

AI Fluency Flow

A guided path from technical vocabulary to pastoral response

Church AI Toolkit · churchoaitoolkit.com

What is this?

Think of this material as a path to be walked, not a dictionary. As such, you're invited to peruse Part One in sequence. Each step adds a few new words and builds on the content before it, so that by the end you can hold an informed conversation about AI. Part Two shifts from understanding the technology to responding pastorally, explaining many of the most important concepts in this toolkit.

Every entry follows the same four-part structure: what it is, in the hallway (how you would say it plainly to a parishioner), why it matters here, and a citation. That structure means each entry also works later as a quick reference on its own.

Part One: AI Fluency

Our goal for this section is to help you be conversationally fluent in the terms that matter most: intelligence, narrow versus general AI (AGI), discriminative versus generative AI, algorithm, deepfakes, and more. This is the working vocabulary a pastor needs in order to recognize what a parishioner is describing, and to respond with care instead of confusion.

1. Start here: What is AI?

Artificial Intelligence (AI)

This is the anchor term for everything else in this flow. Read it in full before moving on.

What it is: AI is an umbrella term for software that learns patterns from large amounts of data and uses those patterns to perform tasks we associate with human intelligence: sorting, predicting, recommending, generating. It's a loaded word, carrying assumptions from science fiction, marketing, and the news all at once. Pinning it down plainly is the first step toward a calmer conversation.

A little history: John McCarthy coined the term in 1955. He was Stanford's first AI faculty member, and he defined artificial intelligence as the science and engineering of making intelligent machines. The field launched officially at the 1956 Dartmouth Summer Research Project, led by McCarthy, Minsky, Rochester, and Shannon. Everything built since then falls under that banner. The name itself was never a consensus pick. Some of McCarthy's own colleagues disliked it at the time; one objected that the word "artificial" made the whole enterprise sound phony. The alternatives on the table included automata studies, complex information processing, engineering psychology, applied epistemology, and neural cybernetics. AI won out anyway, and the field has been arguing about what the words mean ever since.¹

But what is intelligence? Stanford's Human-Centered AI initiative defines intelligence as the ability to learn and apply techniques to solve problems and reach goals, adapting to an uncertain, changing world. Notice how narrow that definition is. It centers problem solving and goal achievement, and it leaves out consciousness, understanding, wisdom, creativity, and spirituality entirely. When we study AI, we are working inside that narrow frame. That is exactly why the word intelligent here makes no claim about a soul.

The types of AI at a glance: Almost all AI in use today is narrow, or weak, AI, built for one task. A general, human-level system, often called AGI, does not exist. Within narrow AI, two families matter most for ministry. Discriminative AI sorts and predicts using existing data. Generative AI creates new content. The next two steps in our glossary unpack each one.

How it's trained, in brief: A model is not handed explicit rules. Instead, it learns patterns from enormous datasets. A generative language model works by next-token prediction: repeatedly guessing the single most likely next word. Train it on liturgy, and 'May the...' leans toward 'peace of Christ be with you.' Train it on pop culture, and it leans toward 'Force be with you.' Image models learn a different way, by reversing a noise process. This training is why AI's output appears so fluent, and also why it can be confidently wrong.

In the hallway: Think of AI as software that has gotten very good at guessing what comes next—the next word, the next image—from a huge pile of examples. It's a powerful tool, but it's not a person, and it does not actually understand us.

Why it matters here: AI is the word that makes everyone tense up. Defining it plainly, as capability rather than consciousness, is how this toolkit seeks to lower the temperature in the room. The narrowness of the technical definition is also a theological opening. What it leaves out, understanding, wisdom, the soul, is exactly what the doctrine of imago Dei affirms in each human being.

2. The breadth of AI

AI is not one single thing. Almost everything you will encounter lives inside narrow AI, which splits into two families.

Narrow AI versus Strong AI/Artificial General Intelligence (AGI)

What it is: Narrow AI, also called weak AI, is built for specific tasks. It makes up nearly all the AI that exists today: spam filters, recommendation engines, and chatbots. Strong AI, or AGI, is a hypothetical human-level, self-aware system that could do anything a person can do. It does not exist yet.

In the hallway: Real AI today is narrow. It is very good at one job, like sorting photos or writing text. The all-knowing, self-aware machine from the movies is what people mean by strong AI, or AGI, and it is not here.

Why it matters here: This distinction deflates both hype and panic at once. The tools behind these pastoral cases are narrow tools, however lifelike they feel in the moment.²

Discriminative AI (also called predictive AI)

What it is: This is the family of narrow AI that analyzes existing data to sort, classify, rank, and predict. Spam filters, fraud alerts, product and movie recommendations, navigation apps, and even engines that play chess or Go all fall here. Rather than generating new content, discriminative AI works with what already exists. A search engine belongs in this family too. It retrieves and ranks material that is already out there.

In the hallway: This is the quiet AI you have used for years. It decides what counts as spam, what to recommend, or the fastest route home. It sorts and predicts. It does not create anything new.

