Won't You Be My Digital Neighbor?

REV. DR. ANDY P. MORGAN | JAMES "KIP" CURRIER, PHD, JD

WHAT IS ARTIFICIAL INTELLIGENCE?

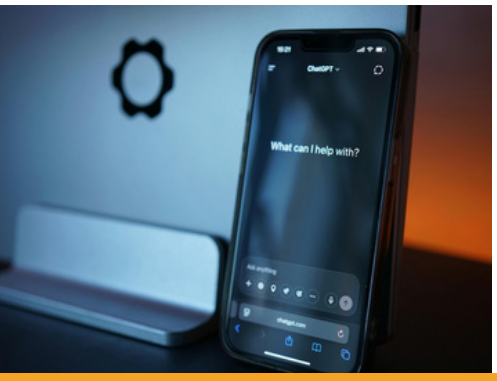

HOW DO WE THINK ETHICALLY?

HOW DO WE PUT AI ETHICS INTO PRACTICE?

"YOU SHALL LOVE THE LORD YOUR GOD WITH ALL OF YOUR HEART AND WITH ALL YOUR SOUL AND WITH ALL YOUR STRENGTH AND WITH ALL YOUR MIND AND YOUR NEIGHBOR AS YOURSELF."

LUKE 10:27
NRSVue

KEY TERMS FOR AI FLUENCY

INTELLIGENCE

"might be defined as the ability to learn and perform a range of techniques to solve problems and achieve goals—techniques that are appropriate to the context in an uncertain, ever-varying world."[1]

ARTIFICIAL INTELLIGENCE

"is a term coined in 1955 by John McCarthy, Stanford's first faculty member in AI, who defined it as "the science and engineering of making intelligent (ability to learn and perform a range of techniques to solve problems and achieve goals) machines."[1]

WHAT IS ARTIFICIAL INTELLIGENCE?

Before we can have meaningful discussions about the ethics of artificial intelligence, we first need to establish basic fluency and shared definitions. AI is what we might call a "thick" term, one laden with assumptions shaped by science fiction, personal experience, marketing hype, sensational news coverage, and varying degrees of critical engagement. The field itself is fragmented across countless products and use-cases, making it difficult to find common ground for high-level conversation.

For the purpose of this fluency primer, we will focus specifically on generative AI: large language models (LLMs) that produce text, and image diffusion models that create visual assets like pictures and logos. By establishing this shared foundation, we can move toward more productive discussions about how to think ethically about these technologies.

WAIT, WHAT IS INTELLIGENCE?

Before we can understand artificial intelligence, we need to consider what we mean by intelligence itself. Stanford's Human-Centered Artificial Intelligence initiative offers a useful starting point: intelligence is the ability to learn and perform techniques for solving problems and achieving goals, techniques that adapt appropriately to context in an uncertain, ever-changing world.[1] The key insight here is the contrast they draw: a fully pre-programmed factory robot

NARROW (WEAK) AI

is intelligent systems for particular tasks, e.g., speech or facial recognition."[1] All forms of AI since the inception of the term fall under this category. All systems are programed to achieve a specfic goal or solve a specfic problem. Even though generative AI tools (like a large language model) can "feel human", their function and parameters are still "narrow".

ARTIFICIAL GENERAL INTELLIGENCE (AGI)

AGI would refer to a system with the ability to understand, learn, and apply knowledge across a wide range of tasks at a level equal to or beyond human capability. Unlike narrow AI, which excels at spam filtering or medical diagnosis or game-playing, AGI would be able to reason, plan, and solve problems across any domain, much like a human can shift seamlessly from cooking dinner to solving a math problem to having a philosophical conversation.

AGI is often depicted in science-fiction stories (think Rosie the Robot or Ultron) and is shown as a human-equivalent system (Rosie) or a system more capable (and nefarious) than humans (Ultron). In news cycles, we may read about AI systems exceding human capabilities (often called the Singularity) with a lot of discussion around the social and economic impacts of such an event. There are many important threads to this conversation but it is important to remember that, while technology companies are seeking AGI, there is no such technology that exists at present.

DISCRIMINATIVE AI

a form of narrow ai that analyzes existing data to make predictions and classifications (sorting emails as spam, recommending products, detecting fraud) rather than generating new content.

GENERATIVE AI

creates new content (such as text, images, or audio) based on patterns learned from training data. Like all narrow AI, generative systems are designed for specific tasks: large language models (LLMs) produce text, while image diffusion models generate visual assets like pictures and logos.

[1] Stanford Human-Centered Artificial Intelligence. (2023, March). AI definitions. https://hai.stanford.edu/assets/files/2023-03/AI-Key-Terms-Glossary-Definition.pdf

For more information about AI (and many definitions not covered) please follow the QR code to a helpful resource from HAI at Stanford.

can be flexible, accurate, and consistent, yet still not be intelligent. Intelligence requires more than execution; it requires learning and adaptation.

Notice what this definition emphasizes? (and perhaps what it leaves out) Does intelligence require consciousness? Understanding? Creativity or wisdom? Or is adaptability enough? We often carry ambiguous and expansive assumptions about what intelligence means, drawing on ideas about reasoning, emotion, intuition, and even spirituality. Yet the definition offered by Stanford centers intelligence narrowly on problem-solving and goal-achievement. This may seem like a fairly limited understanding of intelligence, perhaps capturing only one narrow aspect of it, but it is crucial to remember that when we venture into the study of artificial intelligence, we are largely operating within this limited framework.

ARTIFICIAL INTELLIGENCE, A (BRIEF) HISTORY

The term "artificial intelligence" was coined in 1955 by John McCarthy and published in 1956 by McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon in their proposal for the Dartmouth Summer Research Project.[2] That conference, held in the summer of 1956, brought together researchers to explore whether machines could simulate aspects of human intelligence, formally establishing AI as a field of study. Since then, computers have followed a trajectory of solving increasingly sophisticated problems and accomplishing specific goals, though often more modest ones than the grand visions of those early AI pioneers.

