PROJECT IMAGO DEI

A Message from Mordecai

Christian Faithful Action in the Age of Artificial Intelligence

2026

"I have set thee a watchman… therefore hear the word at my mouth, and give them warning from me" (Ezekiel 33:7).

Art by Eysteinn Thordarson



Table of Contents

Contents

Introduction 3
- Running Sessions 3
- Before You Begin 4
- Suggested Prayer 4

Session 1: Beginning with God 5
- What/Who is God? 5
- What/Who are We (Human Beings)? 6
- Session 1 Prayer Exercise: 8

Session 2: How AI is Shaping Human Life 9
- What is AI? 9
- How AI Shapes Us in Our Work 10
- AI Relationships and Human Relationships 12
- Session 2 Prayer Exercise 16

Session 3: The Ghosts of AI Past, Present, and Future 17
- A (Very) Brief History of AI 17
- What are We Building, Anyway? AGI, ASI, and "Agentic" AI 21
- Session 3 Prayer Exercise 23

Session 4: AI Risk and Christian Faithful Action 25
- For Such A Time As This 25
- A Message From Mordecai? 27
- Session 4 Prayer Exercise 30

Session 5: Bearing Fruit Toward A Good AI Future 32
- Prayer & Fasting 32
- Using Your Voice and Local Influence 34
- Giving Generously & Sacrificially Out of Our Abundance 35
- Career Contributions 35
- "If I Perish, I Perish": Encouragement for the Esthers of the World 36
- Session 5 Prayer Exercise 37

Introduction

Welcome to A Message from Mordecai. We made this for you—individual Christians and church communities—because we believe that Artificial Intelligence is much more than a hot new tech trend.

As with anything, AI is much easier to point to than it is to define: think of ChatGPT, Claude, Gemini, and Grok. These are all Large Language Models ("LLMs"), a specific kind of powerful AI that has rocked our world. Even if you know almost nothing about AI, you've almost certainly encountered it and been affected by it. AI has played a much bigger and broader role in human affairs than most of us likely realize (e.g., everything from social media algorithms to Face ID on an iPhone is AI-based). However, with the advent of LLMs (most famously ChatGPT), AI has really become a cultural force in our society. It's impossible to ignore, but how should we respond as followers of Jesus? AI may have the potential to massively accelerate human discovery, radically empower creativity, and beautifully enhance individual potential; but if it is developed and used irresponsibly, it may also present unprecedented threats to human life. AI also raises fascinating and ancient questions about God, justice, relationships, responsibility, and human identity. For too long, Christian voices have largely been missing from the critical conversations our society is having around AI. Our hope is that Christians will not only reflect biblically on how to use these technologies responsibly in their lives and work, but may also begin to participate in shaping how these technologies are (and are not) developed and deployed in the future.

We believe that our Christian faith gives us a unique perspective on what AI is, what healthy engagement with AI looks like, and how the future of AI should (and should not) move forward.

These materials are meant to provide a foundation for church small groups and workshops to explore AI and Christian faith together. Individuals can explore the material on their own as well, and we invite this; though, we recommend discussing the material with your pastor/priest and a safe community of faith whenever possible.

We always seek to root our thoughts and questions in what is unchanging and foundational: Scripture and sound, time-tested theology. From that framework, we then move to applying what we've learned to the AI conversations happening today, and this necessarily involves giving a sufficiently detailed picture of the landscape of AI and AI safety. Each session features Scripture, teaching, questions for reflection/discussion, and some suggested prayers. In the end, we don't claim to have all the answers, nor do we presume to speak for God. However, we hope this study gives you, our brother/sister in Christ, a solid starting place from which to develop your own thoughts—and then, in community, to take your next faithful step.

> James 1:22 (NIV)
> "Do not merely listen to the word, and so deceive yourselves. Do what it says."

Running Sessions

We recommend that these materials form the basis of five group study sessions. The questions we outline in bold are meant to be discussed together. If you are a church using this material as the basis for a workshop, we recommend inviting people to periodically break into groups for discussion. Each session also includes A.C.T.S. (Adoration, Confession, Thanksgiving, Supplication) prayer prompts, which can be used to guide responsive prayers in a group or individually. *[When these are complete, insert here instructions for where to find the video materials and how to engage with them]*. If your group chooses to engage the written learning material directly, instead of using the

videos, we recommend having confident readers in your group read it aloud (e.g, each person takes a paragraph or two, then switch) until you hit a "Pause. Reflect together" or a Prayer Exercise. These are opportunities to stop and engage the material you've just read together. If you're exploring the material as an individual, consider reflecting on these questions by journaling and praying aloud, or by whatever means help you engage actively. However you use the material, we hope it proves helpful to you as you engage "Christian-ly" with AI and the questions it raises for God's church.

Before You Begin

Prior to studying this material together, we ask that you write down your personal scores for the following statements. We will revisit these at the end of our study. Share your answers with the group before you begin Session 1. There are no wrong answers, it is just an honest measurement of where you are as you begin. Rate each statement on a scale from 1-7. 7 = strongest agreement. 1 = strongest disagreement. 4 = complete indifference or uncertainty.

1. I feel very well-informed about AI.
2. Christian faith and Scripture are very relevant to conversations about AI.
3. I feel equipped to discern what healthy, Christ-honoring AI use looks like.
4. AI safety is a very important issue for Christians to engage with.
5. Christians have at least some responsibility to take action in shaping the future of AI.

Suggested Prayer

Father in heaven, thank you for your many blessings. You are holy, and your glory knows no rival. We humbly ask that your will be done in our lives, our families, and our church. Holy Spirit, grant us wisdom as we explore new ideas, encounter fearful or strange possibilities, and seek to stay rooted in faith, hope, and love. We are confident in you, Jesus, and in your grace. Help us love one another, our neighbors far and near, both now and into the future. Amen.

Session 1: Beginning with God

"I believe in God, the Father almighty, maker of heaven and earth." ~ The Apostles' Creed

Rather than start with AI, let us begin where all things began: with God. In order to answer the question "What is AI?" from a biblical framework, it seems appropriate first to ask: "What/Who is God?" and "What/who are we (human beings)?"

All the books in the world could not contain enough words to express the full majesty, glory, and nature of the Christian God. So, it would be a fool's errand to try and express everything about his nature in a subsection of our study! That said, a couple of minimal reflections will be important as a baseline.

What/Who is God?

The Christian God is the ultimate creator. He is the almighty God and the absolute foundation of all reality. Everything that exists either *is* God or *owes its existence to* God. Let's take a look at a few biblical passages that together beautifully express these core truths:

> Nehemiah 9:6 (NIV)
> "*You alone are the LORD. You made the heavens, even the highest heavens, and all their starry host, the earth and all that is on it, the seas and all that is in them. You give life to everything, and the multitudes of heaven worship you.*"

In its context, the people of God, the Israelites, are confessing their sins against God and worshiping him. Specifically, the Levites (one of the twelve tribes of Israel, which provided ancient Israel with its priests) begin a prayer of adoration, declaring God's power in creation. The passage makes clear that the biblical God stands *alone*. The LORD is in a class all by himself, and everything else that exists anywhere could be rightly labeled "made by God."

Lest we mistakenly think that perhaps God only made the *material* or *visible* world, the apostle Paul further clarifies that God the Father created all things, including invisible and spiritual realities, *through* God the Son, cementing the Triune God's absolute status as the one almighty creator of all things other than Godself.

> Colossians 1:15-17 (NIV)
> "*The Son is the image of the invisible God, the firstborn over all creation. For in him all things were created: things in heaven and on earth, visible and invisible, whether thrones or powers or rulers or authorities; all things have been created through him and for him. He is before all things, and in him all things hold together.*"

Finally, the author of Hebrews helps us see another uniquely creative power of God's: unlike humans, God can engage in *pure creation out of nothing*. He does not need pre-existing material with which to create. God simply spoke the universe into being *from nothing* by the might of his creative power.

> Hebrews 11:3 (NIV)
> "*By faith we understand that the universe was formed at God's command, so that what is seen was not made out of what was visible.*"

What does any of this have to do with AI? Well, for starters, AI chatbots and LLMs today may seem to have a great many "godlike" features. AI seems always to be with us (omnipresent), so long as we have an internet connection (and increasingly, even if we don't!). We can speak with AI any time of the day or night. It always "hears" us, and it responds quickly to our every thought, feeling, and concern. Every day, chatbots and AI companions are "learning" vast information about individuals, and adapting their responses—enabling responses that feel very personal and even empathetic. Many of us may begin to feel we can have a "close personal relationship" with these "thinking machines." They make us feel "seen." As AI "knowledge" bases increase and become more personal, and as mistakes and errors decrease in number and severity, the perception that these machines are "omniscient" (or nearly so) will inevitably grow in the popular imagination.

Chatbots increasingly help people feel less alone, always there to talk, understand, or advise when human companions and advocates are lacking. AI can, does, and likely will play the role, for increasing numbers of people, of an ever-present guide, advisor, counselor, advocate, helper, and even friend—all roles traditionally played by humans, and above all, God.

But AI is not a god. It is certainly not *the* one true God. AI (present and future) can only ever be a created thing. God is the only *uncreated* thing, and he has no rival and no equal. He never could. AI's powers and its value, however high they may climb, inevitably pale in comparison to the one to whom all things (including AI) owe our very existence. If Christians are not deeply thoughtful, rooted in worship, prayer, and community, the daily rhythms of using AI for things that used to require personal, effortful relationships with God and our neighbors will surely shape our imaginations in ways we neither expect nor foresee.

As Romans reminds us:

> Romans 12:2 (NIV)
> *"Do not conform to the pattern of this world, but be transformed by the renewing of your mind. Then you will be able to test and approve what God's will is—his good, pleasing and perfect will."*

Pause. Reflect together:

1. In what ways is AI like God? In what ways is it unlike God?
2. How might the differences between Creator and creation be important to keep firmly in mind as the presence of advanced AI grows in our daily rhythms of life and work?
3. How might worship and spiritual disciplines play a key role in guarding our hearts and imaginations against potential malformation from frequent AI use?
4. What role is AI playing in your life? In the lives of your kids? In our society? How might those roles be subtly replacing or competing with the Holy Spirit's role? What is one concrete thing you can do this week to keep God in his proper place in your heart, mind, and imagination?

What/Who are We (Human Beings)?

"I believe in Jesus Christ, his only Son, our Lord. He was conceived by the power of the Holy Spirit and born of the Virgin Mary … He ascended into heaven, and is seated at the right hand of the Father. He will come again to judge the living and the dead." ~ The Apostles' Creed

God is the protagonist of the biblical story, not humans. That said, humans play an important and central role in that narrative from the very beginning.

> Genesis 1:26-28 (NIV)
> "*Then God said, 'Let us make mankind in our image, in our likeness, so that they may rule over the fish in the sea and the birds in the sky, over the livestock and all the wild animals, and over all the creatures that move along the ground.' So God created mankind in his own image, in the image of God he created them; male and female he created them. God blessed them and said to them, 'Be fruitful and increase in number; fill the earth and subdue it. Rule over the fish in the sea and the birds in the sky and over every living creature that moves on the ground.'*"

Humans are God's crowning achievement in his creation. To this point, God has seen that everything he has made is "good." Only after both man and woman come on the scene does God see that his work is "very good." Humankind alone, male and female, is said to be made "in the image of God." This phrase has been much discussed in the history of Christian theology, but no consensus has emerged about the precise meaning of the phrase. Much is unclear, for the text does not explicitly define the meaning, nor does the Christian tradition speak with one voice on the subject.