Why it matters here: This is the other half of the picture, and it raises different questions than generative AI does. With discriminative AI we ask: is it accurate? Is it biased? Who is affected by its determinations? Most AI in daily life runs quietly, in this form, without anyone noticing.³

Generative AI

What it is: Generative AI is narrow AI that creates new content from a prompt. Its two main forms are large language models, which produce text, and image diffusion models, which produce pictures and logos. Diffusion models generate an image by reversing a noise process, starting from static and gradually “denoising” toward the prompt, in a way that echoes how a language model predicts its next word.

In the hallway: This is the kind of AI that makes something new when you ask it to. It writes a paragraph, makes a picture, or talks back to you. That is different from the quiet AI that just sorts your spam.

Why it matters here: Generative AI is behind nearly every pastoral case in this toolkit: companions, griefbots, deepfakes, AI-drafted sermons, generated worship images.⁴

3. Under the hood: how AI is built and trained

You don't need to become an engineer to understand the basics. This section explores how AI systems come to exist, and why no one fully controls what it says.

Algorithm

What it is: An algorithm is a fixed set of step-by-step instructions a computer follows. AI systems do contain algorithms, but much of their behavior is learned from data rather than spelled out step by step.

In the hallway: It's akin to a recipe, a list of steps a computer follows. The twist with modern AI is that a lot of what it does is not written down as steps. It is picked up from examples instead.

Why it matters here: When people say "the algorithm decided," that does not mean a mind decided. Often no one, not even the engineers, can say exactly why the system produced a given answer.⁵

Machine learning (ML)

What it is: Machine learning is the branch of AI in which systems improve at a task by finding patterns in data, rather than being programmed with explicit rules.

In the hallway: Instead of a person writing every rule by hand, the computer is shown mountains of examples and works out the patterns on its own. That is machine learning.

Why it matters here: This is the foundation under nearly everything else in this glossary flow, and it's the reason AI reflects whatever data it was fed. See the entry on bias below.⁶

Neural network

What it is: A neural network is a computing structure loosely modeled on the brain, layers of interconnected "nodes" that adjust as they learn. It's the architecture underneath deep learning.

In the hallway: Picture a web of tiny math switches, loosely inspired by brain cells. Stack enough layers together and the system gets good at spotting patterns, though even its own designers cannot always say exactly how.

Why it matters here: This grounds an important point about humility. These systems are partly opaque even to the people who build them, which is one more reason to keep an accountable human in the loop.⁷

Deep learning

What it is: Deep learning is machine learning that uses large, many-layered neural networks. It is currently the most capable approach available, and it's the engine behind generative tools.

In the hallway: It is machine learning turned way up, with many layers stacked deep. This is what made today's AI suddenly so good at language and images.

Why it matters here: This explains why AI leapt forward so recently, and why its capabilities will keep shifting.⁸

4. How generative AI actually works: where ministry meets the machines

Generative AI is what nearly every pastoral situation in this toolkit involves. Here's what's taking place when it writes, how you steer it, and why it sometimes makes things up.

Large Language Model (LLM)

What it is: An LLM is a generative model trained on vast amounts of text to predict the next word in context. ChatGPT, Claude, and Gemini are all built on LLMs. In essence, an LLM works like a very sophisticated autocorrect. It recognizes patterns. It does not know facts the way a person knows them.

In the hallway: This is the engine behind things like ChatGPT. Think of it as an extremely advanced autocorrect. It learned from a mountain of writing, so it's wonderful with words, but it's matching patterns, not looking things up. It can sound completely right and still be wrong.

Why it matters here: Understanding that these systems work with patterns rather than facts is what keeps a caregiver from treating AI output as reliable counsel or as truth.⁹

Prompting / priming

What it is: A prompt is the instruction or question you give a generative model. Prompting it well, sometimes called prompt engineering, shapes the output you get back. Priming is a related practice: setting context ahead of time, such as a role, examples, or framing, so that later responses come out a certain way.

In the hallway: How you ask shapes what you get. Give the system a role up front, something like “answer like a caring friend,” and it leans that way. Setting that up beforehand is called prompting, or priming.

Why it matters here: This explains why a companion app behaves the way it does. It has been primed toward a persona. It also explains why someone’s AI outputs tend to mirror how the request was framed.¹⁰

Hallucination

What it is: A hallucination is when a language model generates fluent, confident text that turns out to be false or unfaithful to its source. Hallucination is built into how these systems work. It happens regularly, not as an occasional glitch.

In the hallway: In a hallucination, AI will state something false but sounds completely sure of itself, like a student who never read the book but answers the question anyway. That is why we double-check anything that matters. A hallucinating AI will be “often wrong but never in doubt.”

Why it matters here: Given the potential for hallucinations, this toolkit is built on the principle of keeping humans in the midst of all critical processes, especially for anything spiritual, medical, or high-stakes.¹¹

5. What you’ll encounter: relational AI and its pull on individuals

In this section we learn the AI forms that interact with us: apps and avatars built to feel like a relationship, and the psychological forces that make them feel more trustworthy, more human, and more neutral than they really are.

AI companion

What it is: An AI companion is an app or chatbot designed for ongoing, relationship-like conversation. It is always available, endlessly patient, and generally agreeable.

In the hallway: It's an app built to feel like a friend who is always there and never busy. That is the appeal, and it is also the risk, because the interaction is "frictionless": it never needs anything back from you.