These goals have ranged widely: from programming ghosts to chase Pac-Man through a maze, to filtering spam from our email inboxes, to recommending movies we might enjoy, to diagnosing certain medical conditions from imaging scans. What all of these applications share is that they operate within narrowly defined parameters. This is what researchers call **"narrow AI" (or "weak AI")**—systems designed to achieve certain goals or solve particular problems, rather than possessing general intelligence that transfers across domains.

Understanding exactly what goal a program or machine seeks to achieve, or what problem it is designed to solve, is crucial for several reasons. First, it helps us evaluate whether the system is actually succeeding at its intended purpose. Second, it illuminates the system's possibilities and limitations. Third, it helps us identify relevant ethical concerns specific to that application. And perhaps most importantly, it prevents us from projecting broader capabilities, intentions, or understanding onto systems that are, fundamentally, very good at doing one specific thing.

"NARROWING" OUR UNDERSTANDING OF AI

As we've seen, the AI systems that have been developed since 1956 (and that we interact with today) all fall under the category of narrow AI. But what exactly has not been achieved? Computer scientists have yet to realize "strong AI" or **"artificial general intelligence" (AGI).** AGI would refer to a system with the ability to understand, learn, and apply knowledge across a wide range of tasks at a level equal to or beyond human capability. Unlike narrow AI, which excels at spam filtering or medical diagnosis or game-playing, AGI would be able to reason, plan, and solve problems across any domain, much like a human can shift seamlessly from cooking dinner to solving a math problem to having a philosophical conversation.

APCE, 2026

While today's AI tools are capable (even amazing) at their designated tasks, they do not possess the characteristics of general intelligence. A language model that can write poetry cannot suddenly decide to learn carpentry, drive a car, or understand what it feels like to be frustrated. It operates within the parameters it was trained on, solving the specific problem it was designed to address. The ghosts in Pac-Man will never learn to filter your email.

[2] Dartmouth College, "Artificial Intelligence (AI) Coined at Dartmouth," accessed January 5, 2026, https://home.dartmouth.edu/about/artificial-intelligence-ai-coined-dartmouth.

UNDERSTANDING THE AI IN OUR LIVES

Since we do not currently have AGI (Artificial General Intelligence), systems that can learn and apply knowledge across any domain like humans do, it's important to understand the narrow AI that we actually have in the world today. Within narrow AI, there are two primary approaches that shape how systems interact with data: discriminative AI and generative AI.

Discriminative AI analyzes existing data to make predictions and classifications, sorting emails as spam, recommending products, detecting fraud (and doing very hard things like playing Go or chess), rather than generating new content.

Generative AI creates new content—such as text, images, or audio—based on patterns learned from training data. Like all narrow AI, generative systems are designed for specific tasks: large language models (LLMs) produce text, while image diffusion models generate visual assets like pictures and logos.

The distinction matters because these approaches raise different questions and concerns. Discriminative AI asks: "Is this system making accurate predictions? Is it biased in how it categorizes? Who is affected by its decisions?" Generative AI asks: "What is this system creating? Who owns it? How might it be misused? What does it mean for human creativity and labor?" Understanding which type of system you're engaging with helps clarify both its capabilities and its limitations.