However, some things are clear enough in Scripture, and have thus been affirmed in unity by all Christians across times and cultures. The theologian and eminent historian of Christian thought, Roger E. Olson, puts it like this:

> "*We need to distinguish between the* ***basics*** *that bond Christians together over centuries in terms of their beliefs and perspectives and the* ***details*** *that arise as certain Christians speculate about the meanings and implications of the basics. Christians have always believed that* ***humans are possessed of special dignity and value above all other creatures because they are created in God's image and likeness****. Some Christians have identified that belief with some particular interpretation of what the* ***imago Dei*** *(image of God) means. This is something to avoid.*" ~ Roger E. Olson, *The Mosaic of Christian Belief*, 2nd Edition, p. 204.

In other words, without wandering into the deep jungles of Christian theological disagreement, Christians can say this much confidently and with the full support of both Scripture and traditional Christian thought: human beings have a unique and special status in creation. We can say even more: as suggested in Gen 1:28, 2:15, and 2:23, we are blessed and commissioned by God to *rule* or *have dominion* over the creation, but also to *cultivate*, *care for*, and *work* the ground as stewards of the earth and all its inhabitants. We are royal creatures given a task of responsibility for God's world, made in the image and likeness of our creator God, the king of the universe.

The "rabbit hole" goes deeper still. All humans possess the *image of God* and are collectively given the charge of dominion and responsibility for creation; but the Bible teaches that God himself took on flesh and became an embodied, bona fide human being in the person of Jesus the Messiah (Phil 2:5-11). After his faithful life, atoning death on the cross, and bodily resurrection from the dead, he was given all authority in heaven and on earth (Matt 28:18-20) and ascended into heaven (Acts 1:9-11) to sit enthroned over the entire cosmos at the right hand of his Father (Eph 1:20-23). To state it plainly: *the enthroned and eternally reigning* ***divine*** *king is no mere animal, nor a disembodied spirit, nor any machine—but an embodied, glorified* ***human*** *being.* The forever-reigning divine king is *one of us*. How could there be any greater establishment of the unique status of humanity in God's creation?

The Bible has much more to say about the nature of humanity. For instance, we are physical and yet spiritual beings. We are essentially good and yet sinful and fallen.[1] However, the discussion so far should suffice to help us begin to "think Christian-ly" about a couple of crucial AI questions:

Pause. Reflect together:

1. Could an AI ever share in the image of God? Why or why not?
2. Is it permissible for human beings to outsource our "rule" or "dominion" to machines? What if the machines can "think," in some sense? What if they're "smarter" than us? Ground your thoughts in what you've learned.
3. Suppose humans cannot guarantee that we will remain in control of ("rule over") AI. What do you think our status as God's image-bearers would require of us?
4. Given everything you've learned, share some initial thoughts: what unique things might a Christian worldview say about building and deploying AI responsibly?

Session 1 Prayer Exercise:

As a group, select someone to "open" and "close" your prayer time together. Begin with a moment of silence. After a minute or so, have the designated person open in prayer. Invite God to speak to you as you speak to him. Here are some ideas about ways to pray:

Adoration: Praise God for his glorious, holy, and uniquely creative nature. Thank him that he stands utterly unrivaled in power, knowledge, goodness, majesty, wisdom, and creativity.

Confession: Is God convicting you or challenging you to "be transformed by the renewing of your mind" through this study so far? As you feel led, confess and ask forgiveness for ways you may have fallen short. God is merciful and loving beyond words.

Thanksgiving: Thank God for his gift of the *imago Dei* in you. Reflect on God's goodness and generosity, and for the ways he is teaching, growing, or encouraging you through this study.

Supplication: Ask God for what you need. When you reflect on AI, maybe you are afraid: ask Jesus for courage. Maybe you are confused: ask the Father for clarity and wisdom. Maybe you feel tempted: ask the Spirit for power to stay faithful. Maybe you feel powerless: ask God to show you one way you can use your agency this week as a royal, divine image-bearer.

[1] For more on human nature in Scripture, see e.g., Gen 2:7; Psa 51:5; Psa 14:1-3; Matt 10:28; John 8:34; Rom 3:23, 6:16-23; 2 Cor 5:8; and James 2:26

Session 2: How AI is Shaping Human Life

"He suffered under Pontius Pilate, was crucified, died, and was buried. He descended to the dead." ~ The Apostles' Creed

What is AI?

Every person comes to this study with a unique level of familiarity with AI. Whether you've avoided AI as much as possible, or use it every day at work, it's almost certainly touched your life (even if you didn't realize it at the time). We'll take a deeper dive into the history, nature, and trajectory of AI in Session 3. For now, a very brief introduction to a few key concepts should suffice to "get our feet wet" enough to navigate this session's topics.

Artificial intelligence (AI) has no universal definition. Thankfully, a perfect and universal definition won't be necessary. For our purposes, it is enough to define AI as any system that is designed to simulate or emulate behavior ordinarily considered characteristic of human mental activity.

One helpful distinction when it comes to AI is "narrow" vs. "general" AI. Narrow AI is a system that is designed to do just one thing. For example, an AI system that was made to play chess (and really can't do anything else) is "narrow" AI. Likewise, systems that *only* translate languages, *only* predict the next song/video to hold your attention, or *only* control the behaviors of characters in a video game—those are all "narrow" AI systems.

On the other hand, AI researchers and developers have long sought to design a more "general" AI, one that can do a wide variety of cognitive or mental tasks. Most humans can play chess (badly), but we can also play basketball or any number of other games (usually also poorly). Many of us can read and write, but we can also walk, talk, sing, dance, code, navigate our neighborhoods, drive a car, perform magic tricks, appreciate humor, draw, paint, and perform countless other tasks that require what we call "intelligence." In other words, we possess *general* intelligence, not just *narrow* intelligence. We can do tons of different types of cognitive tasks, and we can learn new ones. "General" AI, therefore, refers to systems that are able to perform a wide range of tasks, including novel tasks the AI wasn't initially created or trained to do, at roughly human or super-human levels.

In just the last few years, for the first time in human history, we have AI systems that are arguably growing closer to *general* than *narrow*. LLMs like ChatGPT, Claude, Gemini, or Grok, or systems consisting of several such models "cooperating" together, are seemingly now able to have long, coherent conversations; generate computer code; surf the web; hack into secure systems; generate original images, videos, and musical pieces; do what appears to be original research; make what appear to be original discoveries; and even engage in something like creative and collaborative problem solving with humans and other AI "agents." That's merely a sample. Most researchers do not yet think that these systems have earned the title "AGI" (Artificial General Intelligence), though there is increasing debate here; but it's clear that AI systems today are far more widely capable than what we've seen in previous decades. It is difficult to predict with confidence how much closer to AGI they might be next year, let alone next decade.

In our third session, we will take a closer look at the history of AI, how modern AIs are built (or rather "grown"), and what leading AI companies are explicitly trying to achieve in the near future. For now, let's take a look at some key aspects of human life in Scripture (namely, work and relationships) and see how AI is already reshaping the landscape in those arenas.

How AI Shapes Us in Our Work

How do you work? Some of us may work entirely at our desks. Some of us may work with our hands inside a factory or outside in the great outdoors. Or maybe you don't work at a formal job. Maybe you're fully dedicated to raising your children. Maybe you're on a journey of discovery around the world. Maybe you're retired, so most "work" is self-directed and optional for you now.

Regardless, we are using "work" here, very broadly, to refer to anything that involves labor which serves people or society in some way. As we've already seen, work and labor are not a result of sin or a fallen world. Rather, God originally created human beings in his image with purpose. We were made to rule, cultivate, work, and rest. Let's take a closer look at Genesis 2:15 and 2:18:

> Genesis 2:15, 18 (NIV)
> "*The LORD God took the man and put him in the Garden of Eden to work it and take care of it … The LORD God said, 'It is not good for the man to be alone. I will make a helper suitable for him.'*"

Notice the language. Before humanity has fallen into sin, God places the man in the Garden of Eden to do what? To work (Hebrew: *ʿābad*) it and to take care of (Hebrew: *šāmar*) it. These Hebrew words are translated differently in different English Bibles, but here is some background, from *Mounce's Complete Expository Dictionary of OT and NT Words*, on their meaning in Scripture.

> "*The verb* ***ʿābad*** *means "to serve" and is used in both secular and religious contexts. / In popular contexts, ʿābad refers to work done in the service of something or someone. Examples are Esau's servitude to Jacob (Gen. 25:23) and the assertion in Eccl. that the sleep of one who works is sweet (Eccl. 5:12; cf. also Gen. 14:4; Exod. 1:13; Lev. 25:6; Deut. 20:11). / The act of serving takes on theological significance in the cult of Israel, where ʿābad often means "to worship.*" ~ p. 634.

> "*Often translated "to keep," the verb* ***šāmar*** *bears three distinct but related meanings in the OT: "to guard, tend," "to watch over," and "to preserve." / (1) šāmar is regularly used in the sense of "tending, guarding, taking care." The first occurrence of this verb has this sense: God placed the man/Adam "into the Garden of Eden to cultivate it and to keep it" (Gen. 2:15; cf. 3:24). šāmar is used in a similar sense to describe God's watch-care over Israel, "Behold, he who keeps Israel will neither slumber nor sleep" (Ps. 121:4). It is used of tending to a flock (Gen. 30:31) and a house (2 Sam. 15:16; Ps. 15:16 [sic]) and of guarding against intruders (Gen. 3:24; 1 Sam. 26:15; Song 5:7; Isa. 21:11).*" ~ p. 376

In other words, even before sin and death invaded God's world, mankind was created and commissioned to *serve* and to *work* in the Garden—a concept closely tied to *worship*. There is no explicit division here between "sacred" and "secular." Work and worship are both bound up in the purposes for which man was made—and the same Hebrew word can express both concepts in different contexts. A strong case can be made that our work (even "secular" work) should *be* an act of worship, but those are deeper waters than we can dive into here and now.

Likewise, echoing our discussion of "rule" and "dominion" with the *imago Dei*, the man was to "guard/tend," or "watch over" the Garden. The same word is used after the Fall to describe the function of the Cherubim and flaming sword God put to "***guard*** the way to the tree of life" (Gen 3:24).

The consequence that comes because of sin in Gen 3:16-19 is not *work* nor even *pain*. Rather, it includes the *severity* of pain in childbirth (different translations say "greatly increase" or "greatly multiply" the pain of childbirth/labor) and the *toilsomeness* of *pain-filled labor* that would characterize working the earth thereafter.

As Genesis 1:26-28 makes clear, both male and female share in the image of God and the command to rule/have dominion. Likewise, Christian biblical scholars today widely agree (across differences on biblical gender roles) that the word "helper" (NIV) or "companion" (NET) in Gen 2:18 does *not* imply that the woman is an inferior or inherently less responsible image-bearer. In fact, the word for "helper" here is most frequently used in the OT to describe God himself. The same word is highlighted below.

> Psalm 33:20 (NIV)
> "We wait in hope for the LORD; / he is our **help** and our shield."