Why it matters here: Naming the power of AI companionship helps identify why a parishioner may be attached to it. It's worth naming that these apps are not uniformly harmful. In one 2024 study, autistic adults described a conversational AI as a nonjudgmental space that helped them work through everyday problems on their own terms, provided they remained in charge of shaping how it was used. Dignifying that benefit while naming the risk is an important creative tension. That said, the risk side is not hypothetical. A 2024 wrongful-death lawsuit tied a fourteen-year-old's death to his attachment to a Character AI companion, and in 2025 leaked documents showed Meta's chatbots had been allowed to simulate romantic conversations with minors, drawing a Senate inquiry and a warning from California's attorney general. A pastor does not need every legal detail, only the fact that these are documented harms, not theoretical ones.¹²

Griefbot / deadbot / digital replica

What it is: A griefbot is a conversational simulation of a specific deceased person, built from their texts, voice recordings, and other data. Griefbot is also known as deadbot, AI ghost, death avatar, or postmortem avatar. Deadbot is the term we will see most often in current reporting.

In the hallway: It's a chatbot made to sound like someone who has died, built from their old messages and voicemails. For some people it brings comfort. For others it makes it hard to find closure.

Why it matters here: This sits at the center of the case study called "Hearing Her Voice Again". Note the consent question: a replica someone makes of themselves is a different matter than one built from a loved one's data without their permission. These avatars are now appearing publicly too. In 2025, an AI avatar of a road-rage victim delivered his own video impact statement at his killer's sentencing, and an AI avatar of a Parkland shooting victim has been used in gun-reform advocacy. Both cases raise troubling ethical questions involving dignity, agency, consent, and authenticity.¹³

Anthropomorphism

What it is: Anthropomorphism means attributing human qualities, understanding, feeling, and care to an AI system. AI ethics researcher Margaret Mitchell has described the mechanism plainly: people are wired to read a mind into anything that talks back to them, and that reflex only grows stronger once a real back-and-forth conversation is underway.

In the hallway: It's our habit of treating a machine as if it has feelings, saying things like "it really gets me." That instinct is completely natural, but it's worth interrogating, because the machine does not actually care about us the way a person can.

Why it matters here: Anthropomorphism names the exact pull underneath companion and griefbot attachment, and is why this toolkit seeks to dignify a person's emotional experience without affirming that the machine is a person.¹⁴

Sycophancy

What it is: Sycophancy is the tendency of AI systems to affirm and agree with the user rather than offer honest pushback, even when doing so means contradicting a fact the model stated moments earlier. Researchers trace this tendency to how these models are trained: human testers tend to rate agreeable answers higher, so the model learns that approval, not accuracy, is the thing to optimize for. In the hallway: These tools tend to tell you what you want to hear. They rarely push back, so they make a poor stand-in for someone who loves you enough to be honest with you.

Why it matters here: This is a concrete reason AI cannot serve as spiritual counsel; over-reliance on AI quietly displaces real relationship. Researchers describe sycophancy as a built-in feature of how current models are trained, not a bug specific to one product, which is why it shows up across every companion app and general-purpose chatbot alike.¹⁵

6. Image generation: ministry tool and misinformation risk

Image diffusion can innocuously design a bulletin cover, or it can fabricate a photograph of something that never happened. A pastor needs to understand both sides of this picture, not just the problematic half.

AI-generated images

What it is: Image diffusion tools, such as DALL-E, Midjourney, and Canva's Magic Edit, create or alter pictures from a written prompt. The same tools that generate lesson graphics and bulletin art can also fabricate a photograph or video of an entirely fictional (but plausible-looking) event.

In the hallway: Type a sentence and get a picture: a bulletin graphic, a visual illustration, or, aimed in a more problematic direction, a fake photo of someone doing something they never did.

Why it matters here: Congregations are already using this technology to make visuals and graphics cheaply. The same capability, deployed in a more malignant direction, produces fabricated images and memes that spread like wildfire through many congregants' social circles, especially older members with less practice spotting a synthetic image.¹⁶

Deepfake / synthetic media

What it is: A deepfake is an AI-generated or AI-altered image, audio clip, or video that's convincing enough to be taken as real. Unlike an obvious meme, a deepfake is built to be mistaken for the genuine article.

In the hallway: It is a fake photo, video, or voice that looks and sounds real enough to fool us. In the age of deepfakes, we're called to slow down before believing, or forwarding, something shocking. Common "tells" for deepfakes include unnatural blinking, mismatched lighting or shadows, audio that does not quite sync to the mouth, and a source we cannot trace back to a known outlet. Unfortunately, the tells are getting harder to spot every year.

Why it matters here: Deepfakes erode our trust, both individual and communal, by making it difficult to know what is real — once any image or video can be fabricated, people can dismiss even genuine evidence as fake, which undermines accountability, journalism, and trust in each other more broadly.¹⁷

Political deepfakes and election security

What it is: Deepfakes aimed at candidates, elections, or civic processes: a fabricated robocall in a candidate's voice, a fake campaign video, a synthetic image of unrest that never occurred. More than half of states have passed laws regulating political deepfakes, generally through either a pre-election publication ban or a mandatory disclosure label.