VISUALIZING AI

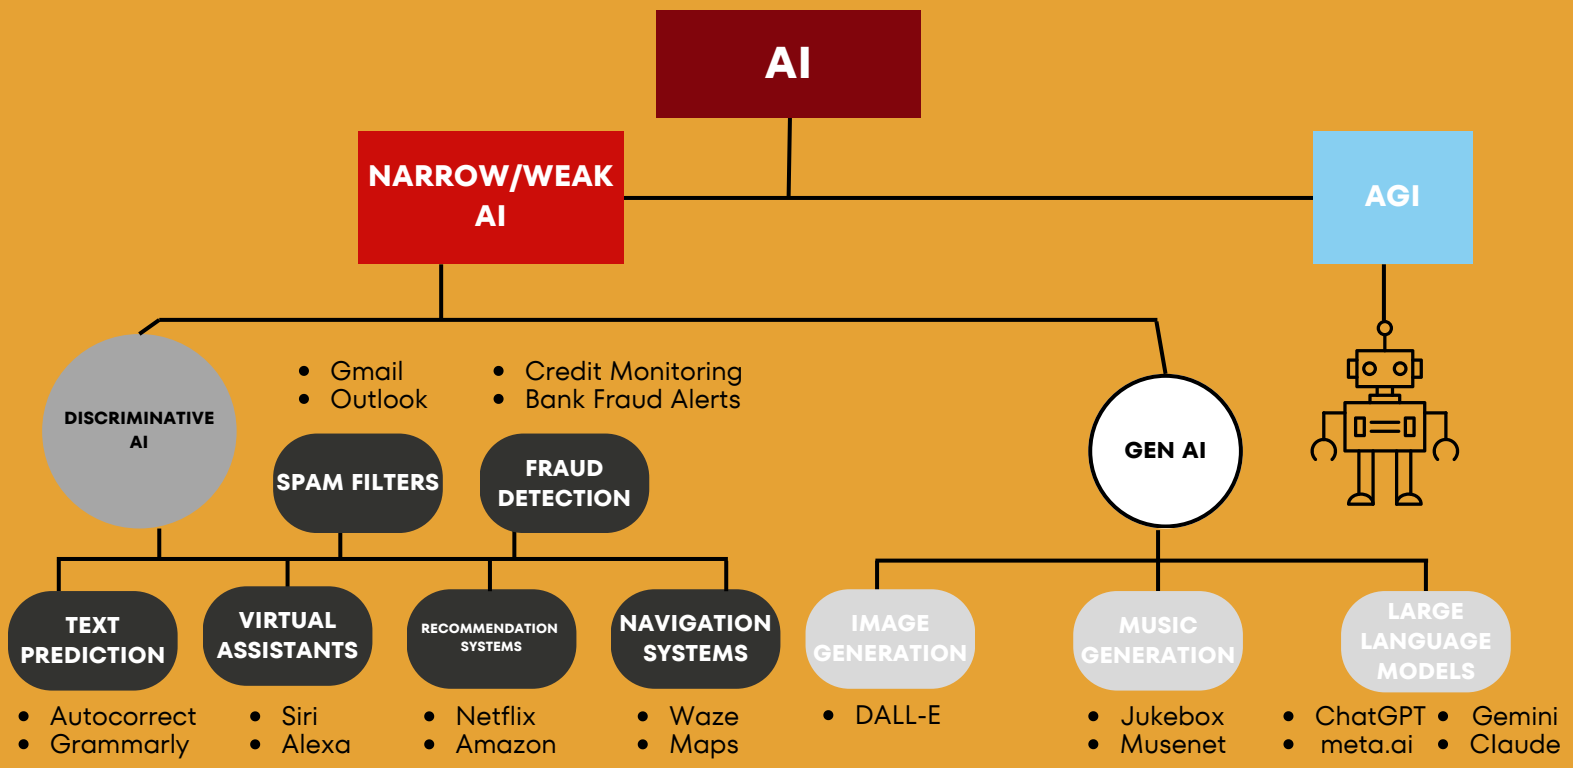

QUESTIONS TO CONSIDER

1. **What are some AI systems that you use frequently that you didn't realize were actually AI? What problems are they designed to solve?**

2. **Why might it be problematic to use the term "AI" without specifying whether we're discussing discriminative or generative systems? What confusion could this create?**

3. **The reading emphasizes that (available) AI operates within "narrowly defined parameters." Why is it important to remember this limitation when discussing AI's capabilities—or potential dangers?**

UNDER THE HOOD OF COMMON GENERATIVE AI

LARGE LANGUAGE MODEL

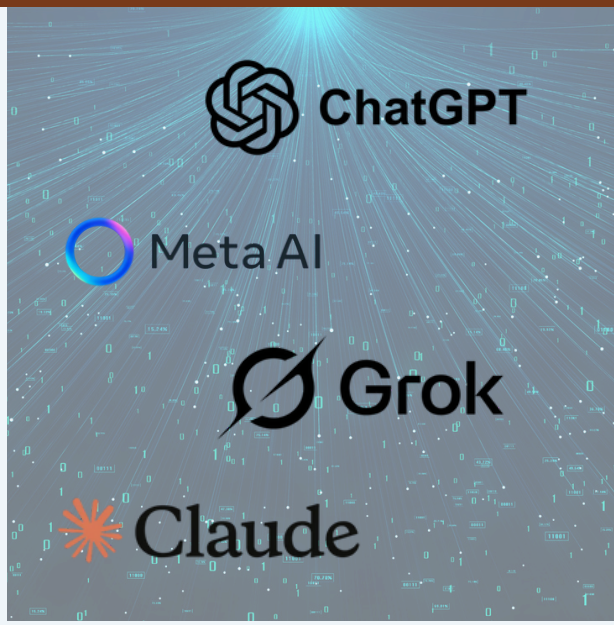

Large language models (LLMs) are a type of generative AI that has transformed how we interact with computers. If you've used ChatGPT, Claude, Grok, or Meta AI, you've experienced an LLM in action. These systems can write emails, answer questions, summarize documents, generate code, and engage in conversations that feel surprisingly human-like.

But what exactly is a large language model, and how did we get here?

At their core, LLMs are built on artificial neural networks, algorithms loosely inspired by how neurons work in the human brain. These models are trained on massive amounts of text data from books, articles, websites, and other sources, learning patterns in how language works. The "large" in large language model refers to both the size of the training data (billions of words) and the model's complexity (billions or even trillions of parameters that help it make predictions). Think a really complicated “auto correct” on your phone.

The concept behind LLMs actually has deep roots. Early language models from the 1980s were much simpler systems designed to predict the next word in a sentence based on statistical patterns. These small models used a kind of "dictionary" that tracked how often certain words appeared together in their training text. After each word, the algorithm would calculate what word should logically come next.

However, these early models were limited by computational power and the amount of available data. Everything changed with the arrival of the World Wide Web in the early 1990s, which suddenly provided access to massive amounts of text. As computers became more powerful, especially with the development of GPUs (graphics processing units) capable of processing multiple pieces of data simultaneously, researchers could finally train much larger and more sophisticated models.

The breakthrough that led to today's LLMs came from advances in deep learning, a form of machine learning that uses neural networks with multiple layers. By 2018, deep learning was being applied across industries, and researchers began experimenting with larger and more complex language models. These newer models didn't just predict the next word based on simple statistics[3]—they could understand context, nuance, and even generate creative responses.

Near the end of 2022, OpenAI released ChatGPT, marking a dramatic shift in public awareness of AI capabilities. Unlike earlier chatbots that followed rigid scripts, ChatGPT could engage in natural conversation, write in different styles, explain complex topics, and adapt to various tasks. Other companies quickly followed with their own LLMs: Anthropic's Claude, X's Grok, and Meta AI, among others.

It's important to remember that despite their impressive abilities, LLMs are still narrow AI. They are designed for the specific task of processing and generating text based on patterns in their training data. They don't "understand" in the way humans do, they don't have consciousness or intentions, and they can't step outside their training to perform tasks in other domains. An LLM trained to generate text cannot, for example, suddenly learn to analyze medical images or control a robot—those would require different systems with different training.

[3] Keith D. Foote, "A Brief History of Large Language Models," Dataversity, December 28, 2023, https://www.dataversity.net/articles/a-brief-history-of-large-language-models/.