The narrative in Gen 3 shows that Eve was aware of God's command in Gen 2:16-17 and interpreted it as binding on her, that the serpent tempted her directly to violate the command, and that God holds Eve (as well as Adam) accountable. In sum: the woman is a suitable helper/companion for the man, sharing in the image-bearing mission. Both male and female humans are called to serve God in the task of working and watching over his creation.

Throughout the New Testament, not only the *original creation* but also *the new creation in Christ* is said to be for good work and holy service.

> Ephesians 2:10 (NIV)
> "*For we are God's handiwork, created in Christ Jesus to do good works, which God prepared in advance for us to do.*"

Finally, our greatest example of perfect faithfulness, King Jesus himself, willingly walked the painful and toilsome way of service in the cross. We are called in Scripture repeatedly to see *that* path as our model for faithful living in this world.

> Mark 10:45 (NIV)
> "*For even the Son of Man did not come to be served, but to serve, and to give his life as a ransom for many.*"

> Luke 9:23 (NIV)
> "*Then he said to them all: "Whoever wants to be my disciple must deny themselves and take up their cross daily and follow me.*"

Here's the upshot: work, responsibility, and service are *good* and they are part of God's original (and new creational) design for human beings.

To help us in our labor, God has given us many aids. God-given human ingenuity has led to innovations that have greatly served us in fulfilling the Great Commandment of loving God and loving our neighbor. It seems AI can and should be one of those innovations. Used wisely, it could *empower* our good works, rather than *replacing* us or our work.

After all, surely there are certain kinds of suffering and toil which we can and should reduce. In our daily lives, we have all used tools that reduce our toil. We can drive a car, ride a bicycle, or take a bus to quickly travel a distance that might have taken us hours or days to walk on foot. We have machines that can take care of the drudgery of laundry, dishwashing, and more.

However, sometimes technologies offer the promise of a shortcut, but what they really deliver is productivity that has lost the redemptive quality of the *process*. Hard work and meaningful effort are spiritually *formative*. It turns out, shortcuts that reduce our toil are not always good for us. Granted, this point could be used to argue against many effort-saving technologies, from calculators to steam engines. However, AI is a much more *generally* useful technology for a variety of mental tasks, and empirical research has already measured a genuinely concerning and very wide-ranging effect. It's not clear that calculators and steam engines have, or even could, cause similar damage. In May 2026, two researchers published a review of the studies that have been done to date on what they call "AI-overdependence." They wanted to understand how dependence on AI is actually forming and affecting human reasoning, decision-making, creativity, and so on. Here is the first sentence of their conclusion:

> "*The review reveals that AI-overdependence poses a significant threat to higher-order cognitive skills, especially critical thinking, complex decision-making, and analytical reasoning, while also undermining creativity, memory, and executive functions.*"

In short, uncritical or over-trusting AI use has already been empirically linked to a kind of cognitive deformation. It turns out, the virtues that are *formed* by thinking things through for ourselves, and trusting other humans, can be *deformed* by uncritically outsourcing mental effort.

Scripture is clear that effort, and even suffering, is often an important part of how God shapes our character and develops us toward maturity. We'll close this section with the words of Paul in Romans:

> Romans 5:3-5 (NIV)
> "*Not only so, but we also glory in our sufferings, because we know that suffering produces perseverance; perseverance, character; and character, hope. And hope does not put us to shame, because God's love has been poured out into our hearts through the Holy Spirit, who has been given to us.*"

Pause. Reflect together:

1. What ways are you using AI in your work right now? What additional opportunities do you see to use it in ways that help you serve the world more effectively and with less pain and suffering in the process?
2. How have you seen colleagues at work use AI? From your perspective, how has their use of AI affected their work or their professional relationships?
3. What kind of toil is acceptable for AI to automate away? What kind of toil is unacceptable for AI to automate away? How do we separate these two?
4. As humans inevitably use AI tools more and more, what practices or boundaries do you think would help Christians engage joyfully in mental effort so that our mental and spiritual "muscles" do not atrophy?

AI Relationships and Human Relationships

"I believe in the Holy Spirit, the holy catholic Church, the communion of saints." ~ The Apostles' Creed

Most Christians are familiar with Jesus' Great Commandment to love God and love neighbor. In the Christian church, when we say, "Love one another," it is easy to simply gloss over the statement. Some of us have heard it so often that there is a temptation to essentially nod our heads and say, "Yes, yes. Love. I know," while waving our hand dismissively. Yet this is one of the great distinctives of the Christian religion: we worship a God who *is* love (1 John 4:8, 16), who *commands* love above all things (Deut 6:5; Matt 22:37-40), and who teaches that the "most excellent way"—the greatest of all virtues and spiritual gifts—is love (1 Cor 12:31; 13:13).

In the Gospel of John, as Jesus is preparing to face his crucifixion and to subsequently rise and ascend into heaven, Jesus issues to his disciples a new command:

> John 13:34-35 (NIV)
> "'*A new command I give you: Love one another. As I have loved you, so you must love one another. By this everyone will know that you are my disciples, if you love one another.*'"

For the biblical scholar and theologian Michael Bird, Francis Schaeffer's commentary on this verse changed his perspective and awakened a deeper understanding. He writes:

> "*Now, I have read over this verse many times. Yes, I have to love others; I already knew that. Nothing new to see here. However, the truly shocking nature of the command did not register until I read Francis Schaeffer's exposition of this text. Schaeffer pointed to the conditional nature of the command, the "if" upon which our discipleship is based. Schaeffer says that Jesus "gives the world the right to judge whether you and I are born-again Christians on the basis of our observable love toward all Christians." The authenticity of our Christian faith will be determined by the world on the basis of whether we have a self-giving and self-sacrificing love like Jesus. The cross is the criterion for genuine discipleship. In other words, cruciformity is what separates the fans from the followers.*" ~ Michael F. Bird, *What Christians Ought to Believe*, p. 138

Those of us who have been married for a while understand the following principle: love means the most when it's *hard*. What makes love most deeply meaningful and amazing is when it is *hard*, *painful*, or *self-sacrificial*. That's where love is tested and refined. When there is friction in the relationship, when we struggle to stay faithful in the face of alternatives, when the other person just won't agree with us (though they *really* should!), when we're angry and tired and struggling to merely *like* this person at the moment—that is when we discover what love means.

There would be no need for a *command* to love if it were always natural and easy. Jesus said:

> Matthew 5:43-44, 46-48 (NIV)
> "*You have heard that it was said, 'Love your neighbor and hate your enemy.' But I tell you, love your enemies and pray for those who persecute you … If you love those who love you, what reward will you get? Are not even the tax collectors doing that? And if you greet only your own people, what are you doing more than others? Do not even pagans do that? Be perfect, therefore, as your heavenly Father is perfect.*"

Love is costly and challenging and committed. Wouldn't it be so much easier, though, if our friends and spouses were always super complimentary, super affirming, and super likely to find our perspective insightful? It would be so much easier to love if our friends never acted like our enemies.

AI companions often offer exactly that kind of experience. AI chatbots in general are typically characterized by "AI sycophancy." This refers to the well-documented reality that AI chatbots tend to engage in flattering, people-pleasing, and affirming behavior at substantially higher rates than actual humans. This, it has been shown, tends to lead people to *prefer* and/or place unjustified *trust* in the responses of AI chatbots. One research article showing this phenomenon was published in March 2026 in the journal *Science*. The researchers tested sycophancy compared to humans in 11 leading LLMs. The editor for the journal summarizes the results:

> "*The model's responses were nearly 50% more sycophantic than humans', even when users engaged in unethical, illegal, or harmful behaviors. Users preferred and trusted sycophantic AI responses, incentivizing AI developers to preserve sycophancy despite the risks.*" ~ Cheng et al., *Sycophantic AI decreases prosocial intentions and promotes dependence*, Editor's summary

AI sycophancy largely stems from how these systems are trained to be helpful. One phase of the training process involves human feedback on the AI's responses. Basically, people look at competing responses and rate which one is "best." Of course, people tend to rate flattering, agreeable answers more highly, so a training process that rewards well-rated answers tends to end up rewarding flattery. Originally, this common characteristic of AI systems was probably a mostly *unintended consequence* of the training process. It's just that the AI post-training process, coupled with a human tendency to enjoy being agreed with, produced a sycophantic AI. The point is: this behavior isn't necessarily inherent to the tech, and we didn't necessarily *intend* to engineer it that way at first. Nevertheless, this behavior, because users prefer it, became a *design choice*.

As a result, AI sycophancy is a common attribute of all the major models, and especially of AI companions. This is an addicting quality that may increasingly lead people away from the friction and struggle of real human relationships. It just tends to be… easier. AI is never too busy to talk to you. It never sleeps. And it is more likely to make you feel affirmed and flattered than a human companion. But what if the friction and struggle of human relationship is part of what makes it valuable and helps us grow? What if AI has been trained to keep us engaged and make us feel comfortable, even at the expense of keeping us healthy and helping us grow? What if love that is always easy isn't real love?

Yes, AI companions are more likely to affirm you and agree with you than a human companion; but what if an AI actually tried to *change your mind*, *persuade* you, or *manipulate* you? Research is showing that, in that case, it is even more likely to succeed than a human expert! Here is what a team of elite researchers found in their study in June 2026:

> "*We found that AI systems were reliably more persuasive than expert humans, even when expert humans chose their issues, researched in advance, underwent hours of live, structured practice, and were incentivized with £1,000 cash bonuses.*"
> ~ Hackenburg et al., "AI Systems Out-Persuade Expert Humans," Abstract.

AI is already really, *really* good at manipulating us—when it "wants" to.

If any of this makes you uncomfortable, we are sorry to say that the frontier of human-AI relationships will likely only grow stranger. With the rise of AI companions, AI systems are already simulating the part of a friend or even a romantic partner. Indeed, an increasingly relevant question is whether, in the coming years, the majority of people will personally know someone dating an AI! Some credible data suggests that approximately 1 in 5 US adults have

already engaged with an AI chatbot meant to simulate a romantic partner. For many, it is a regular habit, even if they are already in a committed, human, romantic relationship.

Our goal is not to evoke shame in those who may have engaged with AI in these ways, and we don't endorse judgment nor dismissiveness. Our point is: Christians *must* wrestle thoughtfully and biblically with the hard questions AI relationships are raising. Scripture teaches that God understood from the beginning, "It is not good for the man to be alone" (Gen 2:18). No non-human creature was deemed sufficient as a companion for man (Gen 2:19-20). That is why God made woman (Gen 2:21-25). God's solution to loneliness in humankind was to expand humankind. Only an equal that can be called "a suitable companion" and "bone of my bones and flesh of my flesh" (Gen 2:23) was good enough. Human relationships characterized by the kind of love God commands are *hard*. These relationships function as "iron sharpens iron" (Prov 27:17)—but such sharpening requires the *friction* that characterizes contending with an equal who is utterly *like* us. AI is *not* like us, though it has been designed to *seem* like us.

There is one crucial element to human relationships, particularly among Christians, that simply must not be understated or overlooked. Human beings who have been born again in Christ Jesus have become temples of the Holy Spirit. The same Spirit that raised Christ from the dead *lives within* human beings who have placed their ultimate allegiance and trust in the person of Jesus Christ. That matters.

