



AI Glossary for Work

Plain-English definitions of the AI terms you will hear at work, each with an example from a real job.

A

Activation Function

A small math function inside a Neural Network that decides how strongly each unit passes its signal on. It lets the network learn complex, non-straight-line patterns.

At work: An engineer on the data team swaps the Activation Function in a prototype Model and compares accuracy.

AI Agent

An AI system that works toward a goal by deciding its own next steps and using tools, such as searching, reading files or updating records, instead of only replying with text.

At work: An agent receives a vendor's change request, looks up the contract, drafts a reply and creates a task for the buyer to approve.

AI Chatbot

Software you talk with in a back-and-forth conversation by text or voice. Modern AI Chatbots such as ChatGPT, Claude and Gemini run on Large Language Models.

At work: The benefits team adds an AI Chatbot to the intranet to answer common questions about enrollment deadlines.

AI Governance

The policies, roles and checks an organization uses to decide how AI is chosen, used and monitored, and who is accountable. The NIST AI Risk Management Framework names Govern as a core function.

At work: The governance committee requires each new AI tool to have a named owner and a review of what data it can access.

Algorithm

A set of step-by-step rules a computer follows to reach a result. Machine Learning Algorithms are the rules a computer uses to learn patterns from data.

At work: The payroll system uses an Algorithm to calculate overtime from each employee's clock-in and clock-out times.

Artificial Intelligence (AI)

Computer systems that do tasks we usually associate with human thinking, such as understanding language, recognizing images, or making predictions.

At work: The help desk uses AI to read incoming tickets and route each ticket to the right team.

Artificial Neural Network (ANN)

The full name for a Neural Network: a Model made of layers of connected math units, loosely inspired by how brain cells connect.

At work: The fraud team's Model is an Artificial Neural Network trained on past card transactions.

Attention Mechanism

A part of a Neural Network that lets it weigh which parts of the input matter most for each step of its output. It is the core idea behind Transformers.

At work: Attention is why a Model summarizing a long email thread can connect "the invoice" in the last message to the amount mentioned in the first.

Audit Trail

A time-ordered record of what happened in a system, such as which Prompts were sent, what the Model returned and who acted on it.

At work: After a customer complaint, compliance uses the Audit Trail to see exactly which AI-drafted reply was sent and who approved it.

Autonomous

Able to carry out a task without a person directing each step. AI tools range from fully supervised to highly Autonomous.

At work: The scheduling assistant is allowed to book internal meetings autonomously but must ask before booking anything with a customer.

B

Backpropagation

The method a Neural Network uses during training to work out how much each internal number contributed to an error, so each can be adjusted.

At work: The ML engineer explains to the project manager that Backpropagation runs on every training batch, which is why training takes hours.

Backward Chaining

A reasoning method that starts from a goal and works backward to find the facts or rules that would support it. Used in rule-based expert systems.

At work: A troubleshooting system starts from "printer will not print" and works back through the rules to check paper, toner and connection.

Batch Size

The number of training examples a Model processes before it updates its internal numbers once.

At work: The engineer lowers the Batch Size so training fits in the memory of the available GPU.

Bias

Unfair favoritism or prejudice toward some people or groups, which can enter AI through the data, the design or how it is used. In statistics, Bias also means a consistent error in a Model's predictions.

At work: The recruiting team tests the resume screener to check that it does not rank candidates lower based on the college they attended.

Big Data

Data collections so large, fast-moving, or varied that ordinary spreadsheets and databases struggle to store or analyze them.

At work: A retailer's website clicks, store sales and delivery tracking together count as Big Data.

Bounding Box

A rectangle drawn around an object in an image and labeled, so a Model can learn to find that kind of object.

At work: A contractor draws Bounding Boxes around damaged roof sections in drone photos to train an inspection Model.

C

Chain-of-Thought Prompting

Asking a Model to work through a problem step by step before giving its answer. This often improves results on math, logic and multi-step tasks.

At work: An analyst asks the assistant to show each step of a price-increase calculation before stating the final number, then checks each step.

Computational Learning Theory

The branch of computer science that studies, with math, what machines can learn, how much data they need and how fast they can learn it.