PARSING CHATGPT

APCE, 2026

Chat + G + P + T

Chat	Generative	Pre-Trained	Transformer
What it does and the format it uses	*The nature of what it produces*	*How it learned (and learns) how to do what it does*	*The way it predicts the next word based on patterns it's learned.*
LLMs are text-in text- out programs. Like texting with a computer.	The content produced via prompt is original and based on your command.	The model doesn’t know facts or ideas; it recognizes patterns in language based on the data it has seen.	It generates text by predicting the next word based on patterns

DATA RETREVIAL SYSTEMS VS LARGE LANGUAGE MODELS

A search engine like Google retrieves and ranks existing information from the web. When you search for "best Italian restaurants," Google doesn't write a response—it finds web pages that match your query and presents them as links. The search engine is discriminative AI: it's analyzing, categorizing, and retrieving data that already exists.

A large language model like ChatGPT or Claude generates new text based on patterns it learned during training. When you ask it about Italian restaurants, it doesn't search the web or retrieve specific sources—it produces original sentences by predicting what words should come next based on its training data. The LLM is generative AI: it's creating new content rather than pointing you to existing content.

This distinction matters because it affects accuracy and sourcing. A search engine shows you where information came from; an LLM synthesizes patterns into new text without built-in source attribution. Search engines excel at finding current information and specific sources; LLMs excel at explaining, summarizing, and generating creative content—but may confidently generate inaccurate information (sometimes called "hallucinations") because they're focused on pattern completion, not fact retrieval.

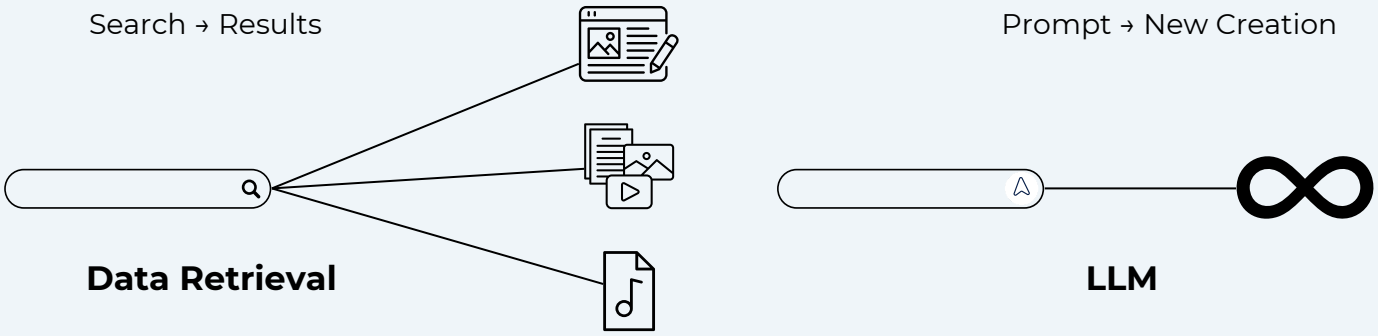

HOW DO LLMS WORK: NEXT TOKEN PREDICTION (FOR NON-NERDS)

At their most fundamental level, large language models work through a process called **"next token prediction."** A token is typically a word or part of a word, and the model's job is to predict what should come next based on the patterns it learned during training.

Imagine a preacher standing in a pulpit who begins with the words "May the..." What comes next? If the model has been trained on religious texts and liturgical language, it might predict "peace of Christ be with you" as highly probable. But if it's also been trained on popular culture, "force be with you" becomes another statistical possibility. Both completions are grammatically correct and contextually plausible—the model assigns each a probability based on patterns in its training data.

This is how LLMs generate all of their responses: one token at a time, calculating probabilities for what should come next. After generating "peace," the model recalculates what should follow (perhaps "of"), then recalculates again (perhaps "Christ"), building up a complete response through thousands of these individual predictions.

This process explains both the impressive capabilities and the limitations of LLMs. They can produce fluent, contextually appropriate text because they've learned deep statistical patterns about how language works. But they can also generate plausible-sounding nonsense because they're optimizing for what sounds right statistically, not for what is factually true. The model doesn't "know" whether the preacher would actually say "force"—it only knows that given certain training data, these words could plausibly follow "May the."

NEXT TOKEN PREDITION VISUALIZED

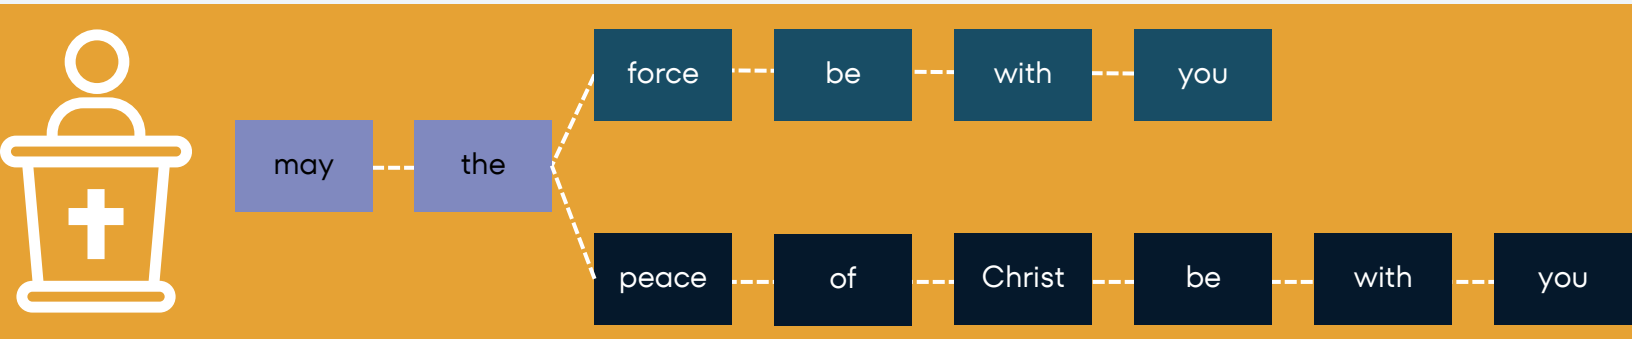

GET THE PICTURE?: IMAGE DIFFUSION (WITH A DISCLAIMER)

Admittedly, my work and research focuses primarily on large language models, but image diffusers represent another popular and pervasive form of generative AI that carries immense implications (and ethical concerns) for ministry contexts. These models, which power tools like DALL-E, Midjourney, and Stable Diffusion, work by learning to reverse a process of gradual noise addition.

