

Artificial Intelligence Vocabulary for Chaplains

A Glossary of Common Terms

For Chaplains in Healthcare and Institutional Settings

CRITICAL PRIVACY REMINDER

Consumer AI tools (including ChatGPT, Claude, Gemini, and Copilot) are not HIPAA-compliant and are not approved for use with patient information or any protected client information.

Entering patient/care recipient data into a consumer AI tool is a potential HIPAA violation regardless of intent. The information entered may be stored, reviewed, and used to train future versions of the model.

For privacy-protected work, use only AI tools explicitly approved by your company.

Introduction

Artificial intelligence is no longer a distant prospect for healthcare chaplaincy. AI has been a feature of products used for daily administrative tasks like email and document drafting for several years. Today, AI is embedded in electronic health records, documentation platforms, clinical decision support tools, and predictive referral systems. Chaplains who understand AI terminology are better equipped to evaluate these tools ethically, advocate for the care recipients and communities they serve, and participate meaningfully in institutional conversations about AI's benefits and its risks to equity and professional ethics. This document is meant to support chaplains in developing a foundation for interdisciplinary conversations about AI, ethics, and human dignity.

How to Use This Dictionary

Terms in this dictionary are organized into three tiers to help readers navigate by familiarity:

Part One — Foundational Terms: Core concepts accessible to all readers regardless of AI background.

Part Two — Ethical Terms: Terms from critical technology studies and healthcare law that highlight justice concerns and engage questions about integrity in building and using AI.

Part Three — Applied and Contextual Terms: More technical or specialized terms for readers who want deeper literacy.

The **In Context** note with each entry points toward professional chaplaincy application.

Part 1 – Foundational Terms

These terms form the basic vocabulary for understanding what AI is, what it can do, and where it operates in healthcare and institutional settings. They are accessible to readers with no prior AI background.

Algorithm

A set of step-by-step instructions a computer follows to complete a task or solve a problem. In AI, algorithms process data and identify patterns to make decisions or generate outputs. (Not all algorithms are AI-based, and not all automated processes are algorithms).

In Context: *When an AI tool suggests a care resource or flags a patient's risk level, it is following an algorithm, a logic you can ask to have explained. It is sometimes helpful to view the pattern of logic used to create outputs. See also: Predictive Analytics; Black Box.*

Artificial General Intelligence (AGI)

A hypothetical form of AI capable of performing any intellectual task a human can, flexibly and autonomously across domains, without being specifically programmed for each task. AGI does not currently exist. Most AI tools are narrow: they excel at specific tasks and fail outside them. The timeline for whether and when AGI might arrive is a matter of genuine, significant disagreement among experts.

In Context: *Public anxiety about AI often conflates current tools with the hypothetical threat of AGI. Chaplains may encounter people whose fear or enthusiasm is shaped by AGI narratives. Staying current on expert debate, without dismissing either camp, is part of pastoral attentiveness. See also: Doomers and Boomers.*

Artificial Intelligence (AI)

- 1) The simulation of human-like intelligence in computer systems, enabling machines to perform tasks such as recognizing language, making decisions, generating text, and learning from experience.
- 2) The ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience.

In Context: *AI is not a single technology; it is an umbrella term for a wide range of tools and systems. AI is increasingly present in healthcare settings, from predictive analytics in the EMR to voice-transcription tools to documentation assistants. See also: Types of AI Tools; Consumer AI; Enterprise AI; EHR-Embedded AI.*

Automation

The use of technology to perform tasks with minimal human intervention. There are also non-AI algorithms that can adapt to context. AI-powered automation can learn and adjust, while traditional automation follows fixed rules.

In Context: *Some repetitive documentation tasks may be automated with AI. See also: Deskilling.*

Chatbot

A software program designed to simulate conversation with a human user. Simple chatbots follow scripted decision trees. AI-powered chatbots use natural language processing to generate more dynamic, context-sensitive responses. Not all chatbots are AI-based.

In Context: *Some healthcare systems are deploying chatbots for patient check-ins, mental health screening, or care navigation. Chaplains may be asked to evaluate the spiritual or emotional*

appropriateness of these tools: appropriateness in terms of accuracy, sensitivity to religious and cultural identity, and the risks of substituting automated response for human presence.

See also: Consumer AI; Spiritual Care and AI.

Generative AI

A category of AI that creates new content: text, images, audio, code, or video, in response to a prompt. Examples include ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google), and Grok (xAI).

In Context: *Generative AI is the category most immediately relevant to chaplains exploring documentation support, reflection writing, and education development. Outputs require pastoral review and professional judgment before use. Patient data must stay out of consumer generative AI tools. A concern for chaplains is understanding the difference between generative and non-generative AI, so as not to apply blanket skepticism toward every use of AI in healthcare. See also: Consumer AI; Enterprise AI; Types of AI Tools.*

Hallucination

When an AI model generates information that is plausible-sounding but factually incorrect, fabricated, or entirely invented. Hallucinations are often stated with apparent confidence. Hallucinations occur because AI models predict likely word sequences rather than retrieving verified facts. According to studies by AI companies themselves, a significant portion of AI output may be hallucinated. The model has no relationship to truth, only to pattern completion.

In Context: *AI tools may confidently generate inaccurate religious information, fictitious citations, or incorrect details about a patient's faith tradition. Never rely on AI outputs for spiritual care guidance without independent verification. See also: Training Data.*

Large Language Model (LLM)

A neural network trained on massive amounts of text data to understand and generate human language. LLMs such as GPT (OpenAI), Claude (Anthropic), Gemini (Google), and Grok (xAI) are utilized for AI writing, conversation, and summarization tools. They predict the most statistically likely next word or phrase; they do not 'think,' 'understand,' or 'care' in any meaningful sense. The same underlying LLM may power both a consumer product (not HIPAA-compliant) and an enterprise-licensed product (which may be compliant), the model is the engine; compliance status depends on the contract and platform.

