions are also affected by sin. It is notable that although AI and robots are *affected* by sin, if they do not have moral agency, then they cannot *initiate* sin on their own. Being held morally responsible for an act implies freedom and choice, and AI programs simply follow a program. Thus sin becomes evident in machines when humans develop and employ them in ways that go against God’s intent for his creation. Those that hope AI will somehow enable us to surpass ourselves, producing an intelligent, yet sinless creature, are mistaken. The whole creation is groaning, including our machines, and no human effort can erase that taint. Our redemption does not lie in our technology.

4.3. Redeemed Human

As Christians, we believe that God created all things good, including humans. We believe that with our first parents, we fell into sin, tainting all creation in the act. Only through Christ do we have redemption: it is by grace we have been saved. As Christians, we are called to be Christ's redemptive agents in the world. As engineers, we are especially called to use our creative gifts to develop redemptive technology. What would such tools look like? Such tools would enable us to fulfill our purpose better. They would aid us towards shalom. Following his work with ELIZA, Joseph Weizenbaum reflected on the appropriate role for computers. He concluded that computers ought not to be used for tasks that require wisdom.⁴⁵ Weizenbaum goes on to conclude that “there are limits to what computers ought to be put to do.”⁴⁶ Our purpose is to love God and our neighbor, to fill the earth and steward it, to act justly, love mercy, and walk humbly with our God. These are things that we ought not offload to machines. Why? Regardless of the question of whether machines could actually do these things, humans ought not delegate those tasks that form our very purpose. Tools that aid us in our purpose are commendable, but tools that purportedly perform our purpose instead of us are condemnable. Imagine inventing a machine that rather than helping us pray or worship, instead did our praying or worshiping for us, so that we no longer felt the need to do it ourselves. Such a machine would be completely misguided, and the users of such a machine would be truly deluded. For the remainder of this section we thus explore how AI could aid (but not replace) us in pursuing our purpose. Micah 6:8 (NIV) can be a helpful guide for engineers to ensure our project requirements conform with God's requirements:

He has shown you, O mortal, what is good.
And what does the Lord require of you?
To act justly and to love mercy
and to walk humbly with your God.

We could use AI technology to help us seek *justice* by enabling an attorney to help less fortunate members of society at a reasonable cost, by using an expert AI system as a first contact “help desk.” However, it would be important that the attorney does not simply sit back and let the expert system provide the only advice to the client. Rather, the attorney ought to use the expert system as an assistant to do a first interview, so that her or his in-person follow-up meeting with clients is more effective. Such combination of humans and machines could help the rural poor get at least initial legal advice remotely over the Internet or phone. A negative example related to justice would be handing over life and death decisions on the battlefield to AI programs. One principle of just war is that someone should be held justly responsible for any deaths that occur, an ethical requirement that cannot be offloaded to machines.⁴⁷ We have faced the moral question of machines replacing humans to do physical labor for many years. Some might argue that manual labor is drudgery and the machines free us to do more creative work. Now with AI, we have a new variation of this dilemma as expert systems are developed in order to replace experts such as doctors, lawyers, or perhaps even engineers. But does freeing us from labour free us from the very activity that makes us human? Work is not a result of the fall; both manual work

and knowledge work are a legitimate part of our creational calling in this world. In his book *The Glass Cage*, Nicholas Carr provides a nuanced discussion on the many effects of automation illustrating that it is an ethical choice since it shapes our lives and our place in the world.⁴⁸

We could use AI technology to help us love *mercy*. For example, AI image processing systems already exist today that can detect certain types of breast cancer in mammogram images better than human doctors can alone.⁴⁹ As another example of mercy, we could enable caring by providing a first contact for call centers with AI natural voice recognition so that trivial tasks could be completed routinely, while ensuring a human operator smoothly steps in for more creative and service-oriented needs. Home automation systems could enable the elderly to maintain independence longer by assisting them with everyday tasks such as cleaning and meal preparations. However, it is important that such systems do not entirely replace humans. For instance, the design of AI programs and robots should recognize social norms and not be employed to replace human care and companionship. Sherry Turkle observes that any relationship with a robot is a relationship only about one person.⁵⁰