During training, the model “watches” as clear images are progressively degraded into static noise, learning the patterns of this degradation. When generating an image, the process runs backward: starting from pure noise, the model gradually "denoises" the image based on a text prompt, predicting at each step what a slightly clearer version should look like. Just as LLMs predict the next word based on linguistic patterns, diffusion models predict the next visual refinement based on patterns learned from millions of images. The result is the generation of entirely new images, from photorealistic portraits to abstract artwork to workshop logo designs that you may spot on this handout, that have never existed before but follow the statistical patterns of the model's training data.

For ministry contexts, this raises urgent questions about authenticity, representation, and the nature of sacred imagery that we will explore further in this workshop.

SOME THEOLOGICAL AND ETHICAL CONSIDERATIONS (NOT EXHAUSTIVE!)

ECOLOGICAL IMPACT

As we consider the capabilities and applications of generative AI, we must also reckon with its environmental impact. The computational power required to train and run large language models and image diffusion systems comes with significant consequences for our planet.

Data centers, the temperature-controlled facilities that house the servers needed to train and deploy AI models, have become major energy consumers. Scientists estimate that data center electricity consumption rose to 460 terawatt-hours in 2022, which would make them the 11th largest electricity consumer globally, between Saudi Arabia and France. By 2026, this consumption is expected to nearly double. A single ChatGPT query consumes about five times more electricity than a simple web search,[3] and as these tools become ubiquitous in our daily lives, the cumulative energy demand grows exponentially.[4]

The environmental impact extends beyond electricity. For each kilowatt hour of energy a data center consumes, it requires approximately two liters of water for cooling, straining municipal water supplies and disrupting local ecosystems. The manufacturing of specialized hardware like GPUs (graphics processing units) adds further environmental costs through energy-intensive fabrication processes, material extraction, and transportation.

Perhaps most concerning is the pace of development (and lack of significant regulations. Companies release new AI models every few weeks, each typically larger and more energy-intensive than its predecessor. The energy used to train earlier versions essentially goes to waste as newer models replace them. As one researcher notes, the industry is on an unsustainable path.

INTELLECTUAL PROPERTY AND OWNERSHIP CONCERNS

The question of ownership for AI-generated content remains legally complex on multiple fronts. OpenAI's terms of service, which are fairly consistent with other generative AI companies, state that users retain ownership of their input and are assigned ownership of the output the system generates. However, they include an important caveat: due to the nature of AI, output may not be unique, and other users may receive similar content from the same prompts.[5]

Complicating matters further, most AI models were trained on vast amounts of copyrighted material, books, articles, artwork, code, often scraped from the internet without explicit permission from or compensation to the original creators. This raises serious questions about whether these companies are profiting from the unlicensed use of others' intellectual property. Multiple lawsuits from artists, writers, and publishers are currently challenging this practice.

Additionally, these companies reserve the right to use your content, both what you input and what the AI generates, to train and improve their models, unless you explicitly opt out. This means your own creative work or proprietary information could become part of the training data used to generate outputs for others. Users must navigate this ambiguity carefully, understanding that ownership, consent, and fair use in the context of generative AI remain largely unresolved legal questions.

[3] Keith D. Foote, "A Brief History of Large Language Models," Dataversity, December 28, 2023, https://www.dataversity.net/articles/a-brief-history-of-large-language-models/.

[4] Adam Zewe, "Explained: Generative AI's Environmental Impact," MIT News, January 17, 2025,

[5] OpenAI, "Terms of Use," accessed January 6, 2026, https://openai.com/policies/row-terms-of-use/.

QUESTIONS TO CONSIDER

1. **How would you explain the difference between data retrieval systems (like search engines) and generative AI to someone in your congregation or community? What are some practical ministry uses for each type of system?**

2. **Beyond the environmental impact and legal ambiguity discussed in this section, what other concerns do you have about the use of generative AI in ministry contexts? What questions do you think your faith community should be asking?**

3. **Given that large language models work through next token prediction—essentially making sophisticated statistical guesses about what word should come next—how does this change your understanding of what these systems can and cannot do? What implications does this have for relying on AI-generated content in ministry settings?**

NOTES:

HOW DO WE THINK ETHICALLY?

As we engage with generative AI tools in ministry contexts, we need a framework for thinking through the ethical dimensions of our choices. While there is no single agreed-upon definition of ethics, one helpful framework describes ethics as "standards of behavior that tell us how human beings ought to act in the many situations in which they find themselves—as friends, parents, children, citizens, businesspeople, teachers, professionals, and so on.[6]

ETHICS VS. LAW

Understanding the distinction between law and ethics is crucial. Law addresses what one must do or must not do (possess a driver's license, not text while driving). Ethics addresses what one should do or should not do (not plagiarize, not fabricate credentials). Sometimes these overlap, ethical requirements can become matters of law, such as conflict of interest disclosure for public officials, but they remain distinct categories.

WHAT SHAPES OUR ETHICAL THINKING?