In Context: *LLMs are increasingly used in healthcare for clinical documentation, patient communication, and administrative support. They are powerful tools with significant limitations, including bias and the potential for hallucination. See also: Generative AI; Consumer AI.*

Machine Learning (ML)

A subset of AI in which systems learn from data without being explicitly programmed for each task, identifying patterns in training data and applying them to new inputs.

In Context: *Predictive tools in palliative care, such as models that estimate prognosis or identify patients who may benefit from goals-of-care conversations, typically use machine learning. Understanding this helps chaplains interpret how referral triggers are generated. See also: Predictive Analytics; Clinical Decision Support AI.*

Model

The trained AI system itself. The output of a machine learning process applied to a dataset. Different models may produce very different outputs from the same input. There are millions of AI models available.

In Context: *Not all AI tools are equivalent in quality, bias level, or suitability for chaplaincy contexts. Asking 'What model is this tool built on, and how was it trained?' is a legitimate professional question.*

Prompt

The question, instruction, or request input that a user provides to a generative AI system. The quality and specificity of a prompt significantly shapes the quality of the AI's response. **Prompt engineering** refers to the practice and skill of deliberately crafting prompts to guide AI toward more accurate, relevant, and useful outputs.

In Context: *Learning to write clear, specific prompts is a practical AI literacy skill. A vague prompt produces a generic response; a well-crafted prompt that includes context, audience, and purpose produces a more useful output. Chaplains can improve AI-assisted documentation and educational materials by developing these skills.*

Training Data

The large dataset used to teach an AI model which determines what the AI 'knows,' how it reasons, and where its biases originate. For consumer AI tools, training data includes vast amounts of publicly available text from the internet. Additionally, when users interact with consumer AI tools, their conversations may (depending on platform settings and terms of service) be incorporated into future training data.

In Context: *When a chaplain types a patient's spiritual history, diagnosis, or identifying information into a consumer AI tool, that information may become training data for future AI models. The patient did not consent and the institution did not authorize this use of protected health information; in most cases, this constitutes a HIPAA violation. See also: Consumer AI; BAA; Data Retention.*

Part Two — Ethical Terms

These terms address the ethical, justice, and values dimensions of AI use in chaplaincy and healthcare contexts.

Algorithmic Bias

When an AI system produces systematically skewed, incomplete, or discriminatory outcomes as a result of biased, incomplete, or non-representative training data. Because AI learns patterns from historical data, it frequently encodes and amplifies the inequities already present in that data. Real-world examples include: diagnostic tools that underperform for Black patients, leading to higher rates of missed diagnoses; image generation tools that produce more sexualized results for certain demographic groups; and language tools that associate certain names or dialects with negative sentiment. Algorithmic bias is distinct from intentional prejudice; it is structural and often invisible to those deploying the tool. Note, however, that some algorithmic bias is intentional in origin.

In Context: *Algorithmic bias is a patient safety and justice issue. Chaplains are positioned to name it in interdisciplinary care conversations, advocate for patients receiving inequitable AI-influenced care, and ask: 'Has this tool been validated on the population we actually serve?' See also: Racial Bias; Fairness; i-Methodology.*

Black Box

A term describing AI systems with opaque internal decision-making processes. The output can be observed, but the reasoning behind it cannot be easily explained or audited. Conventional algorithms can in principle be explained if source code is available, but how large language models (LLMs) arrive at their outputs is a black box even to the AI companies themselves. Moreover, LLMs are probabilistic: the way a result is computed differs every time. It is extremely difficult to explain why exactly an AI system produced a certain output.

In Context: *When AI influences clinical care decisions affecting your care recipients, asking 'How did the system arrive at this?' is both ethically sound and professionally appropriate. If the answer is unavailable, that is not a technical detail; it is a limit with moral weight. See also: Explainability; Transparency.*

Confidentiality and AI

The risk that information shared with AI tools may be stored, analyzed, retained, or used to train future models, compromising privacy and creating potential legal and ethical violations.

In Context: *This risk is especially acute with consumer-grade AI tools, which are not covered by HIPAA's protections by default. Do not enter identifying patient information into any AI tool that has not been specifically approved by your institution's IT and compliance teams. This includes the standard public-facing versions of ChatGPT, Claude, Gemini, and others. Approval requires a Business Associate Agreement, not merely a vendor's general privacy policy. See also: BAA; Protected Health Information; Training Data; Data Retention.*

Deskilling

The process whereby acquired professional skills atrophy because practitioners increasingly rely on machines to perform tasks that once required human expertise. Related terms are 1) Non-skilling: skills never developed because AI handles the task from the outset in professional practice. 2) Mis-skilling: investing in AI tool use at the expense of other skills essential to professional success. 3) Cognitive offloading: the practice of delegating cognitive tasks (such as recall, analysis, or synthesis) to external tools, including AI, which may accelerate deskilling when used habitually.

In Context: *Deskilling is a professional formation concern. For chaplains, mis-skilling may involve investing in AI tool use at the expense of relational attunement, theological depth, and unmediated clinical judgement. If chaplains consistently use AI to draft reflections, formulate assessments, prepare public religious/spiritual gatherings, or identify spiritual distress patterns, the question must be asked: What capacities are we maintaining? What are we allowing to atrophy?*

Digital Divide

The gap between those who have access to digital technology and those who do not. This gap is intersected by race, socioeconomic status and education, age, geography, and disability. The digital divide also encompasses achievement gaps and unequal access to AI-enhanced education and services.

