AI TERMINOLOGY

GLOSSARY

Every AI word in the course — in plain language

Get this at the START — so a word you don't know never stops you

AI IN HEALTHCARE | Youssef Hagag

How to Use This Glossary

You're getting this at the very start of the course — on purpose. Keep it open beside the videos. Whenever you hear a word you're not sure about, pause and look it up here. Never let one unfamiliar term cost you the next five minutes of a chapter.

The terms are ordered from the most foundational (what AI even is) to the practical vocabulary you'll use daily (prompts, hallucination, agents, PHI). Definitions are deliberately plain — no computer science required.

The 36 Terms

Artificial Intelligence (AI)	Software that performs tasks we associate with human intelligence — understanding language, recognizing patterns, reasoning, making predictions.
Machine Learning (ML)	A branch of AI where systems learn patterns from examples/data instead of being programmed with explicit rules.
Deep Learning	ML using large multi-layer neural networks; the approach behind modern image and language AI.
Neural Network	A model loosely inspired by brain neurons, made of layers of simple units that together learn complex patterns.
Large Language Model (LLM)	An AI trained on huge amounts of text that predicts the next word/token — the engine behind ChatGPT, Claude, Gemini, and DeepSeek.
Generative AI	AI that CREATES new content — text, images, audio, code — rather than only classifying or predicting a label. LLMs are generative AI for text.
Token	A word-piece: the unit AI reads and writes in. 'Hypertension' may be several tokens.
Context Window	How much text an AI can hold and work with at once — its working memory. More context in usually means better answers out.
Prompt	The instruction/brief you give the AI. Its quality largely determines the output's quality.
Prompt Engineering	The skill of writing effective prompts — the one AI skill that transfers to every tool.
Zero-shot	Asking the AI to do a task directly, with no examples provided.
Few-shot	Giving the AI 1-3 examples of what you want before asking it to produce a new one.
Chain-of-Thought	Prompting the AI to reason step by step and show its logic before the final answer.
Hallucination	When AI states false information confidently as if it were true — its most dangerous failure mode in medicine.

Pretraining	The first training phase: the model reads enormous text to learn language and knowledge.
Fine-tuning	Further training on narrower data to specialize a model or make it more helpful.
RLHF	Reinforcement Learning from Human Feedback — humans rank answers to teach the model what 'good' looks like.
Knowledge Cutoff	The date after which the model has no built-in knowledge; anything newer must be verified with live sources.
Multimodal	An AI that handles more than text — images, audio, and/or video as input or output.
Retrieval-Augmented Generation (RAG)	Connecting an LLM to a document/knowledge source so it answers from those specific materials (e.g., NotebookLM).
Grounding	Tying an AI's answers to specific, verifiable sources rather than its general memory.
System Prompt	The permanent instructions that define a custom assistant's role, rules, and behavior.
Custom GPT / Gem	A pre-configured assistant (in ChatGPT / Gemini) with a fixed role, rules, and optional knowledge files.
Artifact (Claude)	An interactive mini-tool or document Claude builds inside the chat — e.g., a calculator or decision aid you can share as a link.
AI Agent	AI that doesn't just answer but WORKS — it plans steps, uses tools, browses, and produces finished output toward a goal.
Automation	Chaining steps so a task runs with little or no manual effort; agents are one way to automate knowledge work.
Temperature	A setting controlling how 'creative' vs. deterministic an AI's output is; lower = more focused and repeatable.
Parameters	The internal values a model learns during training; more parameters can mean more capability (and cost).
Transformer	The 2017 neural-network architecture ('Attention Is All You Need') behind virtually all modern LLMs.
Inference	Running a trained model to get an answer (as opposed to training it).
API	A way for software to talk to an AI model programmatically, so tools can build AI into their own products.
Open-source Model	A model whose weights are publicly available, so it can be self-hosted — enabling on-premises, private deployments.

PHI (Protected Health Information)	Any data that can identify a patient — names, IDs, dates, locations, photos of reports. Never put PHI into consumer AI.
De-identification	Removing every identifier so information can no longer be traced to a person — do this BEFORE any AI prompt.
Human-in-the-Loop	The principle that a qualified professional reviews and owns every AI-assisted output before it affects care.
Guardrails	Rules and constraints (in a prompt or a tool) that keep AI output safe and within scope.

Remember the one rule that outranks every term here: never put real, identifiable patient data into a consumer AI tool. De-identify first — always.

AI IN HEALTHCARE — From Zero to Empowered | Youssef Hagag | Version 3.1 — July 2026