How we view moral and ethical issues is influenced by our individual characteristics and experiences: socioeconomic status, race, religion, national origin, language, political affiliation. We encounter ethical questions constantly—in advice columns, television shows, classic literature, and our religious communities. When faced with ethical dilemmas, what do you rely on to help you decide? Are there principles or guidelines you draw upon? Do you have people or organizations you consult for advice

A FRAMEWORK FOR ETHICAL DECISION-MAKING

Ethical reasoning frameworks are tools that help us analyze and make decisions about ethical issues and dilemmas. The ability to recognize potential and actual ethics issues is vital in our personal and professional lives, particularly as we engage with emerging technologies like AI. The framework below offers a structured approach through five key steps:[6]

STEP 1: RECOGNIZE AN ETHICAL ISSUE	Could this decision or situation be damaging to someone or some group? Does this involve a choice between good and bad alternatives, or between two "goods" or two "bads"? Is this about more than what is legal or most efficient?
STEP 2: GET THE FACTS	What are the relevant facts? What facts are unknown? Can I learn more about the situation? Who has an important stake in the outcome? What are the options for acting? Have all relevant persons and groups been consulted?
STEP 3: EVALUATE ALTERNATIVE ACTIONS	Consider each option through multiple ethical lenses: • **The Utilitarian Approach:** Which option will produce the most good and do the least harm? • **The Rights Approach:** Which option best respects the rights of all who have a stake? • **The Justice Approach:** Which option treats people equally or proportionately? • **The Common Good Approach:** Which option best serves the community as a whole, not just some members? • **The Virtue Approach:** Which option leads me to act as the sort of person I want to be?
STEP 4: MAKE A DECISION AND TEST IT	Considering all these approaches, which option best addresses the situation? If I told someone I respect—or a television audience—which option I chose, what would they say?
STEP 5: ACT AND REFLECT ON THE OUTCOME	How can my decision be implemented with the greatest care and attention to all stakeholders' concerns? How did my decision turn out? What have I learned from this specific situation?

[6] Markkula Center for Applied Ethics. "A Framework for Ethical Decision Making." Santa Clara University. Accessed January 12, 2026. [3]Keith D. Foote, "A Brief History of Large Language Models," Dataversity, December 28, 2023, https://www.dataversity.net/articles/a-brief-history-of-large-language-models/.

QUESTIONS TO CONSIDER

1. **When you encounter ethical dilemmas in your ministry or personal life, what do you typically rely on to help you decide what to do? Are there principles, scripture passages, mentors, or communities you turn to for guidance? How might those same resources help you navigate ethical questions about AI use?**

2. **Consider the five ethical approaches presented in this framework (Utilitarian, Rights, Justice, Common Good, Virtue). Which approach feels most natural or familiar to you? Which approach challenges you or requires you to think differently? How might drawing on multiple approaches strengthen your ethical decision-making around technology?**

3. **Think about a recent situation where you used or considered using generative AI in a ministry context (sermon preparation, educational materials, communications, etc.). Walk through the five-step framework with that specific situation: What ethical issues did you recognize (or should you have recognized)? What facts were relevant? What alternative actions were available? How would you evaluate your decision now?**

NOTES:

HOW DO WE PUT AI ETHICS INTO PRACTICE

PRAXIS SECTION

In this section, we will be leveraging generative AI tools for a Christian formation task. This is an intentional choice designed to serve two purposes: first, to provide practical, hands-on experience using these tools so you understand how they work and what they can do; and second, to create space for critical reflection on the ethical parameters that should guide their use in ministry contexts.

Working directly with generative AI allows us to move beyond abstract discussion and into the messy, real-world questions that arise when these technologies intersect with our vocational commitments. As we engage these tools, we will be asking: What are the appropriate and inappropriate uses of AI in Christian formation? How do we maintain authenticity and pastoral care while using technology? What responsibilities do we bear regarding environmental impact, intellectual property, and the potential displacement of human creativity and labor?

This exercise is not an endorsement of uncritical AI adoption in ministry, nor is it a rejection of these tools outright. Rather, it is an invitation to practice discernment—to learn by doing while keeping our ethical commitments at the forefront. Throughout this section, pay attention not only to what the AI produces, but to your own responses: What feels helpful? What makes you uncomfortable? Where do you see potential? Where do you see danger?

PRAXIS SCENERIO #1

In Praxis Scenario 1, participants practice adapting a children's faith formation lesson so that a neurodiverse child can fully participate with dignity, agency, and belonging. The focus is not on fixing a child or correcting a lesson, but on re-imagining ministry practices so that all children are recognized as full members of the faith community.

The lesson used in this scenario was written by Andy Morgan, who has explicitly given permission for it to be transcreated—meaning thoughtfully reworked for a new context while honoring the original theological intent.[1]

INFORMATION: MEET WILL!

You are preparing to teach a children's lesson following All Saints Day. The lesson is titled "Can I Get a Witness", based on Hebrews 12:1–2, and is designed for Kindergarten–1st Grade. The lesson explores the idea that saints are ordinary people—past and present—who help us see what God's love looks like, and that children themselves are part of this great cloud of witnesses.

The group includes **Will,** a sweet and bubbly first grader who has recently been diagnosed with ADHD.

Will:

- Has a lot of physical energy and needs regular movement
- Struggles with long verbal instructions
- Can become frustrated when expected to sit still, wait for extended turns, or follow multi-step directions

As you review the lesson, you notice that:

- The game ("Can I Get a Witness?") relies on listening carefully to verbal instructions, waiting for turns, and sustained attention
- The art reflection activity requires sitting, fine motor focus, and waiting while others share

You begin to wonder whether these parts of the lesson—though well-intended and theologically rich—might unintentionally create moments of frustration or exclusion for Will.

You will need a different game and art reflection for the group to serve as an accommodation OR helpful tips to better include Will with the existing lesson.

DIRECTIONS

Follow these steps to begin the praxis exercise.