In Context: *Not all patients, and not all chaplains, have equal access to or comfort with digital tools.*

Ethical AI

AI designed, deployed, and governed in ways that are fair, transparent, accountable, and respectful of human dignity. Ethical AI frameworks address bias, privacy, explainability, safety, and the protection of vulnerable populations.

In Context: *Interfaith chaplaincy's core commitments to human dignity, justice, and the protection of the vulnerable align directly with the principles of ethical AI. Chaplains are natural and credible voices in institutional AI ethics conversations.*

Explainability (Explainable AI)

The degree to which the reasoning behind an AI output can be understood and communicated to humans, allowing clinicians, patients, and administrators to evaluate and trust the system's recommendations.

In Context: *When an AI tool generates a risk score or care recommendation, explainability allows you to evaluate whether it reflects your patient's actual story, or a statistical average that may not apply to them. See also: Black Box; Transparency.*

Fairness (in AI)

The principle that AI systems should produce equitable outcomes across demographic groups. Fairness in AI is contested: different mathematical definitions of fairness can conflict, meaning a system can be 'fair' by one measure while remaining unjust by another.

In Context: *An AI system that performs 'fairly' on average may still produce inequitable outcomes for specific communities your chaplaincy serves. Aggregated fairness can mask individual injustice. See also: Algorithmic Bias; Racial Bias.*

Ghost Work

The invisible, often underpaid human labor that underlies generative AI systems, particularly the work of content moderation, data labeling, and output review. Ghost workers review vast quantities of violent, sexually exploitative, or otherwise harmful material in order to train AI systems not to reproduce it. This labor is frequently performed by workers in precarious economic situations and is structurally invisible in public narratives about AI. Term coined by Mary L. Gray and Siddharth Suri in *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass* (2019).

In Context: *Ghost work is a pastoral and justice concern. When institutions celebrate AI efficiency, chaplains can ask who bears the hidden cost.*

Human-in-the-Loop (HITL)

A design principle requiring that a human reviews, approves, or can override AI-generated decisions at critical points. Human-in-the-loop systems are considered safer for high-stakes environments. The concept is also relevant in considering when AI operates autonomously; for example, some algorithms deliver orders without human oversight.

In Context: *Human-in-the-loop design is especially important in healthcare, where AI-generated decisions can affect patient care. Ethical chaplaincy practice demands that the chaplain remain responsible for their work as the human in the loop. AI may assist with documentation, resource identification, or communication drafting, but spiritual judgment, accountability, and relational care remain irreducibly human responsibilities. See also: Deskilling.*

i-Methodology

A design practice, documented in science and technology studies, in which developers assume their own interests, experiences, and contexts align with those of a universal user. Most AI tools have been designed predominantly by young, Western, male, technically trained engineers whose default assumptions shape the systems they build. The scope of i-methodology extends beyond interface design to include the assumption that AI should be everywhere in the first place. i-Methodology is the structural opposite of participatory design, in which intended users are meaningfully involved in shaping the systems built for them. Term coined by Nelly Oudshoorn, Els Rommes, and Marcelle Stienstra in "Configuring the User as Everybody: Gender and Design Cultures in Information and Communication Technologies" (2004).

In Context: *Administrators eager to incorporate AI in numerous workflows can ask whether their own motivations (enthusiasm for novel technologies, economic incentives, etc.) align with care recipient desires for more or less AI exposure in their medical care. See also: Algorithmic Bias; Ghost Work.*

Informed Consent (AI Context)

The process of ensuring that patients understand how AI tools are being used in their care, including what data is collected, how it is used, and what decisions it influences.

In Context: *Patients have a right to know when AI is shaping their care. Chaplains can advocate for that transparency as part of their patient rights role. See also: Transparency; Confidentiality and AI.*

Racial Bias (in AI)

A specific and well-documented form of algorithmic bias in which AI systems produce less accurate, more harmful, or inequitable outputs for people of color. Documented in clinical risk scoring, imaging analysis, dermatology AI, pain assessment tools, and facial recognition.

In Context: *Racial bias in clinical AI is a patient safety issue. Chaplains are positioned to name AI bias as a justice concern in interdisciplinary settings, and to advocate for patients whose care may be compromised by tools not validated on them. See also: Algorithmic Bias; i-Methodology.*

Spiritual Care and AI

The evolving relationship between professional spiritual care and the increasing use of AI tools in healthcare, including AI-assisted documentation, predictive referral tools, and chatbot-based supplements to mental health and spiritual support.

In Context: *Opinions vary in the professional chaplaincy community about appropriate use cases and the extent of human involvement that various AI applications require. AI cannot provide spiritual care. Through supporting chaplains in administrative, analytical, and communicative tasks, AI can potentially augment and amplify spiritual care provider presence. Clarity about this distinction is foundational to continuing conversations about ethical AI use in clinical spiritual care contexts.*

Technocratic (approach to AI ethics)

An approach to AI ethics that frames moral and social concerns as technical problems to be solved by experts, metrics, and system design. The technocratic approach assumes that ethical issues can be adequately addressed through better data, better models, and better guardrails. It prioritizes optimization, efficiency, and capability, and tends to locate authority over AI ethics in engineers, data scientists, and technical governance frameworks rather than in affected communities, ethicists, or practitioners.

In Context: *Much of the dominant discourse around AI ethics is technocratic in character. Recognizing technocratic framing helps chaplains identify when ethical questions are being narrowed into engineering problems, and advocate for the kinds of human and moral considerations that metrics cannot capture. See also: Ethical AI; i-Methodology; Algorithmic Bias; Human-in-the-Loop.*

Transparency (AI Transparency)

The principle that AI systems, their developers, and their institutional deployers should be open about how AI tools work, what data they use and store, what decisions they influence, and what their known limitations are.

