PLAIN-LANGUAGE PRIMER · AI TERMS

You keep hearing the words. *Here is what they mean.*

Every AI term you hear in a vendor call, in one table — plus where HIPAA comes in and which parts of this you can safely ignore. Written for independent practices, not for health systems with an informatics department.

What this is

Every AI term you are likely to hear in a vendor call, defined in one table, with the difference between them stated plainly. Then where HIPAA actually applies, and what open, closed and open-weight mean for who holds your data.

How to use it

Read it once before the first vendor conversation. Keep it beside you during the call and check any term you are asked to accept.

Here is what you have heard, and here is what it means.

The words, in one place

THE TERM	WHAT IT ACTUALLY IS	WHAT IT MEANS IN YOUR PRACTICE
Large language model (LLM)	Software that has read enormous amounts of text and predicts the most plausible next words. The model itself looks nothing up, though a product can wire it to your documents.[1]	The engine under your scribe and your portal-message drafter. It predicts rather than inherently retrieves, which explains most of the failures below.
Frontier model	A policy term for models at the leading edge of capability or scale.	Not a purchasing category for a practice. You buy products built on these, and “built on” tells you whose engine is in the car, not who is responsible for the car.
Predictive AI	It scores and ranks. It has been in health IT for years, often unlabeled.	No-show likelihood, coding suggestions, sepsis alerts. It sorts. It does not write.
Generative AI	It creates new content — text, images, audio, video. This is what most people mean by “AI” today.	You will meet it as writing: draft notes, referral letters, patient instructions, portal replies.
Ambient documentation (the “AI scribe”)	Generative AI plus ambient audio capture. Many can be wired into your documentation workflow.	The realistic first purchase, and the one where the microphone is in the room with the patient.
Agentic AI	It does not stop at a draft. Give it a goal and it plans, takes actions in other software, and keeps going until it decides the goal is met.[2]	dX80's position: use an agent only if you can say exactly what it will do before it does it. Its strength and its danger are the same thing, which is that it follows the instruction completely and without pausing to check.

THE TERM	WHAT IT ACTUALLY IS	WHAT IT MEANS IN YOUR PRACTICE
Hallucination (NIST calls it confabulation)[3]	Confidently presenting false content. The confidence is the dangerous part.	In ambient scribing, omission may be the more common failure — 18% of notes in one 2026 pilot, against 11.5% with hallucinations.[4] It cannot document what never reached its input.
Human in the loop	A qualified person reviews the output before it counts, at a defined point.	You hold the license and the AI does not. Staying in the loop is how you confirm that what it suggested, wrote, or did is actually correct.
Training vs. inference	Training changes the model. Inference runs the trained model, and ordinarily leaves it unchanged.	This sits behind "we do not train on your data." A real promise — but it says nothing about storage, logging, or retention. Ask those separately.

Open, closed, open-weight: who actually holds your data

"Open" gets used several different ways and none of them means "safe."[5] The only question underneath all of it is who ends up holding your patient data, and who has to keep it safe. With a hosted product that is the vendor, under an agreement you can hold them to. An open-weight model has to be hosted and configured by someone, and that someone is you: the data stays in your building, and so does every decision that protects it.[6]