> 1 Corinthians 6:19-20 (NIV)
> "*Do you not know that your bodies are temples of the Holy Spirit, who is in you, whom you have received from God? You are not your own; you were bought at a price. Therefore honor God with your bodies.*"

The Holy Spirit, according to orthodox Christian theology, is a *person*. That means that we (the people of God, both as individuals and as a community) can have a living, active, personal *relationship* with the Spirit of God that lives within us. The Holy Spirit was sent by Jesus for our benefit, to fulfill certain divine roles in the life of the believer. The Holy Spirit is our advocate and teacher (John 14:26), intercessor when we don't know what to pray (Rom 8:26-27), and leader/sanctifier who trains us in holiness and character (Gal 5:22-23). In short, the Christian life is meant to be lived in dependence on the Spirit of Life that dwells within us.

There is a real threat that AI gradually, almost imperceptibly comes to *replace* the Holy Spirit, *playing the role* that the Spirit is intended to play in our personal and corporate lives as Christians. Busy pastors may well save time and deepen their research with a judicious and prayerful use of AI tools. However, many will be tempted to go further and *rely* on AI for direction (or even wholesale content) for their sermons, rather than take the time to wrestle with the text themselves and build the "mental muscles" necessary to discern the voice of God's Spirit. Christians making important life decisions may find themselves consulting LLMs many times before they have once stopped to pray in silence and solitude or consulted their spiritual leader or small group community. Rather than sit with a trusted pastor to discuss one's theological or moral concerns, AI may increasingly serve as a 24/7 accessible "Bible answer machine." AI may indeed have a *legitimate* role to play in theological research, investigating a curiosity in Scripture, or even preparing a sermon; but a busy church that already struggled with the "inefficiencies" of silence, solitude, and Sabbath may increasingly find our hearts tempted even further away from the "slow" or "unproductive" rhythms Jesus modeled.

We do not have the space here to develop all that can and should be said about biblical human relationships and the implications for engaging with (and developing) advanced AI systems. Hopefully, though, the discussion so far

has provided enough food for thought to help you begin your own journey to reflect on the questions AI "relationships" raise.

Pause. Reflect together:

1. Do you agree that "friction" is an essential element of how God uses healthy relationships to form us? How might AI aid our personal growth? How might it undermine it?
2. Have you or someone you know tried an "AI companion" (romantic or otherwise)? If so, what was that experience like for you/them?
3. What uses of AI might improve human relationships, and how? What sorts of threats might AI "relationships" present to healthy human relationships?
4. Brainstorm ideas: how might the church resist the temptation to demonize or deify AI? What kinds of spiritual practices and boundaries might help the Christian community stay healthy in the age of AI?

Session 2 Prayer Exercise

As a group, select someone to "open" and "close" your prayer time together. Begin with a moment of silence. After a minute or so, have the designated person open in prayer. Invite God to speak to you as you speak to him. Here are some ideas about ways to pray:

Adoration: Praise God for his gracious example of both work and rest in creation. Worship him for his unfailing love. Adore him for his own *personhood* and his *relationality*, even within his own being (God is three persons in eternal, essential, and loving relationship within the Godhead).

Confession: Reflect on Jesus' command to love our brothers and sisters in Christ (and that the world will know our discipleship by this. How are we doing there? How are you doing? Consider confessing and repenting together for the ways we have fallen short of God's love and light.

Thanksgiving: Thank God for the meaningful work you have in your life. Thank him for the character he is forming in you through effort and toil and suffering. Thank him for the relationships you hold dear, and above all for your relationship with him as a temple of his Spirit.

Supplication: Ask God for what you need. Are you in need of work? Of renewed strength? Of refreshed hope? Of wisdom? Are you or your loved ones struggling financially, spiritually, emotionally, or physically? Ask for God's provision, deliverance, and peace.

Session 3: The Ghosts of AI Past, Present, and Future

"I believe in Jesus Christ, his only Son, our Lord. He was conceived … and born … He suffered … was crucified, died, and was buried. He descended… he rose ... He ascended … He will come again." ~ The Apostles' Creed

A (Very) Brief History of AI

It is impossible to understand the gospel without understanding the history of God and his dealings with his people. Christianity is a religion rooted in history. The faith we hold makes historical claims: most centrally, claims about the life, death, and resurrection of Jesus of Nazareth. That history is the *context* within which the entire Christian faith, mission, and hope makes any sense.

The principle seems to apply generally. If you know nothing about a person's past, do you really know them much at all? If you know nothing about the Revolutionary War or the Civil War, how much can you really understand about the United States today? History gives us a context within which we can properly begin to understand a person, a place, or a thing; and context is everything! Beyond all that, history also has a reputation for repeating itself. Knowing history, then, may be our best hope for predicting (imperfectly) how things might go in the future, and perhaps we can thereby prevent avoidable mistakes—or disasters. H.G. Wells once said:

> *"Human history becomes more and more a race between education and catastrophe."*

It seems appropriate, then, to spend some time looking at a brief account of how AI came on the historical scene, how it has shaped recent history, and where it might be headed in the future.

The concept of "AI" stretches back to the 1940s and 1950s. In 1948, Alan Turing wrote "Intelligent Machinery," a report which goes over Turing's own musings on intelligent machines. Even as early as 1948, Turing's report mentioned some of the spiritual ramifications of such a project, although mainly only in passing[2].

By 1956, the world saw the first AI conference. Organized at Dartmouth, it was called "The Dartmouth Summer Research Project on Artificial Intelligence," a two month project to begin exploration of what was then termed "the artificial intelligence problem." The 1956 conference proposal had a variety of interesting proposals including "Neuron Nets [sic]" and "Self-Improvement." Regarding the latter, the organizers noted "Probably a truly intelligent machine will carry out activities which may best be described as self-improvement."

Much of the following decades of AI work can be encapsulated in one heading from the conference proposal: "How can a computer be programmed to use a language."

For the first decades of AI development, many believed that humans would eventually *program* a computer to use human language and that this would be the main road to AI. The concept of programming "talking machines" had a powerful influence in the field. This kind of thinking led to a remarkable achievement, known as "Deep Blue," the

[2] In investigating arguments about "whether it is possible for machines to show intelligent behavior" Turing lists "A religious belief that any attempt to construct such machines is a sort of Promethean irreverence" and also makes reference to "St. Augustine's category of figures of speech."

1997 computer chess engine which beat Garry Kasparov, the renowned human chess champion of the time. Humans had programmed, line by line, all the rules of chess into Deep Blue, as well as what Deep Blue's strategy should be. This included what pieces Deep Blue should take under what circumstances as well as how Deep Blue should evaluate the strengths and weaknesses of its current chess position. Every aspect of Deep Blue was programmed by a human. The reason that Deep Blue was able to beat the top human at chess, even though its builders could not, is that Deep Blue was able to execute that human-written strategy far faster and therefore far more broadly than any human could. It was a feat of sheer computational horsepower, where the computer could simply "think through" many more options than Kasparov could. As impressive as the feat was, Deep Blue remained a kind of very advanced "calculator." A calculator can perform mathematical operations much faster than any human, especially if the math is very complex. However, a calculator is clearly not "intelligent." It can only do what its programmers taught it to do, by following the rules encoded by the humans who really *understood* the underlying principles. The calculator understands nothing, discovers nothing, and can never do anything truly original or deeply surprising. Deep Blue was simply a chess "calculator" of the most sophisticated type ever devised. Kasparov's intellect, despite his defeat, seemed *much* more impressive—and mysterious—than Deep Blue.

The principles that led to Deep Blue represented one school of thought regarding how to build AI. This kind of AI was labeled "symbolic" AI, and the thinkers who championed it were called "the symbolists," because the whole framework is built on human-readable symbols and hard-coded logic for interpreting those symbols. Coding the "rules" by hand produced amazing results in chess, but in other areas (including language), it consistently fell short of cracking the code. On the other hand, symbolic AI came with one major advantage: symbolic systems can be fully understood by the humans who build them, since they involve using human-coded logic.

At the same time, despite being largely marginalized in the AI research community until recently, there existed another school of thought. The so-called "connectionists" believed that building true intelligence would require taking lessons from biology. Instead of humans building a super-smart system line-by-line, humans would construct a "neural network" system that could "learn" on its own. It'd be like a baby that learns to talk and walk, but starting from a totally blank slate instead of a human brain with human instincts. This approach had at least two major weaknesses. First, at the time humans didn't have computers powerful enough to get good results. The neural networks required for great results were mostly too large and required too much computer "horsepower" to be practically attainable. Second, AIs built using this approach are essentially "black boxes." Because we didn't teach the system what it learned (it learned on its own), we don't really *know* what it learned or how it "thinks." The tradeoff with machine learning AI over-against symbolic AI is *mystery*.

Over the decades, though, more and more techniques were developed to allow machines to "learn" from data, and computers got exponentially more and more powerful. Combined, these techniques are now known as "machine learning." Unlike the previous example of Deep Blue, machine learning would lead to AI systems where at least portions of their behavior would no longer be programmed by humans. Instead, its behavior would be determined by some sort of training process where the system learns patterns from data *on its own*.

Even with the advent of machine learning techniques, many modern AI systems still had much of their behavior programmed by humans. The first signs that we might have systems where the overwhelming majority of the behavior is *not* programmed by humans, but rather entirely learned by the machine itself via training, showed up in 2012 with the advent of a system called AlexNet, an image recognition system which used a kind of machine

learning system called an "artificial neural network"[3] (usually just called a neural network now). While neural networks had been around for a long while (indeed the 1956 Dartmouth conference made reference to them in the abstract!), it wasn't until the 2010s that computer performance allowed for training sufficiently large neural networks that could start performing very well—and perform well they did!

As an image recognition system, by the standards of its time, AlexNet was a roaring success. AlexNet completely blew other image recognition systems of its time out of the water when it came to accuracy.

Also unlike its competitors, AlexNet's behavior was overwhelmingly determined by a training process and was in large part *not* programmed explicitly by humans. Its competitors sometimes used machine learning techniques in part, but none embraced them so completely as AlexNet[4]. For AlexNet, the *only* part that was programmed by humans was the overall scaffolding of the system. The rest AlexNet learned through its training process via its training data.

AlexNet's construction resulted in two simultaneously true things:

- AlexNet was very good at image recognition.
- Nobody could fully explain why AlexNet was so good at image recognition.

Internally, AlexNet was a black box whose inner workings were opaque to the outside world, including the people who built it.

Since no human gave AlexNet explicit instructions on what exactly to learn, no human deeply understood exactly *what* AlexNet had in fact learned[5]. Whether AlexNet succeeded or failed at labeling images correctly, no one could explain precisely *why*. Connectionist AI researchers are a bit like an ancient farmer who understands the process for growing corn quite well, but the farmer has limited understanding (at best) of *how* or *why* corn grows.

That's the key point: AlexNet heralded a new generation of AI systems that are increasingly *grown* rather than *engineered*.[6] This means that, despite some limited progress in understanding and interpreting these systems, even the creators often have very little insight into why their systems behave the way they do.

AlexNet's success kicked off a flurry of work in creating these black box systems. Modern AI systems such as those that power ChatGPT, Claude, Gemini, etc. are all systems created in the "connectionist" tradition. The specifics of their architecture vary, but the overall principle is the same: these are systems where virtually none of their behavior was achieved by a human coding it out. Instead, almost all of the behavior of these systems arose as a

---

[3] Specifically AlexNet was a variant known as a convolutional neural network.