In Context: *Patients have a right to know when AI is involved in their care. Chaplains modeling transparency about their own AI tool use and advocating for system-wide transparency supports institutional trust and professional integrity. See also: Explainability; Informed Consent.*

Part Three — Applied and Contextual Terms

These terms describe how AI operates in real-world healthcare and institutional contexts, including legal frameworks, tool types, and clinical applications, as well as more technical and specialized concepts for chaplains who want deeper literacy or who are involved in institutional AI governance, research, or education.

AI Disclosure

A transparent statement indicating that AI tools were used in the creation or editing of a document, communication, or professional work product. AI disclosure typically describes what AI was used for, the extent of its contribution, and the degree of human review and oversight applied to the final output. Many academic journals, professional organizations, and institutions now require or recommend AI disclosure. Disclosure protects professional integrity and models the transparency expected of ethical AI use.

In Context: *Chaplains who use AI to assist with documentation, educational materials, or professional writing should be prepared to disclose that use when asked, and to follow any institutional or professional standards that apply. Disclosure is an act of integrity, not a confession of inadequacy. This document includes an authorship and AI disclosure statement. See also: Responsible AI; Ethical AI; Transparency.*

API (Application Programming Interface)

A technical interface that allows different software applications to communicate and share data or capabilities. In the context of AI, an API is what allows ChatGPT, Claude, or other consumer AI models to be integrated into other apps and services without the user necessarily knowing that 'ChatGPT is speaking' in the background. This means AI can be embedded in tools that don't look like AI chatbots at all.

In Context: *Understanding that AI can be integrated via API into products that appear unrelated to AI helps chaplains recognize that AI may be present in tools across their institutional environment, not just in obvious chat interfaces. See also: EHR-Embedded AI; Types of AI Tools.*

Business Associate Agreement (BAA)

A legally binding contract required under HIPAA between a healthcare organization and any vendor or third party that handles protected health information on its behalf. Without a signed BAA, a vendor is not authorized to access, store, or process care recipient data. Note that some companies offer enterprise solutions with improved data protection including enterprise versions of tools whose consumer versions lack these protections.

In Context: *A BAA is the minimum legal threshold for using any AI tool with care recipient data. Consumer-grade AI tools such as the standard versions of ChatGPT, Claude, or Gemini do not have BAAs with individual users. See also: Consumer AI; Protected Health Information (PHI).*

Clinical Decision Support (CDS) AI

AI tools designed to assist clinicians in making care decisions by analyzing patient data and generating recommendations. CDS AI may flag abnormal lab values, suggest diagnoses, predict deterioration, identify medication interactions, or prioritize clinical concerns. CDS tools are typically embedded in or integrated with EHR systems. Examples include sepsis predictors and fall risk scoring algorithms. CDS AI is usually subject to regulatory oversight and institutional validation requirements.

In Context: *CDS AI operates in the background of many clinical environments. It may generate flags that prompt a chaplaincy referral, risk scores visible in a patient's chart, or care pathways being followed by the team. These recommendations are algorithmically generated and can carry all the biases of their training data. See also: EHR-Embedded AI; Predictive Analytics; Algorithmic Bias.*

Consumer AI (Public AI / General-Purpose AI)

AI tools available to the general public, typically through a web browser or app, without requiring institutional licensing or approval. Examples include ChatGPT (OpenAI), Claude.ai (Anthropic), Gemini (Google), and Copilot (Microsoft). Information entered into consumer AI tools may be used to train future versions of the model unless users have configured privacy settings. Some platforms offer a 'temporary chat' feature that limits data storage.

In Context: *A foundational rule: consumer AI tools are not designed for healthcare, are not HIPAA-compliant by default, and a temporary chat feature does not make a tool HIPAA-compliant. Consumer AI tools are appropriate for non-patient work: professional development, writing assistance, research. They are not appropriate for any use involving patient information. The chaplain who pastes a patient's spiritual assessment into ChatGPT to help polish the language has potentially violated HIPAA, contributed to training data, and created an unreportable breach. For patient-related work, use only tools explicitly approved by your institution. See also: Enterprise AI; BAA; Data Retention; Training Data.*

Data

Raw information including text, numbers, images, audio used by AI systems to learn, train, and generate outputs. The quality, diversity, and source of data directly determine the quality and fairness of AI outputs.

In Context: *Chaplaincy data, including spiritual assessment information and care encounter notes, may be included in health system datasets used to train predictive models. Understanding what data your documentation generates is part of ethical practice. See also: Training Data; Unstructured Data.*

Data Minimization

The principle that only the minimum necessary information should be collected, used, or shared for any given purpose. Data minimization is codified in multiple privacy frameworks and regulations.

In Context: *In healthcare, HIPAA codifies data minimization as a requirement for the handling of protected health information. When using AI tools for any work-adjacent task, ask: Does the AI need the patient's name, diagnosis, or facility to help me with this? Usually, the answer is no. Describe care scenarios in general terms without identifying the person. Do not use AI tools to attempt to minimize or de-identify data; use approved institutional processes. See also: De-identification; Protected Health Information; Consumer AI.*

Data Retention

The policies governing how long data is stored by an AI system or platform, and what is done with it before deletion. Consumer AI tools may store conversations indefinitely and use them for model training. Some platforms offer temporary chat features that limit storage. Enterprise tools are governed by institutional contracts.

In Context: *Note that a temporary chat feature does not constitute HIPAA compliance. Before using any AI tool, ask: How long does this platform retain what I type? Can it be used to train future models? For consumer tools, assume data persists. See also: Training Data; Consumer AI; BAA.*

Doomers and Boomers (Techno-Pessimists and Techno-Utopians)

Informal terms for two dominant ideological postures in public discourse about AI. Techno-utopians ('boomers' or 'accelerationists') believe AI will broadly benefit humanity. Techno-pessimists ('doomers') believe AI poses serious existential risks and that unchecked development is reckless. Many serious thinkers, including some who built the AI empire, are now in the 'doomer' camp, and many intelligent people remain 'boomers.' Both camps carry blind spots: utopians minimize risk, while pessimists sometimes put excessive emphasis on catastrophic ends at the expense of more immediate concerns like algorithmic bias and labor displacement.