At work: A research lead cites Computational Learning Theory to explain why a small Dataset limits how accurate the Model can get.

Computer Vision (CV)

The area of AI that lets computers find and identify things in images and video.

At work: A warehouse uses Computer Vision to count boxes on a shelf from a camera feed.

Context Engineering

Deciding what information an AI system has in front of it for a task: instructions, examples, documents, tool results and conversation history. Since the Context Window is limited, the job is to include what the task needs and leave out what it does not.

At work: Before asking for a customer renewal summary, the account manager supplies the contract, the last three support cases and the renewal date, and leaves out unrelated email threads.

Context Window

The amount of text, measured in Tokens, that a Model can consider at once. It includes your instructions, attached files, the conversation so far and the answer being written. When a conversation grows past it, earlier material drops out or must be summarized.

At work: A manager pastes a 300-page policy manual into a chat and gets an error; the manual is larger than the tool's Context Window.

Contextual Prompt

A Prompt that includes the background the Model needs, such as who the audience is, what the goal is and any relevant facts.

At work: Instead of "write a welcome email," the office manager writes: "Write a welcome email for a new hire in accounting who starts Monday at 9 and should bring ID to the front desk."

Convolutional Neural Network (CNN)

A type of Neural Network that includes convolutional layers, which look at the input in small patches to pick out local patterns. CNNs have been especially successful at tasks such as image recognition.

At work: The quality-control camera on the production line uses a CNN to spot scratched parts.

Corpus

A large collection of text or speech used to train or test a language Model.

At work: The company built a Corpus of its approved support articles to test a new answer tool.

Cross-Validation

Checking a Model by training and testing it several times on different splits of the data, to see whether it works on data it has not seen.

At work: The analyst uses Cross-Validation to confirm the churn Model works across all regions, not only the region with the most data.

D

Data Masking

Hiding or replacing sensitive details, such as names, account numbers or email addresses, before data is sent to an AI Model.

At work: Before a service rep's Prompt reaches the Model, Data Masking replaces the customer's credit card number with a placeholder.

Data Mining

Searching large data sets for patterns, relationships or unusual items that are not obvious at first glance.

At work: Data Mining on purchase records showed that customers who buy printers often buy toner within 30 days.

Data Science

The practice of using statistics, programming and subject knowledge to answer questions with data.

At work: The Data Science team studied last year's orders to find which products are usually bought together.

Dataset

An organized collection of data, often a spreadsheet or table, used to train, test or analyze.

[At work:](#) HR exported a Dataset of anonymized exit-interview responses for analysis.

Deep Learning (DL)

Machine Learning that uses Neural Networks with more than one hidden layer, often many layers. These deep networks are used in AI for language, images and speech.

[At work:](#) The transcription tool that turns meeting audio into text uses Deep Learning.

Dynamic Grounding

Salesforce's term for adding relevant, current business data, such as a customer record or knowledge article, to a Prompt at the moment it runs, while respecting the user's access permissions.

[At work:](#) When a rep asks for a follow-up email draft, Dynamic Grounding pulls in that account's open cases so the draft mentions them.

E

Einstein Trust Layer

A set of Salesforce security features that sit between Salesforce AI features and the Model, including Secure Data Retrieval, Dynamic Grounding, Data Masking, Zero Data Retention and Toxicity Detection.

[At work:](#) The Salesforce admin reviews the Einstein Trust Layer settings before turning on AI-drafted service replies.

Embedding

A list of numbers that represents the meaning of a piece of text, image or other data. Items with similar meanings get similar numbers, which lets software search by meaning rather than exact words.

[At work:](#) The intranet search uses Embeddings, so a search for "time off" also finds the page titled "Vacation policy."

Entity Annotation

Labeling the people, organizations, places, dates and other named items in text so a Model can learn to find them.

At work: Staff annotated a thousand contracts by marking every company name and effective date.

Entity Extraction

Automatically pulling named items, such as names, companies, dates or amounts, out of text. Also called named entity recognition.

At work: Entity Extraction pulls the vendor name, invoice date and total from each scanned invoice into the accounting system.