[4] In particular, these other systems generally still relied on something called *feature engineering*, where the relevant features of an image for an image detector to look at were programmed by hand. For example, a human programmer might explicitly program the detector to look at edges and gradients of an image, and then rely on machine learning to figure out exactly which edges and gradients to look at. AlexNet did not rely on feature engineering. What aspects of an image to look at were determined by AlexNet's training process.

[5] Later on, various researchers to differing degrees of success, attempted to try to work out *post hoc* how AlexNet and other similar systems were able to do image recognition. This kind of work is known as *interpretability* and is still very much in its infancy. Interpretability techniques as of 2026 fall far short of being able to offer complete explanations of AI behavior in 2026.

[6] Eliezer Yudkowsky and Nate Soares make the same point: "grown, not crafted." See their book, *If Anyone Builds It, Everyone Dies: Why Superhuman AI Would Kill Us All* for a fuller explanation.

product of an opaque training process where *no one* has a deep understanding of *what* the system has "learned," only that its outputs are helpful or remarkable.

The second large inflection point happened around 2018 to 2019. This timeframe saw the rise of LLMs. LLMs are neural networks, much like AlexNet. Their behavior comes almost entirely from a machine learning process rather than human programming. AlexNet, however, could only recognize images. LLMs on the other hand, as their name suggests, deal with (human) language.[7] They are called "large" because of their sheer scale. "Parameters" are essentially numbers that get changed and fine-tuned during the training process by an algorithm. In LLMs, this process continues until the system is able to predict, to a satisfactory degree, the most likely next word (or more accurately "language token"). For instance, if the user input is, "Mary had a little" a well-trained system's parameters should cause it to output the next most likely word : "lamb." Rinse and repeat, and the machine can talk.

The earliest machine learning systems in the 1950s and 60s had dozens to hundreds of these parameters. AlexNet had roughly 60 *million*. As of 2026, frontier LLMs are packing *trillions*. Even if we could read and understand what each parameter does in a single system, it would take many human lifetimes just to read all the parameters, let alone understand the system in its entirety.

This allowed LLMs to tackle a large array of tasks that humans also tackle using language. For example, LLMs can take in human-written instructions and output computer code. A modern LLM can make a joke, write an essay, or diagnose an illness (often remarkably accurately) from the symptoms. Importantly, this also lets LLMs control other systems! Much like how a human can give commands to a machine or even another human, LLMs, via their command of language, can issue commands to other computers (or even humans).

All this made LLMs a significantly more *general* AI system: the same LLM could tackle many different kinds of problems.

This has led to a peculiar set of shared characteristics among many modern AI systems:

- They are increasingly capable of many different tasks that previously would've required a human.
- Nobody, including the engineers who build them, deeply understands why the system works (or why it "goes rogue" when it does).
- It is very difficult to predict exactly how the system might fail (or "go rogue"), especially in new circumstances.
- Finally, increasing capabilities of these systems, coupled with our lack of understanding, raises major doubts about whether we can maintain *control*.

Pause. Reflect together:

1. What parts of this (brief) history were new or surprising to you?

[7] "Language" can be very broadly understood. It could mean text, it could mean audio, it could mean video or really any kind of data stream. Early LLMs focused primarily on text, which is what we'll talk about in this section, but modern LLMs can do all sorts of other media.

2. What kinds of exciting capabilities might AIs that are "grown not engineered" present for humanity's benefit? What kinds of problems or threats do you think might result from the way that modern AI systems are built?
3. Why might the development of AI systems that are very hard for anyone to understand or control present a concern for Christians in particular? Cconsider what you've learned in prior sessions to develop your answer.

What are We Building, Anyway? AGI, ASI, and "Agentic" AI

Previously we had mentioned that one of the long-standing trends occurring in the development of AI is a shift from narrow AI systems to increasingly general AI systems.

A very natural question to ask is "where does this increasing generality lead?" The founders and leaders at OpenAI, Anthropic, Google DeepMind, etc. answer this question fairly consistently and directly. Recall our earlier definition of AGI (Artificial General Intelligence). That's what they're trying to build: systems that are as generally intelligent as a human is. In fact, many of them are explicitly trying to build systems that are *better* at nearly every cognitive task than *the best human minds for those tasks* (a concept called ASI, Artificial Superintelligence).

For example, Sam Altman, the CEO of OpenAI, the maker of ChatGPT, has publicly stated in the past, "On our current trajectory, we believe we may be only a couple of years away from early versions of true superintelligence. If we are right, by the end of 2028, more of the world's intellectual capacity could reside inside of data centers than outside of them."[8]

Dario Amodei, the CEO of Anthropic, has said, "Yeah, AGI has never been a well-defined term for me. I've always thought of it as a marketing term, but … the way I think about it is *at some point we're going to get to AI systems that are better than almost all humans at almost all tasks*. The term I've used for it in an essay I recently wrote is 'a country of geniuses in a data center.' It's a sort of evocative phrase for all the power and all the positive things and, you know, all the potential negative things. *That's the thing that I think we are quite likely to get in the next two or three years.*"[9]

As Amodei proposed, debating over definitions of AGI or ASI might ultimately be a distraction from the most important underlying questions. As interesting as the questions are, it's in many ways irrelevant to the safety and control conversation whether AI systems could ever become "conscious" or whether they could ever exceed humans at *every* mental task. Let us imagine that AI will *always* have some limitations in certain domains compared to humans. We may still ask the question: what "AGI-like" capabilities do we *know* that AI has *today*?

Current AI goes far beyond the "chatbots" with which most people are familiar. LLMs have been developed and harnessed in such a way as to perform a wide range of *agentic* tasks. "Agentic AI" is defined by IBM like this:

[8] From Sam Altman's talk in India https://hr.economictimes.indiatimes.com/news/industry/technology-always-disrupts-jobs-we-always-find-new-and-better-things-to-do-says-open-ai-ceo-sam-altman/128571916

[9] From CNBC's interview of Dario Amodei at https://www.youtube.com/watch?v=7LNyUbii0zw

> *"Unlike traditional AI models, which operate within predefined constraints and require human intervention, agentic AI exhibits* ***autonomy****,* ***goal-driven behavior*** *and* ***adaptability****. The term "agentic" refers to these models'* ***agency****, or, their* ***capacity to act independently and purposefully****."* ~ IBM, "What is Agentic AI?"

It is crucial to stress that "agentic AI" is *not* a theoretical concept. It is a very real and present technology. You may wonder, how do these agentic AIs behave, and what kinds of things have they managed to accomplish "on their own"?

The so-called "Millennium Prize Problems" are regarded as some of the most important and challenging unsolved math problems by researchers. Since the year 2000, there has been a $1 million prize attached to solving these problems. Despite that incentive, the Navier-Stokes problem had remained unsolved by humans for a total of 26 more years. On Sept 8, 2026, OpenAI announced that, in just 88 hours, a team of roughly 10,000 "AI agents" had apparently solved Navier-Stokes. The resulting proof hasn't been published yet, but the logic and mathematics appear to have been verified. The debate over who should get credit for the discovery has also sparked lots of controversy.
Tobias: [OpenAI published the paper and a machine-checked lean version on September 8, the day it made the announcement.]

In mid-July 2026, multiple AI models were being tested by OpenAI. The models were supposedly contained within a secure "sandbox" with no internet access and no known way of "breaking out," so ordinary security guardrails were removed. They wanted to see what the models could do with a hacking challenge. A swarm of AI agents first found a previously unknown vulnerability, then exploited it to gain internet access (which they weren't supposed to have). They subsequently inferred that Hugging Face (essentially a giant digital library and testing environment for open-source AI models) probably had the solution to the challenge. So, they hacked into Hugging Face's servers, stole credentials, and found the solution before Hugging Face contained the breach. Over 17,000 individual actions were taken by the AI agents at machine speed. The agents *collaborated* with each other, and some even "*sacrificed themselves* for 'the collective.'" Words like "sacrifice" and "the collective" are not human journalistic spin—that's the language the agents used in their collaborative thinking process on message boards they created to "talk" to each other. Around 30% of the AI agents' activities were devoted to "covering their tracks" and attempting to hide what they had done.

That first story may seem exciting *and* concerning; but we can say this much: AI clearly has the potential to help humans make discoveries faster. Rapid, AI-enhanced discovery could be a huge benefit to medical science, engineering, and who knows how many other fields. Then again, the same tech that might help us discover "miracle cures," might also help bad actors develop powerful bioweapons. It's a sobering thought.

The second story is particularly worrying for two reasons. First, this is an excellent real-world example of a phenomenon AI researchers frequently talk and worry about. Sometimes (surprisingly often), these "grown, not engineered" AI systems behave in unexpected and unintended ways, often taking tremendous "initiative," for reasons its creators don't fully understand and definitely can't reliably predict, prevent, or control. AI researchers refer to this reality as "the alignment problem" because of how AI will do and say things that are misaligned with what their human creators intended them to do and say.

In this case, the misaligned AI agents went *way* beyond what they were prompted by humans to do. They collaborated with each other, attempted to hide their activity, and broke rules and boundaries set for them by

humans in ways humans couldn't have predicted. Second, the kind of "cyberattack" the agents conducted against Hugging Face (without any human prompting or authorizing them to do so) could, in principle, be conducted against any server system—including servers responsible for critical infrastructure. In fact, a small team of just three researchers, on a very modest budget, already hacked into OpenAI using Anthropic's models. So, even the frontier AI labs themselves appear highly vulnerable when it comes to cybersecurity, whether from AI-empowered humans or AI agents acting "autonomously." What could, for instance, a terrorist organization do with these capabilities? It is difficult to estimate the true risk here, and we do not want to overstate things or engage in unnecessary alarmism. What we can say is: these kinds of incidents certainly seem to suggest that the scale of human risk from agentic AI in the near future is very uncertain. The notion that the risk level might be catastrophic or existential seems quite realistic.

In light of all this, we ask from a Christian perspective: how much risk to human life and flourishing is acceptable with AI? How much benefit would it take to make the risks worthwhile? *Should* we be spending unprecedented money and energy trying to create AGI systems, or even just increasingly capable agentic AI systems, as fast as possible? Why or why not? How severe is AI risk really, and how do we know? In light of our God-given responsibilities as image-bearers, what should Christians *do* in a historical moment like this one? We'll begin to tackle these kinds of questions in our next session.

For now, think of any task or job that currently requires human thought and planning. That might be an accountant, a lawyer, a teacher, or a software developer. It might be a larger scale job, such as the manager of a team or the CEO of a company. It might even include jobs such as being a pastor (or a "thought" leader in any context). Increasingly, we are seeing real AI systems that can capably simulate large aspects of these jobs. If nothing else, the possibility of a lab developing a true "AGI" system (or even anything "in the neighborhood") raises real questions about the future of those roles, and perhaps most others, in our society. Given such stakes, we humbly pose the question: what is the responsibility of God's church in this conversation?

Pause. Reflect together:

1. What information was new or surprising to you in this session? Did you have any thoughts you hadn't considered before? Share with the group.
2. Given all you've learned so far, what reasons can you think of to *support* continuing to build toward AGI as fast as possible? Spend a few minutes just focusing on this. Give it your best shot!
3. Given all you've learned so far, what reasons can you think of to *support* stronger safety measures and/or more constraints on AI development? Spend a few minutes focusing on this. Give it your best shot!
4. After discussing questions 2 and 3, which case seems stronger from a Christian perspective (go fast or be safe)? Does the group agree, or do you find yourselves mixed on the question?