In Context: *Naming these camps helps chaplains identify which assumptions they carry uncritically and what point on the spectrum of AI enthusiasm colleagues or institutions occupy.. See also: Responsible AI; AGI.*

EHR-Embedded AI

Artificial intelligence tools built into or directly integrated with an institution's electronic health record (EHR) system. EHR-embedded AI operates on institutional data and is governed by vendor contracts and compliance review. Examples include smart documentation assistants, ambient clinical intelligence, and AI-generated care summaries. Because EHR-embedded AI is built into the system, clinicians may not always realize they are interacting with AI-generated content. These tools may appear as a recommended action, a pre-populated field, or a risk flag with no visible AI label.

In Context: *EHR-embedded AI is not voluntary, chaplains interact with it simply by using their documentation system. Key concerns: Does AI-generated content appear in charts I am reading and acting on? Are spiritual care notes being fed into AI-driven summaries? Is the AI flagging my care recipients for care I don't know about? Knowing that AI is present in your documentation environment, and asking where, is part of professional competency. See also: Clinical Decision Support AI; Predictive Analytics; Unstructured Data.*

Enterprise AI

AI tools licensed by an organization specifically for institutional use. These include contractual protections for patient data such as a Business Associate Agreement and HIPAA compliance. Enterprise AI tools are vetted, approved, and monitored by institutional IT and compliance teams.

In Context: *Enterprise AI tools are not the same as the consumer-facing versions of the same products: an institution's licensed Microsoft Copilot may be HIPAA-compliant; the free version on a personal device is not. Chaplains should know which AI tools their institution has explicitly approved for use with patient data, and use only those tools. See also: Consumer AI; BAA; EHR-Embedded AI.*

Fine-Tuning

The process of taking a pre-trained AI model and further training it on a smaller, specialized dataset to improve performance for specific tasks or populations.

In Context: *Healthcare institutions may fine-tune general AI models on clinical data. Whether spiritual care data (care recipient assessments, encounter notes, chaplaincy documentation) is included in institutional fine-tuning datasets is a question chaplains should raise in AI governance conversations. See also: Training Data; Data.*

Model Card / System Card

Published documents that reputable AI developers make publicly available to describe how an AI model was built and how it is intended to be used. A model card describes the model itself: what it was trained on, what tasks it is designed to perform, its known limitations, and performance across demographic groups. A system card describes the real-world deployment of the model: how it was tested for safety, what risks were identified, and what safeguards are in place. Both documents are intended to support transparency and informed evaluation of AI tools before adoption.

In Context: *When evaluating an AI tool for institutional or professional use, asking whether a model card or system card is publicly available is a concrete due-diligence step. These documents can reveal whether the model was tested on populations relevant to your patients and community, and what the developer acknowledges as known risks. See also: Transparency; Explainability; Validation; Ethical AI.*

Natural Language Processing (NLP)

A branch of computer science focused on enabling computers to understand, interpret, and generate human language. NLP is broader than AI; it encompasses non-AI techniques as well, but many current NLP applications are AI-powered. It powers voice recognition, sentiment analysis, chatbots, and automated text summarization.

In Context: *In healthcare settings, NLP tools may scan chaplaincy notes for patterns or generate summaries for the interdisciplinary team. Clinical notes containing spiritual distress language may trigger AI-driven alerts or be included in automated reports. See also: Sentiment Analysis; Unstructured Data; EHR-Embedded AI.*

Predictive Analytics

The use of statistical algorithms and machine learning to forecast future events based on historical data. Not all predictive analytics are AI-based.

In Context: *In healthcare, predictive analytics may identify patients at risk for deterioration, readmission, or unmet palliative care needs. Chaplains may receive referrals triggered or generated by predictive analytics tools. See also: Clinical Decision Support AI; EHR-Embedded AI; Algorithmic Bias.*

Protected Health Information (PHI)

Any information that can be used to identify a care recipient and that relates to their past, present, or future health condition, healthcare treatment, or payment for healthcare services. HIPAA identifies 18 specific types of PHI identifiers, including name, geographic data smaller than state level, dates related to the individual, phone numbers, email addresses, Social Security numbers, medical record

numbers, and photographs. PHI is protected regardless of the medium in which it appears. Entering any PHI into a non-HIPAA-compliant AI tool is a potential HIPAA violation, regardless of intent.

In Context: *When using AI tools for any purpose that touches patient care, default to de-identification. If you are unsure whether something is PHI, treat it as PHI. See also: De-identification; Data Minimization; Consumer AI.*

Sentiment Analysis

A technique that identifies and categorizes emotional tone in text, typically as positive, negative, or neutral.

In Context: *In healthcare, sentiment analysis is applied to patient feedback, clinical notes, and communications. Spiritual distress, grief, lament, and existential suffering do not map neatly onto positive/negative categories. Chaplains should be prepared to challenge sentiment analyses of spiritual care narratives that may miss or misrepresent complex theological and emotional realities. See also: NLP; Unstructured Data.*

Terms of Service (AI Platforms)

The legal agreement between a user and an AI platform that governs how the platform can use data provided by the user. For most consumer AI tools, terms of service explicitly permit the company to use conversations for model training, safety review, and product improvement. Terms can be changed unilaterally by the company, often with minimal notice.