We could use AI technology with *humility* by recognizing our own human limitations. If we are uncertain of the status of our AI creations, then in humility, perhaps we should avoid such pursuits. That is, perhaps there is a line beyond which we may develop technologies for which we no longer fully comprehend the implications. We have a long history of letting the genie out of the bottle, and we know you can never put him back in. This is the risk illustrated in the tale of the “sorcerer's apprentice.” However, even if some or even most agreed to be prudent with research and development, a few might continue these developments. As we consider the AI tools we build, we need to keep in mind our purpose. All technology is utilitarian: we develop tools as means towards ends. But technology has a bias, and this bias shapes us as we use our tools. How can AI help us fulfill our purpose without the means distorting our ends? For one thing, we should not aim to develop thinking machines that replace us, but rather to develop thinking machines that aid us in thinking ourselves, that augment and extend our abilities. On the topic of AI, Fred Brooks suggest that we should explore intelligence amplifying software to work together with humans rather than focusing on building “giant brains.”⁵¹ In a paper on the benefits and risks of AI, the authors conclude that “Some of the most exciting opportunities ahead for AI bring together the complementary talents of people and computing systems.”⁵²

As a final thought regarding redemption as humans, we look forward to restoration. In contrast, consider that Ray Kurzweil looks forward to a day when we can download our brains into computers and thus shed the fragile “hardware” of our brains and bodies.⁵³ Kurzweil goes on to quote Yeats who reflects on our physical self as “but a paltry thing, a tattered coat upon a stick.” We reject this thesis – the biblical story affirms the value of our physical bodies and the physical world. Although our body and soul are temporarily separated at death, God’s redemption plan involves a new heaven and new earth and a bodily resurrection for believers.⁵⁴ We will not be disembodied souls floating in the ether, but rather we look forward to a day when God will renew his creation, which will include both our souls and bodies.

5. Conclusion

The computer scientist Edsger Dijkstra once wrote that the question of whether machines can think is about as relevant as the question of whether “submarines can swim.” Regardless of where one stands on this question, it is clear that AI raises many fundamental questions about what it means to be human. These questions include issues of philosophical anthropology and the notion of the body, soul, and mind. It includes questions about what makes us uniquely human and whether a machine could ever replicate that. Many attributes are associated with being human such as intelligence, emotion, creativity and free will. Views that suggest computers can completely replicate humans are largely based on a materialistic view of humans. We have argued that the implications of materialism lead to a denial of many of these attributes, such as free will, creativity, and the soul. In fact, materialism can lead to a rejection of the body as people seek to shed their mortal bodies and look forward to downloading their brains into virtual environments.

Instead, we have described a view of what it means to be human shaped by the story of Scripture. The creation story describes who we are as image bearers of God who have been granted freedom and responsibility. We acknowledge the fall into sin, which has led to many distortions in the use and application of technology, including AI. In fact, some make an idol out of AI, putting their faith in it as an eventual pathway to immortality. Thankfully there is hope through Christ’s redemption, and we are called to participate in bringing renewal, also in the work of AI and technology.

The question of how should Christians think about thinking machines is not just an academic exercise nor is it just fodder for science fiction movies. This question leads to fundamental beliefs about what it means to be human. As we better understand a biblical view of ourselves we will also better understand our relationship to our machines and technology. We are called to use AI in responsible ways that lead to human flourishing and exercising humility to avoid possible harm. We also recognize that there are some things for which AI ought not to be used and which may require limits. Differing philosophical presuppositions lead to very different conclusions about the place and use of AI. These technologies are not neutral, not only in the presuppositions behind them, but also in their increasing impact on our work, our culture and our world. In humility and in recognition of our fallen state, we should aim to develop tools that ameliorate the effects of sin, that enhance justice, that show mercy. In short, our tools should aid us in working as redemptive agents.

6. Endnotes

¹ Marshall McLuhan, *Understanding Media: The Extensions of Man*, The MIT Press, 1964, p. 7.

² Henk Geertsema, “Cyborg: Myth or Reality?”, *Zygon*, 41, 2006, p. 292.

³ This question is not unlike the “Theseus’ paradox” - a thought experiment which asks whether an object remains the same object if all of its individual components are replaced.

⁴ Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation*, W. H. Freeman, 1976, pp. 3-6.

