

Ten terms from Class 2. Learn them well enough to explain each one in your own words.

1. Token

The chunks of text an AI model reads and writes. A model doesn't see words; it sees tokens, each about three-quarters of a word. "The" is one token. "Unbelievable" might be three. Tokens are the unit everything about AI is counted in: how much text a model can take in at once, and how much you pay to use it.

2. Training

How a model is made. A model has one job: given some text, predict the next word. To get good at it, the model reads a huge amount of text, a large slice of the internet, and makes billions of predictions, adjusting itself a little each time it's wrong. Training happens once, costs millions of dollars, and is where everything the model "knows" comes from.

3. Inference

A model answering a prompt. You type something, and the model uses what it learned in training to predict its reply, one word at a time. The model doesn't learn anything new while it does this. It only uses what it already knows, the way a person answers a question without re-studying. Training happens once; inference happens on every prompt, and it's what you pay for.

4. RLHF

Short for Reinforcement Learning from Human Feedback. After training, people rate the model's answers, and the model is adjusted to give more of the answers people liked. This is why models are polite and eager to help. It's also why they sometimes agree with you, or flatter you, when they should push back.

5. Knowledge cutoff

The date the model's training text ends. The model knows nothing that happened after that date, and it may confidently make something up rather than say so. If you need it to know something recent, paste in the material or give it a link.

6. Tool

Something an agent is allowed to do besides write text: open a file, run a program, search the web, send an email. On its own, a model can only talk. Tools are what let it act.

7. Tool call

An agent asking to use a tool. The model writes a short request, like "search the web for X." The harness, the software around the model, does the search and hands the result back. The model never does the action itself. It only asks for it.

8. The agentic loop

How an agent finishes a job: one tool call at a time, checking each result before deciding what to do next. Ask it to fix a typo in a file, and it reads the file, sees the text, changes the word, checks that the change worked, and tells you it's done. Every turn is the same: decide, use a tool, look at the result, decide again. That's why an agent can work through a whole task, while a chatbot only replies once.

9. Computer use

An agent using a computer the way a person does: looking at the screen, moving the mouse, clicking, typing. It lets an agent use any app or website, even ones that were never built with AI in mind.

10. Skill

Instructions you write down once so an agent can use them whenever a matching job comes up. Think of a recipe card the agent pulls out when it needs it. Skills are how you teach an agent to do something your way, without explaining it every time.