In Context: *Terms of service exist to protect the company, not the user. Consumer AI terms of service are incompatible with the handling of patient information. Even if you could understand and monitor the terms, they cannot replace a HIPAA-compliant BAA. See also: BAA; Consumer AI; Data Retention..*

Token

The basic unit of text that AI language models process, roughly equivalent to a word or word fragment. Models have token limits for both input and output.

In Context: *Token limits explain why AI tools may truncate long documents or forget earlier conversation.*

Unstructured Data

Data that does not fit into predefined categories, such as free-form text, audio recordings, images, or narrative documents, as opposed to structured data in databases or spreadsheets. AI tools are increasingly capable of analyzing unstructured data.

In Context: *In healthcare, unstructured data includes clinical notes, chaplaincy assessments, and free-text patient communications. Chaplaincy narratives are rich unstructured data. AI may mine these notes for patterns or generate summaries. Chaplains should be intentional about how they document, knowing that their language may be processed in ways they did not originally anticipate. See also: NLP; EHR-Embedded AI; Training Data.*

Validation (Model Validation)

The process of testing an AI model against real-world data to evaluate whether it performs accurately and equitably before deployment. Rigorous validation across demographic subgroups is an indicator of AI tool quality.

In Context: *Before an AI tool is deployed to generate spiritual care referrals or influence patient care decisions, it should be validated in the population it will serve. Asking whether a tool has been validated and for what population(s), is a legitimate institutional question. See also: Algorithmic Bias.*

Vendor Agreement (AI Vendor Contract)

The contract between a healthcare organization and an AI tool provider governing terms of use, data handling, security standards, liability, and compliance requirements. A HIPAA-compliant agreement includes a BAA. Vendor agreements are negotiated by institutional leadership, not by individual users, who do not have the authority to modify them.

In Context: *Individual practitioners cannot make a non-compliant tool compliant by being careful with it. If the institution has not executed a proper vendor agreement for a given AI tool, that tool should not be used for patient-related work. See also: BAA; Enterprise AI; Consumer AI.*

Workflow Integration

The degree to which an AI tool is embedded into existing work processes, operating within the normal flow of tasks rather than as a separate, deliberate step.

In Context: *In clinical settings, workflow integration may include generating referrals, populating templates, or flagging patients without a separate action by the clinician. Understanding what is automated in your workflow, and what is not, is essential to maintaining intentional, person-centered practice. See also: EHR-Embedded AI.*

Zero-Shot Learning

An AI model's ability to perform a task it was never explicitly trained on, by generalizing from related knowledge.

In Context: *When you ask an AI to generate a care note or reflection for a faith tradition it was not specifically trained on, you are relying on zero-shot generalization. Results require careful pastoral review, especially for minority or non-Western traditions.*

Authorship and AI Disclosure

This document was developed through iterative collaboration between members of the APC ETHOS Taskforce and Claude (Anthropic), a large language model AI assistant. Claude served as a drafting and research synthesis tool across multiple working sessions. All content was reviewed, substantially revised, and approved by the human contributors, who retain full intellectual responsibility for the final text.

Artificial Intelligence Vocabulary for Chaplains

Bibliography

AI-Assisted Authorship Disclosure: This document was developed through iterative collaboration between members of the APC ETHOS Taskforce and Claude (Anthropic), a large language model AI assistant. All content was reviewed, substantially revised, and approved by the human authors, who retain full intellectual responsibility for the final text.

Access dates reflect March 2026.

Bibliography

Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. "Machine Bias." ProPublica, May 23, 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

Anthropic. "Claude." Accessed March 2026. <https://www.anthropic.com/claude>. Note on AI-assisted authorship: Claude (claude.ai, Anthropic, model versions Claude Sonnet 4.5 and Claude Sonnet 4.6) served as a primary drafting and research synthesis tool throughout the development of this document. Conversations were conducted by members of the APC ETHOS Taskforce between 2024 and 2026. Draft content, term definitions, chaplaincy considerations, setting guides, and document structure were generated through iterative human-AI collaboration and subjected to professional editorial review by the human authors prior to publication. All intellectual responsibility for the final text rests with the human authors. This disclosure follows emerging Turabian and Chicago guidance on AI-assisted authorship.

Anthropic. "Constitutional AI: Harmlessness from AI Feedback." Anthropic Research, December 15, 2022. <https://www.anthropic.com/research/constitutional-ai-harmlessness-from-ai-feedback>.

Association for Clinical Pastoral Education (ACPE). "ACPE Standards and Manuals." Accessed March 2026. https://www.acpe.edu/ACPE/Accreditation/ACPE_Standards.aspx.

Association of Professional Chaplains (APC). "Standards of Practice for Professional Chaplains." Accessed March 2026. <https://www.professionalchaplains.org/content.asp?pl=198&contentid=198>.

Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), 610–623. New York: Association for Computing Machinery, 2021. <https://doi.org/10.1145/3442188.3445922>. Note: The fourth author is listed as "Shmargaret Shmitchell" — the anonymized pseudonym used by the author in the published paper. This is the correct citation form as it appears in the published proceedings.

Braverman, Harry. Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century. New York: Monthly Review Press, 1974.

Broussard, Meredith. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, MA: MIT Press, 2018.