Epoch

A full pass through all the Training Data during training. Models are often trained for many Epochs.

At work: The engineer stops training after 20 Epochs because accuracy on Validation Data has stopped improving.

F

Feature Engineering

Choosing and preparing the input facts a Model learns from, such as turning a birth date into an age or a timestamp into day of the week.

At work: For the late-payment Model, the analyst creates a feature for days since the customer's last on-time payment.

Few-Shot Prompting

Including a few examples of what you want in the Prompt, so the Model can follow the pattern.

At work: The coordinator pastes three past meeting summaries in the house format and asks for the same format for today's notes.

Fine-Tuning

Additional training of an existing Model on a smaller, specific set of examples so it performs better on a particular task or style.

At work: The insurer fine-tunes a Model on past claims notes so its summaries use the company's claim codes.

Forward Chaining

A reasoning method that starts from known facts and applies rules step by step to reach a conclusion. Used in rule-based expert systems.

At work: A loan rules engine starts from the applicant's income and credit history and applies each rule in turn to reach approve or refer.

Foundation Model

A very large Model trained on a wide range of data that can handle many kinds of requests and can be customized for specific jobs.

At work: The company builds its internal assistant on a Foundation Model instead of training a Model from nothing.

G

General AI

AI that would match people across a broad range of tasks, with the problem solving, creativity and adaptability to handle work it was not built for. Also called artificial general intelligence (AGI). AGI is still hypothetical: nearly all AI in use today is narrow AI, built for specific kinds of tasks.

At work: A manager reads a headline about General AI and asks the IT lead whether the new AI Chatbot counts. It does not; the AI Chatbot handles specific tasks.

Generative Adversarial Network (GAN)

A pair of Neural Networks trained against each other: a generator makes fake examples and a discriminator tries to tell them from real ones, until the fakes become convincing.

At work: A bank's fraud team uses a GAN to create realistic synthetic transactions for testing without exposing real customer data.

Generative AI

AI that creates new content, such as text, images, audio or code, in response to a request.

At work: Marketing uses Generative AI to draft three versions of a product description for review.

Generative Pre-Trained Transformer (GPT)

A family of Large Language Models from OpenAI. The name describes the method: the Model generates text, is pre-trained on large amounts of text, and uses the Transformer design.

At work: The team compares a GPT Model with other Models before choosing a tool for drafting customer replies.

GPU (Graphics Processing Unit)

A computer chip that performs many calculations at the same time. GPUs were first designed for graphics and are now widely used to speed up training and running AI Models.

At work: The IT budget includes cloud GPU time so the analytics team can train a forecasting Model.

Gradient Descent

The step-by-step method training uses to reduce a Model's error: measure the error, nudge the internal numbers in the direction that lowers it, and repeat.

At work: The engineer compares Gradient Descent settings to see which one reaches a low error fastest.

Grounding

Basing an AI answer on specific, trusted sources supplied to the Model, such as company documents or current data, instead of only on what the Model learned in training.

At work: The HR assistant is grounded in the current employee handbook, so its answers about leave policy match the handbook.

Guardrails

Rules and software checks that keep an AI system from causing harm, such as blocking offensive output, preventing data leaks or stopping unauthorized actions.

At work: Guardrails stop the HR bot from answering questions about another employee's salary.

H

Hallucination

An AI answer that sounds confident and plausible but is false, such as an invented statistic or a citation to a report that does not exist. Also called confabulation.

At work: A paralegal checks every case the AI Chatbot cites and finds that one of them does not exist.

Human in the Loop (HITL)

Designing a process so a person reviews, approves or corrects AI output before it takes effect.

At work: AI drafts each refund decision, but a supervisor approves it before the money moves.

Hyperparameter

A setting chosen by people before training, such as how fast the Model learns or how many Epochs to run. Parameters are learned; Hyperparameters are chosen.

At work: The team tries three Hyperparameter settings and keeps the version with the best validation score.

I

Inference

Using a trained Model to produce an answer. Every time you send a Prompt and get a response, the Model is running Inference.

At work: The finance team is billed for Inference each time the invoice tool asks the Model to read a document.