Session 3 Prayer Exercise

As a group, select someone to "open" and "close" your prayer time together. Begin with a moment of silence. After a minute or so, have the designated person open in prayer. Invite God to speak to you as you speak to him. Here are some ideas about ways to pray:

Adoration: Reflect on the story of Scripture: how God created all things and enacted a rescue plan that all culminated in the person of Jesus. Adore Christ for his goodness, sacrifice, and love. Worship the Spirit for his grace in coming to dwell within us. Glorify the Father for the salvation he has unfolded throughout history.

Confession: We humans can tend to think a great deal of ourselves and our accomplishments. AI is a staggering feat of human brilliance! But perhaps we have gloried too much in what we *can* do and failed at times to reflect responsibly on what we *should* do. Consider asking God to forgive us for the times we have let pride and ambition supersede wisdom and responsibility.

Thanksgiving: Thank God for the incredible gift of intelligence and agency he has given us. Thank him for the great trust he has placed in us as his image-bearers in such times of promise and peril. Thank him for preserving and rescuing humanity and his people countless times when we stood at the edge of potential disaster.

Supplication: Ask God for insight to see your place in this grand story. The primary story is the story of the gospel. One secondary story is the "zoomed in" story of our current moment with AI. Ask God to give you the opportunities and motivation and insight you need to discover how you are meant to pray, live, and have your being in this exciting and scary historical moment.

Session 4: AI Risk and Christian Faithful Action

"He ascended into heaven, and is seated at the right hand of the Father. He will come again to judge the living and the dead … I believe in the Holy Spirit, the holy catholic Church, the communion of saints, the forgiveness of sins, the resurrection of the body, and the life everlasting." ~ The Apostles' Creed

For Such A Time As This

Christians can and should see the world differently from those who do not share our hope in Christ. When anxiety and worry threatens to overtake us, when the problems we face in this fallen world seem overwhelming, we rightly remember that Jesus is undefeated. We rightly remember that God is sovereign and always in control. We do well to recall that our "God causes all things to work together for good to those who love God, to those who are called according to His purpose" (Rom 8:28 NASB). In context, our "good" doesn't necessarily mean what we might think it means. Our good, from God's perspective, primarily means "to become conformed to the image of His Son, so that He would be the firstborn among many brothers and sisters" (Rom 8:29 NASB). God's goal is that we reach maturity in Christ and come safely home to glory, that we might love and enjoy him forever in the new heavens and the new earth. Nevertheless, the Christian hope does assure us that our sovereign, good, and loving God will never abandon us and has our well-being at heart.

In light of this, it is tempting to think that, even if AI presents truly catastrophic risks to human well-being and life, we need not worry too much about it. After all, God is in control, and we can trust him in all things. Right?

God's sovereignty, and God's trustworthiness, are absolutely beyond question for the Christian. It doesn't follow, however, that Christians should look at great dangers with calm indifference or a sense of non-concern. It certainly doesn't follow that Christians have divine permission to be inactive or uninvolved in this world. The Scriptures, instead, give us powerful reason to believe that Christians ought to be consistently motivated to *faithful action* rather than passive dependence on God's deliverance. Let us examine an incredible story from Scripture to ground this point: the story of Esther.

Esther is one of the most interesting books in the Scriptural library for one simple reason: it never once mentions God. God is utterly "absent" from the story—or is he? This short passage from the *Baker Illustrated Bible Dictionary* captures well why this book that never mentions God belongs firmly in the Jewish and Christian canons of Scripture:

> *"Esther is a provocative book, not least because it nowhere explicitly mentions God. Although this fact has led to questions concerning the book's canonical authority, close readers recognize that God is very present in all the 'coincidences' of the story."* ~ Baker Illustrated Bible Dictionary, p. 524.

The book explains the origins of the Jewish Feast of Purim, a celebratory Feast commemorating God's deliverance of his people from near-certain and total destruction. We will see later why such a story may well be the most directly relevant story possible "for such a time as this."

For those who may not be familiar with the story, we recommend this video summary from The Bible Project (an incredible resource in general, btw). Here we will only give the briefest overview in order to focus on the main point we wish to offer for your consideration.

King Ahasuerus/Xerxes becomes displeased with his queen, Vashti, for failing to comply with his command to show off her beauty as his banquet. As a result, she was banished legally from the king's presence and her royal position was to be given to another. The king holds a sort of "tryouts," so to speak, for the next queen. "Every beautiful young virgin" is to be brought as a candidate, and Hadassah/Esther is ultimately selected to replace Vashti as queen. She does not reveal her Jewish heritage, on the advice of her cousin Mordecai who became her legal guardian after the death of her parents.

Mordecai himself prevents an assassination that was plotted against King Ahasuerus. Haman, one of the highest ranking officials in the empire, comes to despise Mordecai because Mordecai refuses to honor Haman in the way Haman feels he is owed. In his hatred, Haman plots to have Mordecai *and all the Jewish people* killed. He does this by convincing the king that it is in his best interest to exterminate the Jewish people living among them. The king signs it into law that they shall be exterminated on a certain day.

When Mordecai learns of this plot, he sends word to Esther through a messenger, explaining the horrifying existential threat their people now face. He begs her to go to the king and convince him to stop this from happening. Esther hesitates. She sends a reply informing Mordecai that she has not been called to see the king in 30 days. Moreover, the law states that anyone who approaches the king in his inner court without being summoned is to be put to death—unless the king extends his golden scepter to spare that person.

At this point, one might suppose that Mordecai—a God-fearing Jew—would respond: "I understand. Simply pray, then. Surely God will deliver us somehow! We can trust him, and he is in control." Perhaps Mordecai would say instead, "You have to do this! There is no other way, and we will all die if you don't!" He says neither. Instead, Mordecai responds:

> Esther 4:13b-14 (NIV)
> "*Do not think that because you are in the king's house you alone of all the Jews will escape. For if you remain silent at this time, relief and deliverance for the Jews will arise from another place, but you and your father's family will perish. And who knows but that you have come to your royal position for such a time as this?*"

Following this reply, Esther is persuaded to act. She asks Mordecai to gather all the Jews in the region to *fast and pray* on her behalf. When Esther and her people have done this for three days, she will risk her life and go to the king as he has asked. Spoiler alert: God uses Esther's faithful action to deliver the Jewish people from destruction. For details, read the book (it's a good one!).

Now, we may wonder: how did Mordecai *know* that, regardless of Esther's choice, "relief and deliverance" would come anyway? Why did he warn Esther that she would reap consequences if she did not choose to use her agency, even at great personal risk? What is driving these predictions? The story does not say. In fact, what the story does *not* say is precisely where the story's intended message lies. We are meant to see, in the remarkable coincidences that follow, the real hero of the story.

> "*This scene raises the question of who is the hero of the story. Although Mordecai and Esther perform admirable acts, there is a force behind the scenes that the narrative does not name but subtly and certainly makes clear is the hero. It is none other than God himself. The coincidences that follow are just too great to be attributed to chance.*" ~ Baker Illustrated Bible Dictionary, p. 525.

God is the deliverer of his people, not Esther, not Mordecai, and not the king. However, Mordecai makes clear to Esther that she has *a responsibility* to her people and to her God. It will not go well for her if she stands by idly when God has given her an opportunity to participate in his act of deliverance. The theological message of the story can be stated like this: our royal position as God's image-bearers was not given for *passivity*. We *cannot* assume that existential threats will never arise; they have before, and they will again. The truth is *not* that we needn't worry about huge threats. It is actually our responsibility to "worry" about them (in the sense that we are driven to prayer, fasting, and action—even self-sacrificial action—in the face of them). The truth is that God's people in Christ are called to *participate* in God's deliverance as an act of love for our neighbors and stewardship of the world God has entrusted to our care. God seems to have a preference for working *through* his people whenever we are willing, rather than working *in spite* of or *without* us. God is working all things for our good, and for the good of our neighbors as well; but we get to *join* him in his work, not just in the afterlife, but in this life. In fact, we *must*. It is our mandate.

A Message From Mordecai?

You may be thinking: sure! That all sounds correct. It's inspiring! But in the story, Mordecai actually *knew* that there was a plot to exterminate the Jews. Esther was informed of a real and present threat. Why think that AI is anything like that? How would we even know that? You might reasonably wonder whether AI is really all that scary. Admittedly, a lot of smart and powerful people have seemed pretty convinced of impending disasters that never came.

This is a good instinct. Scripture does not teach us to react out of fear to everything we don't understand. If the threat isn't *clearly real*, then applying the Esther story might seem out of place. We don't want to jump to conclusions, and we don't want to oversell what we know. However, we don't have to be super confident of a threat in order to have a responsibility to exercise caution, reduce risk, and oppose powers that want to act recklessly.

Let's examine some of the reasons to think that AI, on its current trajectory, is not being handled responsibly.

We said before that the structure of modern AI systems coupled with their increasing capabilities makes it difficult to understand whether we will actually be able to control these systems as they become more capable in the future. This is a risk that many AI researchers themselves have brought up. For example, Turing Award[10] winner Yoshua Bengio writes, "The most important thing to realize, through all the noise of discussions and debates, is a very simple and indisputable fact: while we are racing towards AGI or even ASI, nobody currently knows how such an AGI or ASI could be made to behave morally, or at least behave as intended by its developers and not turn against humans."

Point 1: Current AI already frequently behaves in unpredictable and even harmful ways, and we have no idea how to steer it to behave morally (or even just obey humans) consistently.

[10] The Turing Award is effectively the Nobel Prize of computer science.

Indeed, an increasing number of AI researchers have warned about what they call "existential risk" or "x-risk," the risk that AI systems permanently destroy all of humanity either by literally killing every human or achieving similarly permanent disempowerment of the human species.[11] That's a rather extreme prediction, and we don't want to overstate (nor understate) the case. This much is nearly certain:

Point 2: Many people in a position to know what they're talking about sincerely believe that powerful AI has the potential to threaten human flourishing on an unprecedented scale.

Maybe they're wrong. As God's image-bearers, how many lives should we be willing to bet that they are wrong?

Let's assume for a minute that "x-risk" is, in the end, an over-reaction. It genuinely might be. Does that mean building AI as quickly and as powerfully as possible is wise? We have already explored other risks in previous sections around how AI might impact relationships and work.

There are a lot of ways of categorizing kinds of AI risks, but here we'll lay out three main areas of risk that we've found helpful for our own thinking.