- Buolamwini, Joy, and Timnit Gebru. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research* 81 (2018): 1–15. <https://proceedings.mlr.press/v81/buolamwini18a.html>.
- Centers for Medicare and Medicaid Services (CMS). "HIPAA Basics for Providers: Privacy, Security, and Breach Notification Rules." Accessed March 2026. <https://www.cms.gov/Outreach-and-Education/Medicare-Learning-Network-MLN/MLNProducts/Downloads/HIPAAPrivacyandSecurity.pdf>.
- CNET. "ChatGPT Glossary: 48 AI Terms That Everyone Should Know." CNET, updated December 2024. <https://www.cnet.com/tech/services-and-software/chatgpt-glossary/>.
- Copeland, B.J. "Artificial Intelligence." *Encyclopædia Britannica*. Last updated April 2, 2026. <https://www.britannica.com/technology/artificial-intelligence>. Note: The definition of artificial intelligence appearing in the main dictionary entry ("The ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings...") is drawn from this source and should be attributed accordingly in the text.
- Dafoe, Allan. "AI Governance: A Research Agenda." Future of Humanity Institute, University of Oxford, 2018. <https://www.fhi.ox.ac.uk/wp-content/uploads/GovAIAgenda.pdf>.
- Department of Defense (DoD). "DoD Directive 1304.19: Appointment of Chaplains for the Military Departments." Washington, DC: DoD, 2004.
- Department of Defense (DoD). "DoD AI Ethics Principles." February 24, 2020. <https://www.defense.gov/News/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>.
- Department of Defense (DoD). "Responsible AI Strategy and Implementation Pathway." June 2022. https://www.ai.mil/docs/RAI_Strategy_and_Implementation_Pathway_6-21-22.pdf.
- Diakopoulos, Nicholas. "Algorithmic Accountability: Journalistic Investigation of Computational Power Structures." *Digital Journalism* 3, no. 3 (2015): 398–415. <https://doi.org/10.1080/21670811.2014.976411>.
- Electronic Frontier Foundation (EFF). "AI and Criminal Justice." Accessed March 2026. <https://www.eff.org/issues/ai-criminal-justice>.
- Epic Systems. "Epic AI." Accessed March 2026. <https://www.epic.com/software/epic-ai/>.
- European Commission. "Ethics Guidelines for Trustworthy AI." High-Level Expert Group on Artificial Intelligence. Brussels: European Commission, April 8, 2019. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>.
- Family Policy Compliance Office, U.S. Department of Education. "FERPA General Guidance for Students." Accessed March 2026. <https://studentprivacy.ed.gov/resources/ferpa-general-guidance-students>.
- Floridi, Luciano, Josh Cowls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, et al. "An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations." *Minds and Machines* 28 (2018): 689–707. <https://doi.org/10.1007/s11023-018-9482-5>.
- Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. "Datasheets for Datasets." *Communications of the ACM* 64, no. 12 (2021): 86–92. <https://doi.org/10.1145/3458723>.

- Google DeepMind. "Gemini: A Family of Highly Capable Multimodal Models." Technical Report. Google DeepMind, December 2023. https://storage.googleapis.com/deepmind-media/gemini/gemini_1_report.pdf.
- Gray, Mary L., and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt, 2019.
- Green, Ben. "The Contestation of Tech Ethics: A Sociotechnical Approach to Technology Ethics in Practice." *Journal of Social Computing* 2, no. 3 (September 2021). <https://doi.org/10.2139/ssrn.3859358>.
- Health and Human Services Office for Civil Rights (HHS OCR). "Summary of the HIPAA Privacy Rule." Accessed March 2026. <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html>.
- IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. *Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems*, First Edition. Piscataway, NJ: IEEE, 2019. <https://standards.ieee.org/industry-connections/ec/ead-v1/>.
- Joint Commission. "Quick Safety 59: Ethical Use of Artificial Intelligence in Health Care." Oakbrook Terrace, IL: The Joint Commission, 2021. <https://www.jointcommission.org/resources/news-and-multimedia/newsletters/newsletters/quick-safety/quick-safety-59-ethical-use-of-artificial-intelligence-in-health-care/>.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Cass R. Sunstein. "Human Decisions and Machine Predictions." *Quarterly Journal of Economics* 133, no. 1 (2018): 237–293. <https://doi.org/10.1093/qje/qjx032>.
- McKinsey Global Institute. "The Economic Potential of Generative AI: The Next Productivity Frontier." McKinsey & Company, June 2023. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>.
- Microsoft. "Microsoft Copilot." Accessed March 2026. <https://www.microsoft.com/en-us/microsoft-copilot>.
- Mitchell, Margaret, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. "Model Cards for Model Reporting." In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT '19)*, 220–229. New York: ACM, 2019. <https://doi.org/10.1145/3287560.3287596>.
- National Institute of Standards and Technology (NIST). "AI Risk Management Framework (AI RMF 1.0)." NIST AI 100-1. Gaithersburg, MD: NIST, January 2023. <https://doi.org/10.6028/NIST.AI.100-1>.
- Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York: New York University Press, 2018.
- Nuance Communications (Microsoft). "Dragon Ambient eXperience (DAX)." Accessed March 2026. <https://www.nuance.com/healthcare/ambient-clinical-intelligence.html>.
- Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." *Science* 366, no. 6464 (2019): 447–453. <https://doi.org/10.1126/science.aax2342>.
- O'Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishers, 2016.
- OpenAI. "ChatGPT." Accessed March 2026. <https://openai.com/chatgpt>.
- OpenAI. "GPT-4 Technical Report." ArXiv preprint arXiv:2303.08774, March 27, 2023. <https://arxiv.org/abs/2303.08774>.