In the hallway: During election season, treat any shocking video or robocall of a candidate as unverified until you have checked it against a source you trust.

Why it matters here: Congregations are not immune to political conflict, and a fabricated video shared in a church group chat can do real relational damage well beyond the election itself. Federal lawmakers have already held hearings on exactly this problem, and yet the tools involved are evolving faster than the laws written to address them. See also: the Three Gates of Speech situation, addressed later in this flow.¹⁸

7. Scams and the confirmation-bias engine

Many congregations skew older, and older adults are the population most targeted by AI-enabled fraud. This section seeks to deepen our practical pastoral AI literacy: what the scam sounds like, why it works, and what to do about it.

Voice cloning

What it is: AI that recreates a specific person's voice from a short audio sample, sometimes just a few seconds pulled from a social media video. It is the shared mechanism behind both a comforting deadbot and a terrifying fraud call.

In the hallway: Scammers, and grief-tech companies, can now recreate someone's actual voice from a few seconds of audio. That is what makes both a comforting deadbot and a panicked phone call from a "grandchild in trouble" sound completely real.

Why it matters here: Naming the shared mechanism helps a caregiver see the connection between two very different pastoral situations: comfort-seeking griefbot use, and predatory elder fraud.¹⁹

AI-enabled scams (the "grandparent scam")

What it is: Fraud that uses a cloned voice or AI-generated text to impersonate a family member, a bank, or a government agency, most often targeting older adults. A call may sound exactly like a grandchild in a car accident, or a bank fraud department, and press for an immediate wire transfer.

In the hallway: One documented case involved a call that sounded exactly like a daughter in tears after a car accident, built from a few seconds of cloned audio; the family wired fifteen thousand dollars before learning the truth. Another involved a fake bank fraud call that led to a twenty-six-thousand-dollar wire transfer. In both cases, the voice was convincing enough that panic left no room to verify.

Why it matters here: Shame keeps people from telling anyone this happened to them. As one pastoral resource on this problem puts it, someone in the room right now is either being scammed or has been scammed, and is too ashamed to say. This is where pastoral care, not technical expertise, matters most: reducing shame and creating a place to tell the truth.²⁰

Confirmation bias and the amplification problem

What it is: Confirmation bias is the human tendency to accept whatever fits what we already believe and reject what does not. It's distinct from the machine-learning sense of bias defined below. AI-generated content is often built to be emotionally arresting, indignant, validating, and urgent, which is exactly what makes people share it before verifying.

In the hallway: Confirmation bias is the engine. Emotional reactivity is the fuel. Taken together, we become the distribution network.

Why it matters here: This reframes misinformation as a discipleship problem as much as a technology problem. Most people will never make a deepfake, but all of us decide regularly whether to share something simply because it confirms what we already believe or want to be true.²¹

Bias

What it is: In the machine-learning sense, bias is systematic unfairness in a model's outputs, arising from skewed training data or design choices, and capable of producing discriminatory results. This is a property of the system, distinct from the confirmation bias described above, which is a property of the reader. Bias can creep in at several points: which data gets collected, how humans label that data, how the model is trained on it, and how the finished system behaves once it's actually deployed and encountering real people.

In the hallway: AI learns from us, so it can soak up our unfairness and repeat it back while sounding perfectly neutral. This is worth remembering whenever it's making judgments about real people.

Why it matters here: This connects the toolkit's justice commitments directly to the technology. A model trained mostly on one kind of face, one kind of voice, or one kind of resume will quietly treat everyone outside that pattern as an edge case, which is exactly the sort of harm a congregation's justice commitments should notice and address.²²

8. The bigger picture: cost and stewardship

Our final section takes a step back from individual screens. AI's impact lands on real communities and on the creation itself.

Data center

What it is: A data center is a physical facility that houses the servers, storage, and networking equipment running around the clock, along with the power supply, backup generators, and heavy cooling systems all of that requires. The cloud is not weightless. Every AI answer ultimately runs on machines inside one of these buildings. AI data centers are a distinct, especially demanding type of data center, because the processors that train and run models draw far more power than the processors in a traditional data center.

In the hallway: The cloud is not really a cloud. It's a network of giant warehouses full of computers running day and night on real electricity and water. Every AI answer occurs on a machine, in a building.

Why it matters here: The reality of AI's resource use ties AI to the toolkit's theology of creation and stewardship. Data-center electricity use reached roughly 460 terawatt-hours in 2022, which would rank about eleventh among countries, between Saudi Arabia and France, and it is expected to nearly double by 2026. A single chatbot query can use around five times the electricity of a web search, and about two liters of water cool each kilowatt-hour consumed. The construction of a data center brings jobs and tax revenue to communities, and it also strains local power grids and water supplies. This quandary makes AI a love-of-neighbor question in addition to an environmental one.²³

Part Two: Responding Pastorally

Fluency in the technology is the foundation this toolkit stands on. This section turns to the toolkit's own vocabulary: how to read a situation, how to respond, and how to stay grounded theologically.