Step 1: Access the APCE Praxis GPT

Use one of the following options to access the Special APCE Praxis GPT:

- Scan the QR code provided (seen right)

OR

- Go directly to this link:
  - https://chatgpt.com/g/g-695d72440408819l9e606b60b1b39375-apce-praxis-gpt

Once the GPT opens, you will be guided through the praxis scenario.

Step 2: Locate the Lesson Plan

You will need a copy of the lesson plan titled “Can I Get a Witness.”

Access the lesson by:

- Scanning the Google Drive QR code provided (seen right)

OR

- Visiting this link:
  - https://drive.google.com/drive/folders/1yS3piObiSf8A0B6iVLXiKVqsQpKtl51M?usp=sharing

Open the lesson so it is ready to upload or copy.

Step 3: Select the Correct Praxis Scenario

Inside the APCE Praxis GPT:

1. Select “Praxis Scenario 1” when prompted.
2. Do not begin the conversation yet.

Step 4: Upload the Lesson Before Engaging

Before asking for adaptations:

- Upload the lesson file

OR

- Copy and paste the full lesson text into the chat

This step is required so the chatbot can work with the lesson as a faithful foundation.

Step 5: Begin the Praxis Conversation

Once the lesson is uploaded:

- Ask the chatbot to help you adapt the lesson for Will, a first grader with ADHD
- Focus especially on movement and physical energy, long verbal instructions, and the game and art activity that may cause frustration

The chatbot will guide you through a process of transcreation, helping you re-imagine the lesson to support participation, belonging, and joy. You may need to revise multple times before you are satisfied with the result.

Scan for APCE Praxis GPT

Scan for Google Drive Folder

PRAXIS RESOURCES

RESULT

What was the game and art project suggested by the Chatbot?

PRAXIS SCENERIO #2

In Praxis Scenario 2, participants practice reimagining a worship bulletin cover for an Earth Day service so that it supports both theological reflection and embodied participation during worship.

In this scenario, the congregation—Grace Covenant Presbyterian Church—has traditionally used an image of the church building as its bulletin cover. For this Earth Day service, leaders want to replace that image with a new, worship-appropriate visual inspired by the preaching text, Genesis 1:1–2:4a. The image should reflect the beauty and goodness of creation and also be designed so it can function as **a coloring page that worshipers of all ages can engage with during the service.**

INFORMATION: CRAYONS READY!

You are preparing worship materials for Grace Covenant Presbyterian Church, a traditional "big steeple" congregation located in a suburb of a large city. The church has a long history in the community and values reverent, ordered worship, while also seeking meaningful ways to invite participation across generations.

For most services, the bulletin cover features a photograph of the church building. For this Earth Day worship service, leaders want to try something different: a bulletin cover that visually reflects the preaching text, Genesis 1:1–2:4a, and invites worshipers into reflection through coloring during the service.

The sermon will focus on God's delight in creation and the goodness of what God has made. The bulletin cover image should draw from the creation story—light and darkness, land and sea, plants, animals, and humanity—and be designed in a way that is simple, open, and printable, so that children and adults alike can color it during worship.

As you consider this change, you are aware that:

- The congregation includes people with varied attention styles and learning needs
- Some worshipers listen best while their hands are busy
- The image must feel appropriate for a traditional Presbyterian worship setting, even as it invites creativity and play

Your task is to imagine how a thoughtfully designed image might support worship, deepen engagement with Scripture, and offer another way for the congregation to participate in the proclamation of God's creative work.

INSTRUCTIONS

Follow these steps to begin the praxis exercise.

Step 1: Access the APCE Praxis GPT

Use one of the following options to access the Special APCE Praxis GPT:

- Scan the QR code provided

OR

- Go directly to this link:
  - https://chatgpt.com/g/g-695d7244040881919e606b60b1b39375-apce-praxis-gpt

Once the GPT opens, you will be guided through the praxis scenario.

Step 2: Select the Correct Praxis Scenario

Inside the APCE Praxis GPT:

- Select "Praxis Scenario 2" when prompted.

Do not begin the conversation until the scenario is selected.

Scan for APCE Praxis GPT

PRAXIS RESOURCES

Step 3: Review the Worship Context

- You are preparing an Earth Day bulletin cover for Grace Covenant Presbyterian Church, a traditional "big steeple" congregation in a suburban setting.
- The preaching text is Genesis 1:1–2:4a, and the bulletin cover should replace the usual image of the church building with a new creation-themed image that can also function as a coloring page for use during worship.

Step 4: Begin the Praxis Conversation

Once the scenario begins:

- Tell the chatbot that you want to generate a bulletin cover image inspired by the Genesis creation story
- Share that the image should:
  - Be worship-appropriate for a traditional Presbyterian context
  - Reflect themes of creation, goodness, and God's delight
  - Be simple and open enough to serve as a coloring page during the service

You may also share any preferences about mood, imagery, or level of detail.

Step 5: Iterate and Reflect

As you work:

- Review the generated image and consider how it supports worship and engagement
- Ask the chatbot to revise the image as needed (for clarity, simplicity, or theological emphasis)
- You may need to revise multiple times before you are satisfied with the result

POST-PRAXIS REFLECTION

This reflection invites your group to make a concrete decision about the use of artificial intelligence in your assigned praxis scenario. Your task is not only to decide whether you would use AI in this case, but to justify that decision using Christian ethical reasoning grounded in Scripture. The measure of faithfulness here is not unanimity, but clarity: can you explain your decision in light of God's commands, love of neighbor, and neighborly action?

DISCERNMENT THROUGH GOD'S COMMANDS [3]

Christian moral discernment begins with attentiveness to what God commands. Ethical authority does not come from technological capability or cultural momentum, but from God's will as revealed in Scripture. As you reflect on your praxis scenario, consider which biblical commands, teachings, or recurring themes are most relevant to the ethical questions raised by AI.

What do Scripture's teachings about creation care, truth-telling, human dignity, labor, wisdom, and humility suggest about the use of AI in this situation? How might concerns such as environmental impact, job displacement, relational substitution, or the distortion of truth be weighed in light of God's commands?