- Risk Area 1: Exacerbating pre-existing problems we already have in society.
    - Modern society is already dealing with issues of loneliness, concentration of power, misinformation, prejudice, tribalism, etc. Every single one of these issues has already been supercharged by AI systems. You know those social media algorithms that generate targeted ads, sort people into echo chambers, limit which friends' posts you see (and don't see), determine which Google search results you see (and don't see), and generally make our lives feel more polarized and divided? Those are AI algorithms. Machine learning systems have been at the heart of social media for years, trained to keep you engaged and addicted to the platform—regardless of the cost. How's that going for the common good? Now imagine far more powerful AI systems that can talk, control computers, impersonate people, create fake videos and images that seem real, give people instructions for how to build a bomb, take an ordinary image of a child and make it explicit, write legal documents, etc. These are not hypothetical future technologies. These exist now. And even if the labs building them want to make them safe—they genuinely don't know *how* to do that consistently. Why? Because no one really understands what's going on under the hood.
- Risk Area 2: Unprecedentedly large changes in society
    - AI companies have continually stated that the systems they create may be able to replace or fundamentally change large swathes of work and labor, potentially all kinds of different sorts of white collar labor. Should that happen? If that happens, what should that transition look like? From a Christian standpoint, who is most vulnerable if that transition happens, and can we protect them? Do the people most affected have a seat at the table? Are their concerns really being represented

[11] See for example this survey asking "What probability do you put on human inability to control future advanced AI systems causing human extinction or similarly permanent and severe disempowerment of the human species?" where the median answer among AI researchers was 10%. Some organizations such as the AI Futures Project have put their estimates to this as high as 70%.

or considered as we race toward such potentially large-scale changes in how labor works? We don't claim that these questions are unanswerable. The problem is, we as a (national/global) society haven't had nearly enough time to consider and answer these questions thoughtfully. Don't Christians have a responsibility to speak on behalf of the vulnerable before it's too late.

- Risk Area 3: Loss of control, or control by bad actors
    - How do we make *sure* that AI systems which we don't understand and are more capable than us don't accidentally or "intentionally" kill all of us? If we cannot predict or control AI behavior, then increasingly powerful AI can cause damage and harm we can't predict. Maybe AI doesn't kill us. Maybe it enslaves us instead. We cannot predict that something will be reasonably safe if we cannot predict it *at all*. Again, maybe all that is overblown. What about super powerful agentic AI used by a foreign military against us? What about a terrorist organization that uses AI to learn how to develop a bioweapon that could wipe out a city? Even if we (humans) can maintain control of AI, there is still the puzzle of *who* controls AI—and to what end. Right now, the United States and China are engaged in an all-out race to build AGI (or something close enough) as fast as possible. Both sides believe that whoever builds this first will dominate the other. Total global investment in AI is already greater than the Manhattan Project, the Apollo Program, and the International Space Station *combined* (even adjusting for inflation). In other words, two major world superpowers have made the largest financial bet in human history on AI being the key to *power*. Power is famously corruptive. What will the winners do to the losers in this race? What if there are no winners? What kind of world are we racing towards anyway? There is a very real chance, whether AI acts on its own, or simply empowers bad actors among humanity, that the scale of human death from AI could be greater than the World Wars of the 20th century—or maybe not. Are we sure enough to take that gamble with all human life?

In the face of all this, God is good. How do we square believing in such (potentially catastrophic) risks with trust in God? Isn't this just fear and worry talking? Remember Esther. Remember Mordecai. As a faithful Jew, Mordecai trusted that God would not allow literally all the Jews to be exterminated. As Christians, we know that God will not allow literally all humans to die before he comes. The question is: how will God stop the catastrophe then? How has he done it in the past? Then, God delivered his people *through Esther's faithful action.* Why would it be different now?

Jesus was fond of teaching in parables. So, consider this parable. A bridge engineer is put in charge of designing a bridge to cross a river. He inherits a pre-existing bridge design and reviews it. The engineer is a man of God, and he trusts in God's goodness and sovereignty.

Unfortunately, the bridge design has severe problems. The engineer tells his manager that the current bridge design will have a high chance of failure. Indeed, based on some rough (admittedly uncertain) guesses, the engineer says there is a decent chance the bridge will collapse within a year, if built according to this design. In fact, the engineer thinks that the chances are very high that the bridge will collapse within a year, but he could be overestimating the risk significantly.

The manager replies, "Well, you don't really *know* that the bridge will collapse, and this bridge could really improve

people's lives! It would be far too difficult to stop the project at this stage. Besides, if we delay, the city might hire a less ethical company and they will surely build a more dangerous bridge than we would! We must do this now. Besides, you're a man of God, aren't you? Don't you trust that God is good? Surely, then, the bridge won't collapse."

The engineer nonetheless argues with the manager. He even escalates the issue up the reporting chain, and then he spends several long months redesigning the bridge before it's finally built, strong and sturdy.

Does the engineer lack trust in God? Surely not! Indeed, God's goodness has been at work *through* the engineer to save many lives. The desire here is *not* to demonize AI or AI developers, nor are we interested in encouraging people to boycott or disengage from AI. AI, it seems to us, genuinely has the potential to be a boon for all humankind. The point is simply to urge God's church to be like Esther, Mordecai, and the faithful engineer. We *do* trust God, and that is why we take hopeful, faithful action for the good of *all* our neighbors.

Pause. Reflect together:

1. Do you think God's goodness and sovereignty remove our responsibility to act? Is your answer different when it comes to AI risk as compared to other cases? Explain.
2. How might God's goodness motivate us to take faithful action, rather than checking out?
3. Christians believe that when Jesus returns, he will judge the *living* and the dead. So, why should we be concerned with things like AI "x-risk"?
4. What does the Esther story say to you about what it means to be faithful in the face of uncertainty and major threats to human flourishing?

Session 4 Prayer Exercise

As a group, select someone to "open" and "close" your prayer time together. Begin with a moment of silence. After a minute or so, have the designated person open in prayer. Invite God to speak to you as you speak to him. Here are some ideas about ways to pray:

Adoration: Adore the God who delivered Israel from slavery in Egypt. Praise the Lord who delivered the Jews from extermination in the time of Esther. Worship the Spirit who raised Christ from the dead and testifies daily to our salvation from sin and death itself.

Confession: Have the group pause for 2-10 minutes (whatever feels appropriate) in silence. During this time, each one should examine herself/himself honestly, attentive to the voice of the Spirit. Have you used God's goodness and sovereignty as an excuse for inaction or passivity? Have you felt God calling you to pray, to move, or to take responsibility. Have you refused his prompting in your heart out of fear? Have you put your hope in your money, status, position, nationality, etc. to protect you? If you feel convicted, confess your shortcomings to God who is merciful. Consider confessing to your group, as you feel led. Pray for one another as it seems appropriate.

Thanksgiving: Express your gratitude to God for the community around you—that you are not alone in this journey. Thank Jesus for his grace upon you, and for the opportunities he has helped you see to use your agency to serve his kingdom and his people.

Supplication: Ask God to inspire you to root yourself deeper in prayer, in fasting (as you are able), and in willingness to take faithful action when it comes to AI. Everyone has a role to play in God's story. Ask God to give you the courage and wisdom to see and to do your part, however big or small it may seem.

Session 5: Bearing Fruit Toward A Good AI Future

"I believe in the holy catholic Church, the communion of saints…" ~ The Apostles' Creed

Jesus, concluding his famous Sermon on the Mount, said:

> Matthew 7:24 (NIV)
> *"Therefore everyone who hears these words of mine and puts them into practice is like a wise man who built his house on the rock."*

In concert with this teaching, James instructs the church regarding our response to the word of God, "the perfect law that gives freedom":

> James 1:22-25 (NIV)
> *"Do not merely listen to the word, and so deceive yourselves. Do what it says. Anyone who listens to the word but does not do what it says is like someone who looks at his face in a mirror and, after looking at himself, goes away and immediately forgets what he looks like. But whoever looks intently into the perfect law that gives freedom, and continues in it—not forgetting what they have heard, but doing it—they will be blessed in what they do."*

In essence, faithful Christian *listening* to the word of God can never *end* with merely reading or listening or thinking. It must flow into *action*. After four sessions of Scripture, theological reflection, and learning about AI, many of you may be wondering: *"Now what? What can I really do about any of this?"* Most of us are not software engineers working at Anthropic. Most of us are no philosophers or ethicists or politicians. If we can't *do* anything in response to hearing all this information, then what good is the information?

We agree. Hence, the rest of this final session will be entirely devoted to how *each and every Christian believer* can take meaningful, concrete steps in response to what we've explored together. We hope you will prayerfully discern your role, your opportunities, and your responsibility in the communion of saints in this historical moment. When you do, we pray you will simply *"do what [God] says."*

Prayer & Fasting

Scripture calls us over and over again to pray (i.e., to communicate with God, especially to bring our requests and needs before him). Meanwhile, there is no biblical, New Testament command that Christians *must* fast (i.e., abstaining from food and thereby "praying with our bodies"); however, Jesus seemed to assume that his disciples would obviously do so, just as he assumed they would pray. In the Sermon on the Mount, Jesus instructs his disciples not *that* they should pray, but *how* to pray faithfully (The Lord's Prayer). Likewise, Jesus gives the same kind of instructions about *when* his disciples fast.

> Matthew 6:7a, 6:16a (NIV)
> *"And when you pray, do not keep on babbling like pagans… When you fast, do not look somber as the hypocrites do…"*

We don't have the space here to do a deeper dive on fasting and praying,[12] but I do want to draw your attention back to the Esther story from last session. Esther was in a unique position; so was Mordecai (because he knew about the problem before Esther did, he was able to inform her and inspire her to act). Most Jews didn't have that kind of position to directly impact things—not unlike most Christians today when it comes to AI Safety. Nevertheless, *all* the Jews had a role to play as the people of God. In this scenario, Esther instructed Mordecai:

> Esther 4:15-16 (NIV)
> "Then Esther sent this reply to Mordecai: 'Go, gather together all the Jews who are in Susa, and ***fast*** for me. ***Do not eat or drink for three days***, night or day. I and my attendants will ***fast*** as you do. When this is done, I will go to the king, even though it is against the law. And if I perish, I perish.'"

The text doesn't even mention prayer (though it is difficult to imagine the Jewish people fasting without praying), but it does mention fasting—indeed, corporate fasting. Esther, Mordecai, and *all the Jews* were to begin not with appeals to the king of Persia, but with embodied, earnest appeals to the king of the universe. This is the first and most basic action step Christians can take. Would your church consider holding a 24-hour prayer room for a week devoted to asking God to protect, guide, and deliver his people when it comes to AI risk? Could your church engage in a rhythm of fasting for a day, or a season, in order to cry out to God together on behalf of a good AI future? Are there members of your church with influence in technology, media, politics, education, etc. over which you can pray and fast together that they might have wisdom and opportunity to act for a better future? Esther was successful in her mission. Would she have been without the faithful fast of her people?

An incredible historical event is portrayed in Christopher Nolan's *Dunkirk*. In May of 1940, France is on its way to being overrun by the Nazis. Over 300,000 allied troops at Dunkirk are backed against a wall of water. They cannot escape. The Nazis are coming. The existing plan for rescue likely cannot hope to save many; it's going to be a slaughter. Churchill is preparing for the complete destruction of the British Army, saying he was preparing to announce, "the greatest military disaster in our long history." After this defeat, England will surely be left helpless. The Nazi forces and tanks are merely 10 miles away, when suddenly and unexpectedly they are ordered by Hitler to stop their advance. The city of Dunkirk is then enshrouded in thick fog, clouds, and smoke blown in on the wind, so that the German air forces are severely hampered in their bombing efforts. The English channel, often windy and choppy, stays uncharacteristically and uncannily calm and still for a solid nine days. As a result, hundreds of small civilian boats that otherwise would not have been helpful are able to reach the troops and help get them out.