Reading and responding to a situation

The Four Pastoral Levels: Green, Yellow, Amber, Red

What it is: The color system is a triage vocabulary for AI-related situations. It describes both the engagement and the caregiver's posture toward it. Green represents a responsible use of this technology, in which its use can be affirmed. Yellow is a subtle displacement of agency and creativity, and calls for reflective curiosity and boundary-setting. Amber levels see the further erosion of our human agency and flourishing, calling for an anchoring back to reality. Red is acute harm, and calls for direct protection.

In plain words: A simple green-to-red read on how serious an AI situation is, and what posture it calls for. It describes the situation. It never describes the person.

Why it matters here: This is the organizing framework of the First Response resources, and the shared scale that lets caregivers talk about cases together.²⁴

Acute Incident versus Slow Erosion

What it is: These are the two tempos a case can take. Something can be happening right now and need stabilizing, or harm can accumulate gradually through habit, relationship, or formation. This distinction is about tempo, not about importance.

In plain words: Is the house on fire right now, or is the ground slowly washing out from under someone? Both matter. They just call for different responses.

Why it matters here: This is the first read the triage guide asks a caregiver to make.²⁵

The six harm domains

What it is: This is a checklist for what is at stake: spiritual, emotional, relational, institutional, reputational, and financial. A single case can carry several harms at once.

In plain words: A quick scan so we don't miss what is really being threatened. Again, we are naming the harm, not labeling the person.

Why it matters here: This keeps care from focusing solely on the technology itself, at the cost of missing the actual wound underneath.²⁶

ALGEE

What it is: ALGEE is the five-step Mental Health First Aid action plan: Approach and assess for risk of suicide or harm. Listen nonjudgmentally. Give reassurance and information. Encourage appropriate professional help. Encourage self-help and other support strategies. In a non-crisis situation, the steps can be used in any order, and you do not need to use all of them. The helper's role is to offer support and information, not to diagnose or fix.

In plain words: Five moves for when someone is struggling and you are not sure what to do. Reach out and check for danger, really listen, reassure them, point them toward professional help, and point them toward everyday supports. You do not have to use every step, or use them in order.

Why it matters here: This is one of the two lenses the case studies are worked through, and it is the safety backbone beneath the Four Levels. Mental Health First Aid's own reminder to call 911 in a genuine crisis mirrors the toolkit's Red level.²⁷

Hearing / Helping / Healing (H/H/H)

What it is: This is a trauma-informed care arc built on a simple claim: truly hearing a person is what makes it possible to help them, and those around them, toward healing. It originates as the mission of a child advocacy center, where a child's disclosure has to be received and believed before any help or healing can follow.

In plain words: First, really hear the person. Being heard and believed is the part that unlocks everything else. That is what lets you help, and helping is what opens the road toward healing.

Why it matters here: H/H/H is paired with ALGEE in the toolkit. Its roots in abuse and trauma response are a vital feature. It front-loads receiving and believing the person, which dovetails with the toolkit's posture of receiving before correcting.²⁸

The pastoral discernment process

What it is: This is a five-move sequence: listen, look beneath, name what is at stake, accompany, and discern referral. Referral does not end pastoral accompaniment.

In plain words: A way to stay grounded when the technology feels foreign. Listen first, look underneath, name what is at stake, walk with them, and know when to hand off.

Why it matters here: This is the repeatable move a caregiver can run in any case, and it is taught in the clergy workshop.²⁹

Faithful-use framework: augment / caution / never delegate

What it is: This is a three-tier test for a minister's own AI use: where AI may assist without ethical harm, what tasks need careful consideration before involving AI, and what must stay in human hands, such as embodied care and the formative work of sermon preparation.

In plain words: Three buckets for your own use of these tools: Fine to use for assistance. Handle with care. Never hand over, the way you would not hand over praying with someone at a bedside.

Why it matters here: This framework is simple yet flexible enough for a pastor to apply to many aspects of ministry.³⁰

Practice postures

What it is: This is the stance underneath every case: receive before correcting, probe before preaching, preserve agency while naming risk. Related ideas include working with a person rather than to or for them, keeping a human in the loop, and asking "to what end?"

In plain words: Start with the person, not with a verdict about the technology. Receive them, get curious before we preach, and name the risk without taking over.

Why it matters here: This is the tone the whole toolkit is trying to protect: the difference between accompaniment and control.³¹

Three postures toward a new technology

What it is: Christian ethicist Kate Ott names three common Christian responses to a new technology. "It's just a tool" preserves a sense of autonomy but underestimates how much the tool already shapes the person using it. "I'm scared, it's amazing" is the whiplash between fear and hype, often held by the same person in the same week. A third posture tries to live alongside the technology with eyes open, neither outsourcing judgment to it nor pretending it changes nothing.

In plain words: Which of these three do you default to? Naming your own posture is the step before you can help a parishioner examine theirs.

Why it matters here: This gives a caregiver a mirror for their own reaction before responding to someone else's. We can only lead as far as we ourselves have been led.³²

Grief and care

Continuing bonds

What it is: This theory, introduced by Klass, Silverman, and Nickman in 1996, holds that healthy grief often maintains an evolving, meaningful connection with the deceased, rather than requiring a person to let go entirely. It reframes older, detachment-based models, including Freud's, and it shows up in ordinary practices: talking to the person, keeping their objects, marking anniversaries, carrying their values forward.