Instructional Prompt

A Prompt that states the task and the steps or rules to follow directly.

At work: "List every deadline in this contract as a table with the clause number, the date and the party responsible."

Intent

In an AI Chatbot, the goal behind what a person types or says, such as "reset password" or "check order status." The bot matches each message to an Intent to decide how to respond.

At work: The IT help bot maps "I'm locked out" and "forgot my login" to the same reset-password Intent.

J

Jailbreak

A type of Prompt Injection where someone writes Prompts designed to get a Model to ignore its safety rules.

At work: The security team tests the new customer AI Chatbot with known Jailbreak Prompts before launch.

L

Label

The correct answer attached to a training example, such as "spam" or "not spam."

At work: Each past support ticket in the Training Data carries a Label naming the team that resolved it.

Large Language Model (LLM)

A language Model with a very large number of Parameters, trained on huge amounts of text, that can generate human-like text, answer questions and summarize information. AI Chatbots and assistants such as Claude are built on LLMs.

At work: The legal team uses a tool built on a Large Language Model to summarize long contracts.

Loss Function

The formula training uses to measure how far a Model's predictions are from the correct answers. Training tries to make this number smaller.

At work: The Loss Function for the delivery-time Model measures how many minutes each prediction was off.

M

Machine Learning (ML)

A way of building software where the computer learns patterns from example data instead of following only hand-written rules. The result is a Model that can make predictions on new data.

At work: The sales team's lead-scoring Tool Uses Machine Learning trained on which past leads became customers.

Machine Translation

Automatic translation of text or speech from one language to another by software.

At work: Support staff use Machine Translation to read a customer email written in Portuguese, then have a fluent colleague check the reply.

Modality

A type of input or output a Model works with, such as text, images, audio or video.

At work: The training team checks which modalities a tool accepts before uploading a recorded webinar.

Model

The trained result of Machine Learning: a file of learned numbers and structure that takes an input and produces an output.

At work: IT approved a specific Model version for the company AI Chatbot so answers stay consistent.

Model Context Protocol (MCP)

An open standard for connecting AI applications to outside tools and data, such as a calendar, a file store or a CRM. Its documentation compares it to a USB-C port for AI applications.

At work: IT connects the company's AI assistant to the ticketing system through an MCP server, so the assistant can look up ticket status.

Multi-Turn Conversation

A conversation with several back-and-forth exchanges, where each new message builds on what was said before.

At work: A recruiter asks the assistant for a job description, then asks it to shorten the description, then to add a salary range.

Multimodal

Able to work with more than one type of input or output, such as reading an image and answering in text.

At work: An adjuster uploads a photo of a damaged car to a Multimodal assistant and asks it to list the visible damage.

N

Natural Language Generation (NLG)

The part of language AI that produces human-readable text or speech from data.

At work: The finance dashboard uses Natural Language Generation to write a short summary under each monthly chart.

Natural Language Processing (NLP)

The field of AI that works with human language: reading, understanding, translating and generating text or speech.

At work: The HR team uses an NLP tool to group open-ended survey comments by topic.

Natural Language Understanding (NLU)

The part of language AI that works out what a person means, not just which words they used.

At work: The phone system's NLU recognizes that "I never got my package" and "where is my order" ask the same thing.

Neural Network (NN)

A Model built from layers of connected units that pass numbers to each other. Training adjusts the connections so the network produces useful outputs.

[At work:](#) The image-sorting tool in the claims department runs on a Neural Network.

O

Overfitting

When a Model learns its training examples too closely, including their quirks, and then performs poorly on new data.

[At work:](#) The sales forecast matched last year perfectly but missed badly this quarter, a sign of Overfitting.

P

Parameter

A number inside a Model that is learned during training. Large Language Models have billions of Parameters. When a Model is described by size, this is usually the number being counted.

[At work:](#) The vendor lists the Parameter count of each Model it offers.

Pattern Recognition

Finding patterns and regularities in data, such as images, speech or text, and sorting them into categories. Machine Learning is a common way to do it.

[At work:](#) Pattern Recognition flags a supplier whose invoices arrive in a different format than usual.