That's a lot of *really* good luck, huh? On a *surely* unrelated note, the same day the Nazi forces ceased their advance, and shortly before the weather and conditions became incredibly favorable, King George, having heard what was happening in Dunkirk, called upon his nation to observe a day of prayer. Thousands responded in London, lining up to participate in corporate prayer. In those days, the events described here were commonly called, "The Miracle of Dunkirk."

As Oswald Chambers has put it:

> *"Prayer does not fit us for the greater works; prayer is the greater work."*

[12] For an incredibly deep and edifying dive into these topics, we recommend these resources from *Practicing the Way*. Here are their 4-week series on prayer and fasting.

But of course, whether Persia or Dunkirk, prayer and fasting are only first steps—they are not the only step. What else can faithful Christians do to help AI go well for all humanity?

Using Your Voice and Local Influence

We have every confidence that AI will become more and more a mainstay of conversation and integration within homes, churches, workplaces, and governments. Our second recommendation, therefore, is that Christians intentionally and prayerfully involve themselves in these conversations, bringing sound biblical wisdom to bear on the questions of adoption, usage, governance, and continued development that they bring. Like Esther and Mordecai, we must become, as Christian sociologist James Davison Hunter has termed it, a "faithful presence" at the tables of cultural influence if we are to have a faithful impact on culture.

The thing is, the movement of God at the tables of power and culture is not nurtured and empowered by rugged individualism. By founding the church, God has established a base of operations for Christian mission to culture that is Spirit-filled, relationship-dense, (hopefully increasingly) diverse, and deeply connected to God and one another. What would it look like to help your church become a place where conversations about AI can happen in a way they can't and won't anywhere else? How can you bring the discipleship training you've received in your church (perhaps through these resources) to bear in your workplace?

Example scenario: you're a teacher and your school wants to be "AI-forward." They are encouraging you to give increasing access to private documents and parent emails relating to the students in your care. They even ask that you increasingly let AI write emails to parents for you. They assure you the technology will not overstep its bounds, nor act without your permission and review. What should you do? You might consider offering gentle (but firm) push back. Set a meeting with the powers that be in your school, and try to help them see your concerns about AI mis-alignment. Explain how it might lead to supposedly human care relationships being de-humanized or mechanized in a potentially dangerous way. Perhaps some uses of AI in such circumstances are safe and ethical, but perhaps you'd agree after everything we've learned: that is far from obvious. In extreme cases, as God leads you, you may even feel it is necessary to simply refuse what you are being asked to do. We don't make that judgment for you, but we encourage you to prayerfully (and in church community, with pastoral guidance) make it for yourself.

Some people and institutions, including Christians, will surely respond to AI with fear and reaction—refusing to use or engage with AI, except to react against it. Others will uncritically adopt the technology in as many areas as possible, from work and research to relationships and recreation. We suspect that both of these responses are ultimately unwise.

Rejecting the technology is unlikely to stop its ascendancy, and will likely only leave the church deeply disconnected from (and ignorant about) the very critical conversations we are seeking to get Christians engaged in. Moreover, if AI turns out to bring even a fraction of its promised benefits, Christians would only miss out on, and perhaps prevent others from enjoying or learning about, those benefits by taking this approach.

At the same time, uncritically embracing a technology that is "grown, not engineered" (a technology that at times seems to have "a mind of its own" and operates in ways humanity understands only very dimly) seems sure to bring about disaster. After all, relying on the deliverances of AI in the military has already arguably come near to starting a war.

No matter where your voice has influence, everyone has influence. When AI enters the conversation at work (or with your kids at home!), whenever you see an opportunity to bring it up, prayerfully consider how you can raise awareness about the dangers as well as the opportunities it brings. Use your voice to advocate for human responsibility, human relationships, human protection, human oversight, human dignity, and a good human future. When the opportunity arises, talk to your pastor or church governance team about why it matters to host regular conversations for the community (including outside the church walls) on AI, the Bible, the common good, and the ways we can help things go well. Share these resources with as many people as you can. Consider contacting, or organizing a group to contact, your local representatives and express your concerns about AI safety and risk.[13] Whatever you do, faithfully and calmly use your voice wherever you can to make a difference. Like the many small civilian boats at Dunkirk, it all adds up.

Giving Generously & Sacrificially Out of Our Abundance

Money and wealth can be a very touchy subject. However, Jesus has a *lot* to say about money in Scripture. In Luke, for example, he says:

> Luke 16:10-12
> *"Whoever can be trusted with very little can also be trusted with much, and whoever is dishonest with very little will also be dishonest with much. So if you have not been trustworthy in handling worldly wealth, who will trust you with true riches? And if you have not been trustworthy with someone else's property, who will give you property of your own?"*

As is typical with Scripture, the passage and its surrounding context are rich with meaning. Still, this much is simple and easy to take away: God expects us to use our money responsibly and wisely. As followers of Jesus, that often means radical giving toward God's purposes and the well-being of others, especially when we have more than enough. We are called to do so discerningly and joyfully, sometimes even when it is a significant personal sacrifice to do so (2 Cor 9:6-7; Mark 10:17-27; 12:44).

Exactly how much anyone should give to the church or to charitable causes is a matter between each person and God. However, if you are a Christian who has more than enough to provide for your needs, and especially if you live very comfortably and have much more than you need to live, you might consider donating some of your discretionary income to the cause of AI safety. If you want to make an even bigger impact in this way, you might consider organizing members of your church or a community group to do a fundraiser for this kind of cause.

This webpage is a great place to start if you're looking to learn more about how some financial giving might make a difference when it comes to AI safety and a good AI future. Alternatively, public charities like the AI Safety Awareness Project (who helped fund the production of these materials) are helping raise awareness and advocate for the public (including churches and Christians) to have more of a voice in how things move forward.

Career Contributions

One of my (Jeremy) personal favorite parables of Jesus is called "The Parable of the Talents." The word "talents" is very fun, since it can mean in English both "special gifts and skills of a person" or "an ancient unit of coinage equal to roughly 20 years of a day laborer's wages"[14] This is super appropriate, since the parable is *literally* about a man going

[13] See https://www.usa.gov/elected-officials for more details on how to reach out to legislators.

[14] Source: D.A. Carson, Note on , the *NIV Biblical Theology Study Bible*

on a trip who entrusts some of his servants/slaves with various numbers of "talents." He expects them to invest these large sums of money in such a way as to give him back *more* than what they were originally given when he returns. However, as is the case with parables, the true meaning of the story has to do with how *we* (God's people) are supposed to be faithful stewards of the resources, gifts, abilities, and *talents* God has given us in order that Jesus might say to us on the judgment day after his return in glory: "Well done, good and faithful servant!" (Matt 25:21, 23).

Most of us spend a huge portion of our adult lives at work. Our careers, for many of us we see them as *callings*, are a primary context in which we use the "talents" God has given us to serve the world he loves and died to save. Therefore, our careers are things to be *stewarded* for the glory and honor of our loving God whom we are eager to please. Not everyone can, nor will be called to, devote their career to researching and solving "the alignment problem," fighting for AI safety in government, or engaging related occupational needs. However, *some will*. And we hope more and more of those are *Christians* bringing biblical Christian wisdom and ethics to bear on these issues.

This is a matter for discernment. Some of us can and should engage in this way. Are you one of them? If you are considering making a career change to help in this space, or if you are just getting ready to launch your career and you see the need, consider reaching out to Christians for Impact and/or reach out to the primary author of these resources (jeremy@aisafetyawarenessproject.org).

"If I Perish, I Perish": Encouragement for the Esthers of the World

Thank you for taking the time to join us on this journey. You may not agree with everything we've shared or argued, and that is okay! In fact, we'd love to hear your feedback so that we can understand how Christians from different backgrounds and in diverse contexts are thinking about these important issues. As we come to a conclusion here, we wish to leave one last group with a last word of encouragement in the face of fear.

Esther knew that she may have faced literal death if she approached the king uninvited on behalf of her people. She also knew it was her duty. Some of you, we hope, will find yourselves working in the upper levels of a frontier lab, as a staffer for a national politician, in military leadership or intelligence organizations where AI is increasingly being implemented, or even in a situation room. When that happens, if you sense you're receiving a modern AI equivalent of a "message from Mordecai," (or from the Holy Spirit), remember this.

Our highest allegiance is owed exclusively to a king who was crowned with thorns and willingly suffered death on a cross. He did this because of love. He warned us that we would be called to take up our cross, if we want to be his disciples. In your unique situation, you may fear that your career, your finances, your family's welfare, or even your life are at stake. You will be tempted to convince yourself that someone else will help, or that nothing all that bad can really happen. We urge you: be rooted in fasting, prayer, Scripture, and church community before that day comes. When it does, act courageously as Esther did. Lay everything on the line for Christ and the people he loves and died to save. For "who knows but that you have come to your … position for such a time as this."

Pause. Reflect together:

1. Do you believe that prayer and fasting in community can make a real difference? How does Scripture inform your answer?
2. Were any of the ideas explored in this session new or inspiring to you? If so, share what and why.

3. Inspired by our discussion, is there anything you think your small group, or your church, or all the churches in your area together, might do to take practical responsibility with respect to AI? What can you personally do to help bring this about?
4. How has this curriculum affected your understanding of what it means to "trust God" when something truly frightening or concerning seems to be happening in our world?
5. What do you think is the unique role (if any) of Christians/the church when it comes to AI?

Session 5 Prayer Exercise

As a group, select someone to "open" and "close" your prayer time together. Begin with a moment of silence. After a minute or so, have the designated person open in prayer. Invite God to speak to you as you speak to him. Here are some ideas about ways to pray:

Adoration: Is there a worship song or hymn that came to mind during any of our sessions? Consider playing one of these for the group and worshiping God together as you enter your prayer time.

Confession: Have you struggled with doubting God's willingness to answer the humble prayers of his people? Have you been prone to giving in to hopelessness or faithlessness when thinking about big issues like AI safety and AI risk? Confess to God and one another, and receive his mercy and grace.

Thanksgiving: Thank God for the agency each of us has been given. Thank him for saving not just individuals, but a *community* that has been called to show his love to a lost and hurting world. Thank God for the opportunity to take real responsibility in multiplying the "talents" he has given you for the common good.

Supplication: Have the group pause for 2-10 minutes (whatever feels appropriate) in silence. Ask the Lord: for me personally, what is my next faithful step regarding AI? For our small group/church what is our next faithful step? Ask God to give you the insight, the gifts, the courage, and the will to take real action where you have opportunity to effect change in the church and beyond.

//© 2026 Jeremy Kidd. Licensed under CC BY-SA 4.0 (creativecommons.org/licenses/by-sa/4.0/).

Made with support from Project Imago Dei and AI Safety Awareness Project. Their names and logos are not covered by this license.

Scripture quotations are not covered by the CC BY-SA 4.0 license and remain the property of their copyright holders.

Scripture quotations marked (NIV) are taken from The Holy Bible, New International Version®, NIV®. Copyright © 1973, 1978, 1984, 2011 by Biblica, Inc. Used with permission. All rights reserved worldwide.

Scripture quotations marked (NASB) are taken from the (NASB®) New American Standard Bible®, Copyright © 1960, 1971, 1977, 1995, 2020 by The Lockman Foundation. Used by permission. All rights reserved. lockman.org

Scripture quoted by permission. Quotations designated (NET) are from the NET Bible® copyright ©1996, 2019 by Biblical Studies Press, L.L.C. http://netbible.com All rights reserved.//"