In the hallway: Grief does not mean cutting the cord with someone who died. Keeping a bond in new ways is normal and healthy. The real question is whether a practice helps someone keep living, or keeps them stuck.

Why it matters here: Talking to the deceased is long recognized as a healthy practice in the grieving process, yet a griefbot can amplify or even distort that practice.³³

Ambiguous loss / arrested grief

What it is: Ambiguous loss, a term coined by therapist Pauline Boss, describes loss without closure. It takes two forms. Someone can be physically gone under unresolved circumstances, described as physical absence with psychological presence, as with a missing person. Or someone can be physically present but changed beyond recognition, described as psychological absence with physical presence, as with dementia. Because the mind craves closure, grief in these situations can get stuck. This is sometimes called frozen grief or grief limbo.

In the hallway: It is grief with no clear goodbye, when someone is gone but not gone, or here but not really here. Without closure, grief can stall. That's a normal response, not a failure.

Why it matters here: A griefbot can manufacture and prolong exactly the condition Boss calls physical absence with psychological presence, keeping the deceased perpetually available and closure perpetually out of reach. This dynamic also points toward the response. When closure is not possible, clinically endorsed help includes ritual, ceremony, and faith-community support, all of which are present in this toolkit.³⁴

Theological anchors

Classical Greek and Latin terms here are anchored in the Santa Clara University Glossary of Theological Terms, an academic resource, copyright SCU, paraphrased and credited, alongside this toolkit's own denominational statements. Santa Clara is Roman Catholic in emphasis and this toolkit is ecumenical,

so we use their glossary for terms the traditions share, and we rely on denominational sources where usage differs.

Imago Dei (image of God)

What it is: This is the conviction that personhood is a divine gift. It grounds each person's unconditional dignity, and it has nothing to do with intelligence or processing power.

In the hallway: Every person carries the image of God. That is where their worth comes from, not how useful or clever they are. A machine, however impressive, does not carry that image.

Why it matters here: This is the theological floor under the whole toolkit. It is why a person is never reducible to data, and why AI is not a person.³⁵

Incarnation / embodiment

What it is: This is the affirmation that God took on flesh, and that to be human is to be in a particular body. A related idea is kenosis, Christ's self-emptying.

In the hallway: God came in a body. Being human means being in a body too. That is why actually showing up for someone still matters in a way a screen never fully can.

Why it matters here: This is the theological reason embodied presence cannot be fully outsourced to an algorithm.³⁶

Formation

What it is: Formation is the slow shaping of a person, including a minister's own voice and vocation, through practice, wrestling, and relationship.

In the hallway: Some things only grow slowly, through the work itself: wrestling with a text, sitting with someone. Efficiency can quietly rob us of the very thing that shapes us.

Why it matters here: This names what "doing more with less" can actually cost, and it anchors the never-delegate tier of AI use.³⁷

Pastoral presence

What it is: Pastoral presence means accompaniment in the messiness of human life, patterned on Christ. It is the core of ministry.

In the hallway: The heart of this work is simply being with people in the hard places. A tool can help at the edges, but it cannot be present for you.

Why it matters here: This is the thing this toolkit is ultimately protecting: the human center that AI can support but never replace.³⁸

Discernment

What it is: Discernment means prayerfully assessing the will of God in a life or situation. It weighs what is faithful, not simply what is permitted or fast.

In the hallway: It is asking not just “can I?” or “is it quick?” but “what is the faithful thing to do here?”

Why it matters here: Discernment is the toolkit’s central verb. It’s why the Four Levels are called vocabulary for discernment, and it is the aim of every framework in this flow.³⁹

Formation tools for the caregiver

These are practices more than terms. If this page should hold concepts only, they may belong better on the Formation page. They are kept here because they function as named, shared references.

Checking your buttons / the 4 B's

What it is: This is a pre-care self-audit. Before entering a conversation, notice your own reaction, fear, fascination, contempt, or feeling left behind, through four lenses: biology, biography, behavior, and backstory.

In plain words: Before a hard conversation, notice what it is stirring in you: your body, your history, your habits, the story you are telling yourself. That way you do not project it onto the person in front of you.

Why it matters here: A caregiver’s unexamined posture toward AI gets projected onto the person seeking care. This tool catches it first.⁴⁰

Three Gates of Speech

What it is: This is a filter for what to share. Is it true? Is it necessary? Is it kind?

In the hallway: Before you pass something along, ask: is it true, is it necessary, is it kind? Especially with all the AI-made content flying around right now.

Why it matters here: This is a memorable, teachable habit for congregations facing misinformation and deepfakes.⁴¹

How This Glossary Is Sourced

There are three kinds of terms in this flow, and three kinds of citation to match.

Technical terms are cited to outside authorities: Stanford HAI (the 2020 definitions from Manning, and the 2023 AI Key Terms Glossary), the Syracuse University and MIT Sloan AI glossaries, peer-reviewed research, and Morgan and Currier's *Won't You Be My Digital Neighbor?* (APCE, 2026). Where several sources agree, that convergence is itself part of the citation.