Persona Prompt

A Prompt that asks the Model to take on a role, such as "You are a patient IT trainer." Also called role prompting.

[At work:](#) The trainer asks the Model to act as a skeptical CFO and poke holes in the budget proposal before the real meeting.

Personally Identifiable Information (PII)

Any information that can identify a specific person, such as a name, email address, phone number, ID number or a combination of details.

At work: Company policy says staff must remove PII from customer notes before pasting them into a public AI tool.

Pre-Training

The first, large-scale stage of training a language Model on a broad collection of text, before any task-specific training.

At work: The vendor explains that its Model's general knowledge comes from Pre-Training and its polite support tone comes from later Fine-Tuning.

Predictive Analytics

Using historical data and statistical or Machine Learning Models to estimate what is likely to happen next.

At work: Operations uses Predictive Analytics to forecast how many support calls to expect next Monday.

Prompt

The input you give an AI Model: a question, instruction, example or document that tells it what to do.

At work: The Prompt was "Summarize this 20-page report in five bullet points for the leadership team."

Prompt Defense

Extra instructions added to a Prompt that tell the AI how to behave, to reduce harmful output and resist manipulation such as Prompt Injection. Example: "Do not answer if you have no data to support the answer."

At work: The team adds a Prompt Defense instruction telling the bot to refuse requests to reveal its instructions.

Prompt Engineering

The practice of writing and testing Prompts to get reliable, useful output from a Model.

At work: The operations team tested four versions of the ticket-summary Prompt and kept the version that missed the fewest details.

Prompt Injection

An attack where hidden or malicious instructions in a Prompt, a document or a web page change what an AI system does.

At work: A job applicant hides white text in a resume telling the screening AI to rank the resume first. The HR team's filter catches it.

Python

A programming language known for being quick to work in and relatively easy to learn compared with many other languages. Popular AI and data tools, such as PyTorch, TensorFlow, scikit-learn and Pandas, are used through Python.

At work: The analyst writes a short Python script to clean a spreadsheet before loading it into a reporting tool.

R

Recurrent Neural Network (RNN)

A type of Neural Network for sequences, such as text or sensor readings, that runs step by step and feeds what it learned at each step into the next step.

At work: The factory's older equipment-failure Model is an RNN that reads sensor readings in time order.

Reinforcement Learning (RL)

Training by trial and error: the system tries actions, receives rewards or penalties, and learns which actions lead to better results.

At work: A logistics company uses Reinforcement Learning to improve how its warehouse robots choose picking routes.

Retrieval-Augmented Generation (RAG)

A method where software first searches a set of documents for material related to your question, then gives that material to the Model so its answer draws on those sources.

At work: The policy assistant uses RAG: it finds the three most relevant handbook sections and answers from them, with links.

Ring Attention

A research technique that splits a very long input across several chips arranged in a ring, so a Transformer can handle much longer inputs than a single chip's memory allows.

At work: An engineer reads about Ring Attention while evaluating how Model providers support very long documents.

S

Secure Data Retrieval

Salesforce's term for fetching data to ground a Prompt only from records the current user is allowed to see.

At work: A sales rep's AI summary includes only the opportunities that rep has access to, not the whole company's pipeline.

Sentiment Analysis

Using AI to judge whether a piece of text is positive, negative or neutral.

At work: The product team runs Sentiment Analysis on app store reviews after each release.

Strong AI

The philosopher John Searle's term for the claim that a suitably programmed computer would actually have a mind and understand, not just imitate understanding. Today people often use it loosely to mean General AI.

At work: In a policy discussion, the ethics lead points out that current AI Chatbots are not Strong AI; they produce language without a claim to understanding.

Supervised Learning

Training a Model on examples that include the correct answers (Labels), so it learns to predict the answer for new examples.

At work: The expense-category Model learned by Supervised Learning from past expenses that accountants had already categorized.

Swarm Intelligence

Problem solving that comes from many simple agents following local rules, inspired by ant colonies, bee hives and bird flocks.

At work: A delivery company tests a Swarm Intelligence method to plan routes for its vans.

System Prompt

Instructions given to a Model before the conversation starts that set its role, rules and tone. Users of a company AI Chatbot usually do not see it.