- Oudshoorn, Nelly, Els Rommes, and Marcelle Stienstra. "Configuring the User as Everybody: Gender and Design Cultures in Information and Communication Technologies." *Science, Technology, & Human Values* 29, no. 1 (2004): 30–63. <https://doi.org/10.1177/0162243903259190>.
- Partnership on AI. "Responsible Practices for Synthetic Media: A Framework for Collective Action." Accessed March 2026. <https://partnershiponai.org/wp-content/uploads/2023/02/PAI-Synthetic-Media-Framework.pdf>.
- Prison Policy Initiative. "AI and Criminal Justice." Accessed March 2026. https://www.prisonpolicy.org/research/artificial_intelligence/.
- Raji, Inioluwa Deborah, and Joy Buolamwini. "Actionable Auditing: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products." In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 429–435. New York: ACM, 2019. <https://doi.org/10.1145/3306618.3314244>.
- Raji, Inioluwa Deborah, Morgan Klaus Scheuerman, and Razvan Amironesei. "You Can't Sit With Us: Exclusionary Pedagogy in AI Ethics Education." In *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 515–525. New York: ACM, 2021. <https://doi.org/10.1145/3442188.3445914>.
- Russell, Stuart. *Human Compatible: Artificial Intelligence and the Problem of Control*. New York: Viking, 2019.
- Selbst, Andrew D., Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. "Fairness and Abstraction in Sociotechnical Systems." In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT '19)*, 59–68. New York: ACM, 2019. <https://doi.org/10.1145/3287560.3287598>.
- Shneiderman, Ben. *Human-Centered AI*. Oxford: Oxford University Press, 2022.
- Society for Human Resource Management (SHRM). "Using Artificial Intelligence in the Workplace." Accessed March 2026. <https://www.shrm.org/topics-tools/tools/toolkits/using-artificial-intelligence-workplace>.
- Turing, Alan M. "Computing Machinery and Intelligence." *Mind* 59, no. 236 (1950): 433–460. <https://doi.org/10.1093/mind/LIX.236.433>.
- U.S. Department of Health and Human Services. "Guidance on HIPAA and Cloud Computing." Accessed March 2026. <https://www.hhs.gov/hipaa/for-professionals/special-topics/cloud-computing/index.html>.
- U.S. Department of Health and Human Services. "The HIPAA Privacy Rule." Accessed March 2026. <https://www.hhs.gov/hipaa/for-professionals/privacy/index.html>.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. "Attention Is All You Need." In *Advances in Neural Information Processing Systems* 30 (NeurIPS 2017), 5998–6008. <https://arxiv.org/abs/1706.03762>.
- Verbeek, Peter-Paul. *Moralizing Technology: Understanding and Designing the Morality of Things*. Chicago: University of Chicago Press, 2011.
- West, Sarah Myers, Meredith Whittaker, and Kate Crawford. "Discriminating Systems: Gender, Race, and Power in AI." *AI Now Institute*, April 2019. <https://ainowinstitute.org/discriminatingystems.html>.
- Winner, Langdon. "Do Artifacts Have Politics?" *Daedalus* 109, no. 1 (1980): 121–136. <https://www.jstor.org/stable/20024652>.
- xAI. "Grok." Accessed March 2026. <https://x.ai/grok>.

Zuboff, Shoshana. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: PublicAffairs, 2019.

Artificial Intelligence Vocabulary for Chaplains

Appendix A: Quick Reference — Types of AI Tools in Healthcare Settings

Type	Examples	HIPAA Compliant?	Use With Patient Data?	Chaplain Encounters How?
Consumer / General-Purpose AI	ChatGPT, Claude.ai, Gemini, personal Copilot	No	NEVER	Personal use, writing assistance, research, non-patient tasks only
Enterprise / Institutional AI	Epic AI, Azure OpenAI (enterprise), institutional Copilot	Yes, if BAA in place	Only as approved by institution	Via approved institutional tools
EHR-Embedded AI	Epic predictive models, AI-generated chart prompts, ambient recording documentation tools	Yes (institutional)	Yes — you use this by charting	Automatically, within charting workflow
Medical / Specialist AI	Aidoc (radiology), Nuance (transcription), sepsis predictors	Yes (institutional)	Yes — operates on clinical data	Via outputs in charts or clinical conversations
Departmental / Workflow AI	Department-built or adopted tools for documentation, scheduling, communication	Verify before use	Only if institutionally approved	Via departmental adoption — verify compliance

Note: Some companies offer enterprise versions of their consumer AI tools with improved data protection. Consumer and enterprise versions of the same product are not interchangeable. Always verify which version your institution has approved.

About the APC Ethical Technology for Human-Oriented Spiritual Care (ETHOS) Task Force

The Association of Professional Chaplains Ethical Technology for Human-Oriented Spiritual Care (ETHOS) task force is a multidisciplinary team of chaplains, clinical ethicists, educators, researchers, and technology professionals dedicated to the responsible integration of artificial intelligence within spiritual care.

The taskforce seeks to promote the ethical, compassionate, and human-centered use of emerging technologies in ways that honor human dignity, support professional spiritual care practice, and uphold the values of presence, trust, confidentiality, and equity. The APC ETHOS task force emphasizes that artificial intelligence is not a substitute for the relational, compassionate, and meaning-centered care provided by spiritual care professionals. We encourage readers to use this resource as a foundational reference while continuing to engage critically, ethically, and thoughtfully with emerging technologies.

APC Ethical Technology for Human-Oriented Spiritual Care (ETHOS) Task Force

- Eric Ayala, D.Min., HEC-C, BCC
- Sara Bozeman, MS
- Ali Candir, M.A., BCC
- Nicholas Collura, D.Min., BCC
- Jill Fisk, MATM
- David Fleenor, PhD, M.Div., BCC, ACPE Certified Educator
- Paul Galchutt, M.Div., MPH, BCC
- Kristen "Jo" Nivling Izquierdo, D.Bioethics, M.Div., HEC-C
- Justin Martin, M.Div., BCC (*Co-Facilitator and APC/BCCI Board Member*)
- Stefano Mugnaini, M.Div., MS, RCC, HEC-C
- Terra Tindle Williams, PhD, M.Div.
- Fabian Winiger, PhD, MSc
- Sheri Winesett Campbell, M.Div., BCC (*Co-Facilitator and APC/BCCI Board Member*)

For more information about the Association of Professional Chaplains and its work in ethical technology and spiritual care, visit www.apchaplains.org.