-
- ⁵ Ernest Davis. "The Singularity and the State of the Art in Artificial Intelligence: The technological singularity", *Ubiquity* (November 2014).
- ⁶ *ibid*, pp.3-4.
- ⁷ Ackerman, E., "A better test than turing", *IEEE Spectrum*, Vol. 51, Issue 10, October 2014, pp. 20-21.
- ⁸ W. Daniel Hillis, "Intelligence as an Emergent Behavior or, The Songs of Eden," *Daedalus*, vol. 117, no. 1, Winter 1988, pp. 175-189.
- ⁹ Ray Kurzweil. *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, Penguin, 2000. p. 2.
- ¹⁰ Tom Dietterich and Eric Horvitz, "Benefits and Risks of Artificial Intelligence".
<https://medium.com/@tdietterich/benefits-and-risks-of-artificial-intelligence-460d288cccf3>
- ¹¹ Bill Joy, "Why the future doesn't need us," *Wired*, April 2000.
- ¹² Rory Cellan-Jones, "Stephen Hawking warns artificial intelligence could end mankind," BBC News, 2 Dec 2014, <http://www.bbc.com/news/technology-30290540>
- ¹³ Samuel Gibbs, "Elon Musk: artificial intelligence is our biggest existential threat," *The Guardian*, 27 Oct 2014, <http://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat>.
- ¹⁴ Matthew Dickerson, *The Mind and the Machine: What It Means to Be Human and Why It Matters*, Brazos Press, 2011, pp xvi-xvii.
- ¹⁵ Note that the term dualism has different meanings in different contexts. In this context we are referring to the general notion that human beings are somehow comprised of two parts.
- ¹⁶ John W. Cooper. *Body, Soul, & Life Everlasting*, Eerdmans Publishing, 1989, pp. 17-18.
- ¹⁷ *ibid*, p. 120.
- ¹⁸ *ibid.*, p. 9.
- ¹⁹ *ibid*. pp. 107-108.
- ²⁰ Gilbert Ryle, *The Concept of Mind*, Hutchinson, 1949, p. x.
- ²¹ Douglas R. Hofstadter. *Godel, Escher, Bach: an Eternal Golden Braid*, Basic Books, 1979, pp. 315-316.
- ²² Hillis, p. 175.
- ²³ Dickerson, pp. xxi-xxii.
- ²⁴ Dickerson, Chapters 2-4.
- ²⁵ It should be noted that physicalism does not necessarily imply determinism. Some physicalists suggest that quantum effects could account for randomness which cannot be predetermined.
- ²⁶ Dickerson, p. 27.
- ²⁷ Kurzweil, p. 57.
- ²⁸ *ibid.*, p. 58.
- ²⁹ Dickerson, pp. 58-59.
- ³⁰ C.S. Lewis, *Miracles*, Harper Collins, p. 27.
- ³¹ Kurzweil, p. 142.
- ³² Turing, Alan. "Computing Machinery and Intelligence." *Mind*, 59, 1950, pp. 433-60.
- ³³ "Computer AI passes Turing test in 'world first'", BBC, June 9, 2014. <http://www.bbc.com/news/technology-27762088>
- ³⁴ Searle, John R. "Minds, Brains and Programs." *Behavioral and Brain Sciences*, 3, no. 3, 1980, pp. 417-57.
- ³⁵ Richard J. Mouw, "The Imago Dei and Philosophical Anthropology", *Christian Scholars Review*, XLI:3, Spring 2012, p. 259.
- ³⁶ Mouw, p. 265.
- ³⁷ Colin E. Gunton, *The One, the Three and the Many: God, Creation and the Culture of Modernity*, Cambridge University Press, 1993, p. 60.
- ³⁸ Dickerson, pp. 43-44.
- ³⁹ *ibid.*, p. 56.
- ⁴⁰ Henk Geertsema, "Cyborg: Myth or Reality?", *Zygon*, 41, 2006, p. 324.
- ⁴¹ Anthony A. Hoekema, *Created in God's Image*, Eerdmans, 1986, p. 4.
- ⁴² Frederick P. Brooks, "The Computer Scientist as Toolsmith II," *Communications of the ACM*, vol. 39, no. 3, March 1996, p. 63.
- ⁴³ *ibid.*, p. 64.
- ⁴⁴ Albert M. Wolters, *Creation Regained: Biblical Basics for a Reformational Worldview*, Eerdmans, 1985.
- ⁴⁵ Weizenbaum, pp. 8, 227.

⁴⁶ *ibid.* p. 11.

⁴⁷ Noel Sharkey, "Automated Killers and the Computing Profession," *Computer*, vol. 40, no. 11, 2007, p. 122.

⁴⁸ *ibid.* p.17.

⁴⁹ <http://www.breastcancer.org/research-news/20130425-4>

⁵⁰ Sherry Turkle, *Alone Together*, Basic Books, 2011, p. 56.

⁵¹ Brooks, p. 64.

⁵² Dietterich.

⁵³ Kurweil, p. 129.

⁵⁴ For a good discussion of this, see: J. Richard Middleton, *A New Heaven and a New Earth*, Baker Academic, 2014.