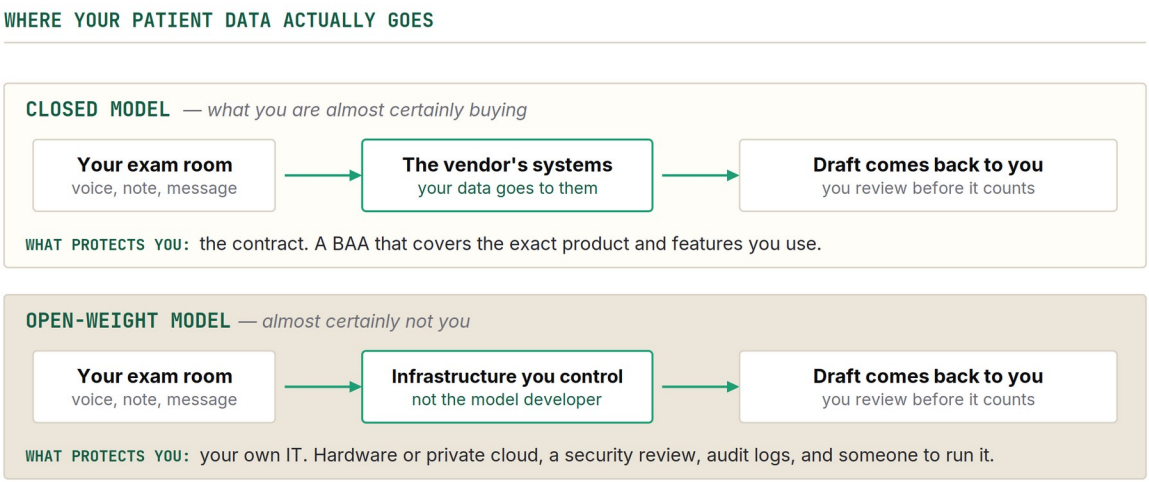

Open source is stricter than open weights, and rarer than either is claimed to be. Neither one means safe, and closed does not mean unsafe.

What you probably do not need to worry about

THE THING YOU KEEP HEARING	WHY IT IS NOT YOUR PROBLEM
Open-weight models and self-hosting	Far more technical, security, and maintenance responsibility than buying a hosted product. Not the practical first choice for an independent practice, and declining it is not a compliance failure.
Frontier model policy and compute thresholds	Governs what the model companies must disclose. Entirely upstream of you.
AI practicing medicine	No AI is licensed as a physician, and states are moving in opposite directions — Delaware barred it in April 2026,[7] Utah is testing AI prescription renewal under a phased regulatory pilot.[8,9] Assume the licensed clinician remains responsible unless a law says otherwise.[10]
AI replacing you	Some narrow AI-enabled or autonomous diagnostic services, such as autonomous diabetic-retinopathy screening, have their own payment pathways.[11] There is no general reimbursement model for an AI to replace the treating clinician.
Building anything yourself	You are a buyer, not a builder. Every term here is a question to ask, not a thing to implement.

REFERENCES

Check our work.

1. US Food and Drug Administration. FDA digital health and artificial intelligence glossary—educational resource. Updated September 26, 2024. Accessed August 29, 2026. https://www.fda.gov/science-research/artificial-intelligence-and-medical-products/fda-digital-health-and-artificial-intelligence-glossary-educational-resource
2. National Institute of Standards and Technology. Agentic AI. Updated August 14, 2026. Accessed August 29, 2026. https://www.nist.gov/agentic-ai
3. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. US Dept of Commerce; July 2024. Accessed August 29, 2026. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
4. Taylor SL, Jost M, MacDonald S, et al. Quality of clinical notes created by ambient listening generative AI: pragmatic prospective pilot study. JMIR Med Inform. 2026;14(4):e86474. doi:10.2196/86474
5. Open Source Initiative. The open source AI definition, version 1.0. October 28, 2024. Accessed August 29, 2026. https://opensource.org/ai/open-source-ai-definition
6. National Telecommunications and Information Administration. Dual-Use Foundation Models With Widely Available Model Weights. US Dept of Commerce; July 2024. Accessed August 29, 2026. https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf
7. Del HB 191, 153rd Gen Assem, Reg Sess, 85 Del Laws ch 250 (2026) (enacted April 23, 2026).
8. Utah Code Ann § 13-72-401 (2024) (regulatory mitigation agreements, Artificial Intelligence Policy Act).
9. Utah Office of Artificial Intelligence Policy. Authorized AI pilots: Doctronic regulatory mitigation agreement. Accessed August 29, 2026. https://commerce.utah.gov/ai/regulatory-relief/authorized-ai-pilots/doctronic/
10. American Medical Association. Augmented Intelligence Development, Deployment, and Use in Health Care. November 2024. Accessed August 29, 2026. https://www.ama-assn.org/system/files/ama-ai-principles.pdf
11. Medicare Program; CY 2022 Payment Policies Under the Physician Fee Schedule and Other Changes to Part B Payment Policies. 86 Fed Regist 64996 (2021) (establishment of values for remote retinal imaging, CPT code 92229).