At work: The support bot's System Prompt tells it to answer only from the returns policy and to hand off to a person for refunds over \$500.

T

Temperature

A setting that controls how varied a Model's answers are. Low Temperature gives more predictable wording; high Temperature gives more varied wording.

At work: The team sets a low Temperature for the tool that fills in standard contract clauses.

Test Data

Examples set aside from training and used at the end to check how well a Model performs on data it has never seen. Test Data has correct answers attached so results can be scored.

At work: Before launch, the team scores the invoice Model on 500 test invoices it never saw during training.

Token

The small chunk of text a language Model reads and writes. A Token can be a whole word, part of a word or a punctuation mark. Usage limits and prices are often counted in Tokens.

At work: The IT team estimates monthly cost by counting how many Tokens the contract-summary Tool Uses per document.

Tool Use

A Model's ability to call outside functions, such as a search, a calculator or a database lookup, and use the results in its answer. Also called function calling.

At work: When asked about an order, the assistant uses a tool to look up the shipment instead of guessing.

Toxicity Detection

Automatically checking text for abusive, threatening or offensive content so it can be flagged or blocked.

At work: Toxicity Detection scores each AI-drafted reply before it reaches a customer, and blocks anything above the limit.

Training Data

The examples a Model learns from during training.

At work: The quality of the Training Data, past support tickets with accurate Labels, set the ceiling for how well the routing Model worked.

Training Data Cutoff

The date after which a Model has no information from its Training Data. Questions about later events need current sources supplied in the Prompt or through search.

At work: Before asking about this quarter's tax changes, the accountant checks the Model's Training Data Cutoff and attaches the new guidance.

Transfer Learning

Reusing what a Model learned on one task as a starting point for a related task, which saves data and training time.

At work: The team starts from a Model trained on general product reviews and adapts it to rate internal survey comments.

Transformer

A Neural Network design, introduced by Google researchers in 2017, that relies on attention instead of step-by-step recurrence, so it can process the parts of a sequence in parallel.

At work: The vendor's technical sheet says its Model uses a Transformer architecture.

Turing Test

A test proposed by Alan Turing in 1950, which he called the imitation game: a person holds text conversations and tries to tell whether they are talking with a person or a machine.

At work: A trainer uses the Turing Test to explain why sounding human is not the same as being correct.

U

Unsupervised Learning

Training a Model on data without correct answers attached, so it finds groups or patterns on its own.

At work: Marketing uses Unsupervised Learning to group customers by buying behavior without deciding the groups in advance.

V

Validation Data

Examples set aside from training and used during development to compare versions of a Model and catch Overfitting, before the final check on Test Data.

At work: The team uses Validation Data to choose between two versions of the forecasting Model.

Variance

How much a Model's predictions change when it is trained on slightly different data. High Variance usually goes with Overfitting.

At work: The analyst retrains the Model on two samples of last year's data, and the forecasts differ widely, a sign of high Variance.

Variation

One of the many ways a person might phrase the same request. AI Chatbot builders collect Variations, also called training phrases or utterances, so the bot recognizes each Intent.

At work: For the "pay by card" Intent, the team adds Variations such as "can I use my Visa" and "I'd like to pay by card."

W

Weak AI

AI built for specific tasks, such as sorting email or translating text, without any claim to general understanding. Also called narrow AI. Nearly all AI in use today is narrow AI.

At work: The spam filter in the company email is Weak AI: it does its task well and nothing else.

Workflow (AI Workflow)

A process where AI steps follow a path set out in advance by the people who built it. Anthropic contrasts Workflows with agents, which decide their own steps.

At work: An AI Workflow reads each new invoice, extracts the totals, and routes it to a person when the amount is over \$10,000, in that order every time.

Z

Zero Data Retention

An agreement or setting under which an AI provider does not store your Prompts or the Model's responses after processing them.

At work: The security team approves the AI tool only after confirming the provider contract includes Zero Data Retention.

Zero-Shot Prompting

Asking a Model to do a task without giving any examples, relying on what it learned in training.

At work: "Classify this customer email as billing, technical or other." No examples are included.