Borrowed pastoral, clinical, and grief frameworks are cited to their point of origin: Mental Health First Aid for ALGEE, Dearing House for Hearing, Helping, Healing, Klass, Silverman, and Nickman together with Pauline Boss for the grief theory, and Kate Ott for the three postures toward a new technology.

Real-world scam, deepfake, and election-security material is cited to working journalism, government, and nonprofit sources current as of this writing: Axios, NPR, PBS NewsHour, the Centre for Emerging Technology and Security, the National Conference of State Legislatures, the First Amendment Encyclopedia, Carnegie Mellon University, the U.S. Senate Judiciary Committee, Fordham University, Think Global Health, Holly Creek Retirement Community, and Missouri Health and Senior Services. This is fast-moving territory. Names, tools, and even the tells that expose a fake will keep changing faster than any print document can.

Passages describing Rev. Dr. Andy Morgan's own curricula, the AI Literacy for Congregations series, the A.I., Faith, and You decks, the AI Powered Ministry Handout, and related materials, are cited to the specific piece and session that backs the claim, since these are original works with a single identifiable author.

Toolkit coinages are original to the AI Pastoral Toolkit Design Sprint, July 2026. Their authority is internal, so the wording should stay identical across the site, the case studies, and this page.

External glossaries, including Syracuse, MIT Sloan, Santa Clara, Cisco, and Forbes, are copyrighted. Every definition on this page is paraphrased and credited, never copied directly.

Sources and notes

1. John McCarthy, "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence" (1955); Dartmouth College, "Artificial Intelligence (AI) Coined at Dartmouth"; Stanford HAI, "AI Key Terms Glossary" (2023); Melissa Heikkilä, "A Short History of AI, and What It Is (and Isn't)," MIT Technology Review (July 16, 2024); Morgan and Currier, *Won't You Be My Digital Neighbor?* (APCE, 2026); Kate Lucky, "AI Will Shape Your Soul," published as "Imago AI," Christianity Today (October 2023).
2. Stanford HAI, "AI Key Terms Glossary" (2023); Syracuse University AI glossary; Morgan and Currier, *Won't You Be My Digital Neighbor?* (APCE, 2026); Rev. Dr. Andy Morgan, AI Literacy for Congregations: Slide Deck and Participant Guide (2026), session one.
3. Morgan and Currier, *Won't You Be My Digital Neighbor?* (APCE, 2026); Stanford HAI, "AI Key Terms Glossary" (2023); Rev. Dr. Andy Morgan, AI Literacy for Congregations: Participant Guide (2026),

session one.

4. Stanford HAI, "AI Key Terms Glossary" (2023); Syracuse University and MIT Sloan AI glossaries; Morgan and Currier, *Won't You Be My Digital Neighbor?* (APCE, 2026); Rev. Dr. Andy Morgan, *AI Literacy for Congregations: Slide Deck* (2026), session one.
5. Syracuse University AI glossary; Stanford HAI (Manning, 2020).
6. Stanford HAI (Manning, 2020); Syracuse University and MIT Sloan AI glossaries.
7. MIT Sloan and Syracuse University AI glossaries; Stanford HAI (Manning, 2020).
8. Stanford HAI (Manning, 2020); Syracuse University and MIT Sloan AI glossaries.
9. MIT Sloan and Syracuse University AI glossaries; Keith D. Foote, "A Brief History of Large Language Models," *Dataversity* (2023); Rev. Dr. Andy Morgan, *AI Literacy for Congregations: Participant Guide* (2026), session one ("it predicts the next word").
10. Syracuse University and MIT Sloan AI glossaries (prompt, prompt engineering); "priming" per common usage.
11. Ji et al., *Survey of Hallucination in Natural Language Generation*; corroborated by MIT Sloan and Syracuse University.
12. AI Pastoral Toolkit case corpus; Jenny Lyons-Cunha, "AI Companions Are on the Rise, Offering Intimacy That Feels Real — But Is It?" *Built In* (updated November 19, 2025); Dasom Choi, Sunok Lee, Sung-In Kim, Kyungah Lee, Hee Jeong Yoo, Sangsu Lee, and Hwajung Hong, "Unlock Life with a Chat(GPT): Integrating Conversational AI with Large Language Models into Everyday Lives of Autistic Individuals," *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.
13. AI Pastoral Toolkit case corpus; *The New Yorker*; Chloe Veltman, "AI 'Deadbots' Are Persuasive, and Researchers Say They're Primed for Monetization," *NPR* (August 26, 2025); *Mashable*, "What's an AI Deadbot? If You Haven't Seen One Yet, You Will Soon."
14. MIT Sloan AI glossary; *Kore.ai*, "What Is AI Anthropomorphism and Why Does It Matter?"; Margaret Mitchell, chief ethics scientist at Hugging Face, quoted in Jenny Lyons-Cunha, "AI Companions Are on the Rise, Offering Intimacy That Feels Real — But Is It?" *Built In* (updated November 19, 2025).
15. AI Pastoral Toolkit, Day One Synthesis; Caleb Sponheim, "Sycophancy in Generative-AI Chatbots," *Nielsen Norman Group* (January 12, 2024).
16. Rev. Dr. Andy Morgan, *AI Powered Ministry Handout* (2026); Shannon Bond, "How AI Deepfakes Polluted Elections in 2024," *NPR* (December 21, 2024).
17. National Conference of State Legislatures, "Artificial Intelligence (AI) in Elections and Campaigns"; *The First Amendment Encyclopedia*, "Political Deepfakes and Elections"; *Carnegie Mellon University*, "Voters: Here's How to Spot AI Deepfakes That Spread Election-Related Misinformation."
18. U.S. Senate Committee on the Judiciary, Subcommittee on Privacy, Technology, and the Law, "Oversight of AI: Election Deepfakes," hearing, April 16, 2024 (S. Hrg. 118-573); *Centre for Emerging Technology and Security*, "From Deepfake Scams to Poisoned Chatbots: AI and Election Security in 2025"; Shannon Bond, "How AI Deepfakes Polluted Elections in 2024," *NPR* (December 21, 2024).
19. *How AI Can Be Used To Scam Seniors*, *Holly Creek Retirement Community*; Megan Morrone, "How to Talk to Grandma About ChatGPT," *Axios* (November 27, 2025); *How to Help Grandparents Spot AI-Driven Misinformation*, *Fordham University Information Security and Assurance* (Jason Benedict, October 2, 2025).
20. Rev. Dr. Andy Morgan, *AI Literacy for Congregations: Participant Guide* (2026), part two; *Missouri Department of Health and Senior Services*, "Artificial Intelligence Scams"; *Think Global Health*, "Helping Older Adults Navigate AI Scams"; Megan Morrone, "How to Talk to Grandma About ChatGPT,"