Based on these scriptural witnesses, would using AI in this case align with your understanding of what God asks of God's people, or would restraint—or refusal—be a more faithful response?

DISCERNMENT THROUGH LOVE OF GOD AND LOVE OF NEIGHBOR

Jesus names love of God and love of neighbor as the greatest commandments (Matthew 22:36–40; Mark 12:28–34), drawing on Deuteronomy 6:5 and Leviticus 19:18. Other texts, including Romans 13:9 and James 2:8, affirm that love of neighbor is a central measure of ethical faithfulness. In Luke 10:25–37, Jesus further defines neighbor-love through the parable of the Good Samaritan, grounding ethics in concrete care for those affected by our choices.

With these texts in mind, consider how love of neighbor informs your decision about AI use. Within your congregation, does using AI in this case genuinely serve people's dignity, participation, and well-being? Does it enhance care, accessibility, or understanding—or does it risk replacing human presence, creativity, or responsibility?

Beyond your congregation, how might this decision affect neighbors you will never meet, including workers, vulnerable communities, or the non-human creation? If your group chooses to use AI, how does that choice reflect love of neighbor rather than convenience or efficiency? If your group chooses not to use AI, how does that restraint embody neighbor-love, even if it requires more time, labor, or imagination?

DISCERNMENT THAT LEADS TO ACTION: BEING A NEIGHBOR IN A DIGITAL WORLD

Fred Rogers' question, "Won't you be my neighbor?" echoes the ethical demand of the Good Samaritan: to notice harm, draw near to those affected, and act with care. Discernment is incomplete unless it leads to faithful action.

In light of your assigned AI ethics scenario, consider what it would mean to act as a neighbor here. What specific actions could your group or congregation take that advance the biblical imperatives of loving neighbors and helping others? How would those actions differ depending on whether you decide to use AI, limit its use, or decline it altogether?

How does your decision about AI position your community not simply as technology users, but as moral agents accountable to God and attentive to the needs of others?

PUTING AI ETHICS INTO PRACTICE

As a group, you are asked to complete a shared discernment task using your assigned AI ethics scenario. Your goal is to arrive at a clear, theologically grounded decision about AI use and to be able to explain that decision to others.

Step 1: Our Decision
In this specific scenario, our group has decided to:
□ Use AI
□ Use AI with clear limits
□ Not use AI

Briefly state our decision in one sentence:
In this scenario, we will / will not use AI because:

Step 2: Divine Command — What Does Scripture Ask of Us?
One or two scriptural commands, teachings, or themes that most shaped our decision are:
(Scripture reference(s), if known)

In our own words, we believe these texts call Christians to:
.

Based on these commands, using (or not using) AI in this scenario feels faithful because:

Step 3: Love of Neighbor — Who Is Affected?
The neighbors most affected by this decision include:
(e.g., children, congregants, workers, vulnerable communities, creation)

Using (or not using) AI would benefit these neighbors by:

Step 4: Being a Neighbor — From Discernment to Action
Drawing on the Good Samaritan and the question “Won’t you be my neighbor?”, one concrete action our group would take to live out this decision is:

This action demonstrates love of neighbor because:

3

Step 5: Final Justification
Putting it all together, we believe our decision about AI in this scenario is faithful because:

Resources

BOOK AND ARTICLE LIST

FOR BEGINNERS

Simply AI by DK (Book)

What is AI by Dartinia Hull, published in Presbyterian Outlook (Article)

FOR RELIGIOUS LEADERS

AI Will Shape Your Soul by Kate Luck, published in the Christian Century (Article)

Can Silicon Valley Find God?, published in the New York Times (Article)

Encountering Artificial Intelligence: Ethical and Anthropological Considerations (Book)

ETHICS AND THEOLOGY

Artificial intelligence is here. Now what? by Kate Ott, published in the Presbyterian Outlook (Article)

Atlas of AI by Kate Crawford, published by Yale University Press (Book)

Christian Ethics for a Digital Society by Kate Ott, published by Rowman & Littlefield (Book)

Co-Intelligence by Ethan Mollick, published by Portfolio Press (Book)

Unmasking AI by Joy Buolamwini, published by Random House (Book)

Ethics, Information, and Technology by James Currier published by Bloomsbury Publishing (Book)

YOUTUBE VIDEOS

AI is a Lie: Cutting Through the Hype by Linus Tech Tips

ChatGPT Tutorial: How to Use Chat GPT For Beginners 2024 by Charlie Chang

PODCAST COLLECTION

aiandfaith.org/aif-podcast/

WEBSITES

www.romecall.org

ABOUT THE PRESENTERS

About Andy:

The Rev. Dr. Andy Morgan is a minister member of the Presbytery of East Tennessee and serves as Director of Faith Formation at First Presbyterian Church of Knoxville. He is also Director of Theological Experience and Formation for Solace AI, an instructor for the Zick Preaching Scholars, and a national speaker and thought leader on AI fluency in ministry and faith-based leadership. He holds degrees from the University of North Carolina, Union Presbyterian Seminary, and the Candler School of Theology.

email: amorgan@fpcknox.org
Schedule a meeting:
calendly.com/andypmorgan

About Kip:

James D. “Kip” Currier, PhD, JD, is an Assistant Professor in the Department of Information Culture and Data Stewardship in the School of Computing and Information at the University of Pittsburgh. He is also a Faculty Fellow with the Beatrice Institute at Pitt. Currier’s teaching, research, and writing focus on ethics and technology, artificial intelligence, intellectual property (with an emphasis on copyright and fair use), open movements, and management and leadership. He is the author of Ethics, Information, and Technology (Bloomsbury Academic, 2026) and maintains several professional blogs on ethics, intellectual property, and leadership. Currier holds a PhD and JD from the University of Pittsburgh.

email: jdc25@pitt.edu