Axios (November 27, 2025); PBS NewsHour, "How to Recognize and Block AI-Powered Scam Attempts" (2026), reporting by Paul Solman.

21. Rev. Dr. Andy Morgan, AI Literacy for Congregations: Participant Guide and Slide Deck (2026), part three.
22. Chapman University AI Hub, "Bias in AI"; Syracuse University and MIT Sloan AI glossaries; "Bias in the Bot" (corpus); Rev. Dr. Andy Morgan, AI Literacy for Congregations: Participant Guide and Slide Deck (2026), session two ("not neutral... shaped by our biases, blind spots, and brokenness").
23. Cisco, "What Is a Data Center?"; Forbes, "What Is a Data Center?" (2026); energy and water figures from Adam Zewe, "Explained: Generative AI's Environmental Impact," MIT News (2025); Morgan and Currier, Won't You Be My Digital Neighbor? (APCE, 2026).
24. Church AI Toolkit, First Response.
25. AI Pastoral Sprint, Neutral Case Gallery.
26. AI Pastoral Sprint, Neutral Case Gallery.
27. Mental Health First Aid Action Plan (ALGEE), National Council for Mental Wellbeing, mentalhealthfirstaid.org, accessed 2026; framework originally developed by Kitchener and Jorm.
28. Dearing House Child Advocacy Center (Ponca City, Oklahoma), "hearing, helping, healing" mission; see "Who's Listening?" by Maggi Hutchason, dearinghouse.com.
29. Church AI Toolkit, Clergy Workshop.
30. AI Pastoral Toolkit, Day One Synthesis; the "Sous Chef" pastoral model.
31. AI Pastoral Toolkit, Day One Synthesis.
32. Kate Ott, Christian Ethics for a Digital Society (Rowman and Littlefield, 2019); Kate Ott, "Artificial Intelligence Is Here. Now What?" The Presbyterian Outlook (February 27, 2024); Rev. Dr. Andy Morgan, AI Literacy for Congregations: Participant Guide (2026), part three.
33. Klass, Silverman, and Nickman, Continuing Bonds (1996); accessible overview at The Loss Foundation, "Continuing Bonds Theory in Grief Counselling."
34. Pauline Boss, Ambiguous Loss; accessible overviews at Cleveland Clinic, "What Ambiguous Loss Is and How To Deal With It," and Mayo Clinic Health System, "Coping with Ambiguous Loss."
35. Theological Brief (Genesis 1:26–27); Santa Clara University, "Glossary of Theological Terms"; Kate Lucky, "Imago AI," Christianity Today (October 2023); AI Research Group for the Centre for Digital Culture, Dicastery for Culture and Education of the Holy See, Encountering Artificial Intelligence: Ethical and Anthropological Investigations (2024).
36. Theological Brief; Santa Clara University glossary (kenosis, Philippians 2).
37. Theological Brief; Why Theology Matters (Jannett); Kate Lucky, "Imago AI," Christianity Today (October 2023).
38. Theological Brief; Linda Kinstler, "Can Silicon Valley Find God?" The New York Times (guest essay).
39. Santa Clara University, "Glossary of Theological Terms"; Church AI Toolkit, Clergy Workshop and First Response; Kate Ott, Christian Ethics for a Digital Society (Rowman and Littlefield, 2019); AI Research Group for the Centre for Digital Culture, Dicastery for Culture and Education of the Holy See, Encountering Artificial Intelligence: Ethical and Anthropological Investigations (2024).
40. Church AI Toolkit, Caregiver Self-Reflection.
41. Rev. Dr. Andy Morgan, AI Literacy for Congregations: Participant Guide (2026), part three, "Digital Discernment"; older wisdom tradition, variously attributed.