Machine Learning Model Glossary

Key model architectures and techniques, defined simply with examples

Contents

Agentic AI

AI systems that go beyond answering a single question: they can plan a sequence of steps, use tools, and take actions on their own to reach a goal, checking in with a human only when needed.

Example: *An agentic AI assistant that checks your calendar, books a flight, and emails you a confirmation — completing the whole chain of tasks itself — is acting as an agent.*

Artificial Intelligence (AI)

The broad field of building computer systems that can perform tasks normally requiring human intelligence, such as understanding language, recognising images, or making decisions. Machine learning, deep learning, and generative AI are all sub-fields within AI.

Example: *A voice assistant that understands a spoken question and gives a sensible answer is an application of AI.*

Autoencoders

A type of neural network that learns to compress data down to its essential features and then rebuild it back to the original. Because it learns purely from the structure of the data itself, with no labelled answers, it is a form of unsupervised neural network.

Example: *An autoencoder trained on photos of faces can learn to remove noise or grain from a blurry image by reconstructing a cleaner version.*

Classic Models

Traditional predictive (discriminative) models trained to sort or classify existing data into categories or outcomes, rather than to create new content. They answer questions like "what category does this belong to?" Decision trees and random forests are common examples.

Example: *A spam filter that labels an incoming email as "spam" or "not spam" is a classic model.*

Clustering Models

An unsupervised technique that groups similar data points together without being told in advance what the groups should be. The model finds natural patterns or clusters on its own.

Example: *A retailer uses clustering to group customers into segments like "bargain shoppers" and "brand loyalists" based on purchase history, without predefining those categories.*

Collaborative Filtering

A recommendation technique that predicts what a person might like based on the preferences and behaviour of similar users, rather than analysing the content itself.

Example: *Netflix recommending a show because "users who watched what you watched also enjoyed this" is collaborative filtering in action.*

Convolutional Neural Networks (CNNs)

A neural network architecture designed to recognise patterns in grid-like data, most commonly images, by scanning small sections at a time to detect shapes, edges, and textures.

Example: *A CNN powers the feature that automatically tags faces in your photo library.*

Decision Trees

A model that makes predictions by asking a series of yes/no or if-then questions, branching like a tree until it reaches a final decision or answer.

Example: *A decision tree used by a bank might ask: "Is income above £30,000? Then check credit history — any missed payments?" to decide whether to approve a loan.*

Deep Learning

A subset of machine learning that uses neural networks built from many stacked layers (hence "deep") to automatically learn complex patterns directly from raw data, without needing those patterns hand-designed by a person.

Example: *Deep learning lets a photo app recognise a dog in a picture without anyone manually describing what a dog looks like.*

Expert System

A rule-based AI approach that encodes human expert knowledge as a set of explicit if-then rules, using that fixed rulebook to reason through problems and reach conclusions, rather than learning patterns from data. Three main elements of expert systems are: **knowledge base**, **inference engine** and **user interface.**

Example: *A medical diagnosis expert system that asks "Does the patient have a fever?" and "Is there a rash present?", then follows a doctor-written rule like "if fever AND rash, then suggest measles" to reach its recommendation.*

Generative AI

A branch of AI focused specifically on producing new, original content — text, images, music, or code — modelled on patterns learned from existing examples, rather than sorting or predicting a fixed label.

Example: *An app that writes a birthday poem or creates an illustration from a text description is using generative AI.*

Generative Architectures

A broad category of models designed to create new content — text, images, audio, or video — rather than simply classifying or predicting a fixed label. Transformers, LLMs, and autoencoders can all be used generatively.

Example: *A tool that writes an original poem or generates a picture of "a cat wearing a spacesuit" is powered by a generative architecture.*

Generative Models

Models that learn the underlying patterns in a dataset well enough to create new, original examples, in contrast to classic models, which only sort or label existing data. GPT-style transformers and diffusion models are common examples.

Example: *A generative model can write a brand-new short story in the style of a favourite author, based on patterns it learned from that author's existing work.*

Language Models

Models trained specifically to understand and generate human language, spanning a huge range of sizes — from small, lightweight versions to today's largest LLMs.

Example: *Both a small on-device autocomplete tool and a large model like Claude are language models, just built at very different scales.*

Large Language Models (LLMs)

Very large transformer-based models trained on vast amounts of text, capable of understanding and generating human-like language across a wide range of topics and tasks.

Example: *ChatGPT and Claude are LLMs that can answer questions, write essays, and hold conversations.*

Large Models

A general term for large-scale machine learning models — often, though not only, language models — with vast numbers of parameters. Their size gives them broader reasoning ability, at the cost of speed and compute expense.

Example: *GPT-4-class models, with hundreds of billions of parameters, are large models capable of handling open-ended, complex tasks.*

Machine Learning

A subset of AI in which systems learn patterns from data and improve their performance over time, rather than being explicitly programmed with rules for every situation.

Example: *A machine learning model that gets better at recommending products the more purchase data it sees is learning from experience rather than fixed rules.*

Multimodal Models

Models that can understand and work with more than one type of input or output at once, such as text combined with images, audio, or video.

Example: *A multimodal model can look at a photo of a receipt and answer written questions about the amounts on it.*

Neural Networks

A general family of models loosely inspired by the brain, made up of layers of connected “neurons” that adjust their connections as they learn from data. Many other models on this list, such as CNNs and transformers, are specific types of neural network.

Example: *A neural network can learn to recognise handwritten digits by training on thousands of labelled examples of 0–9.*

Open-Source Models

Models whose underlying weights are published publicly, so that anyone can download, run, or fine-tune them without going through the original developer's product or API.

Example: *Meta's Llama models can be downloaded and run by developers on their own servers, unlike a closed, proprietary model.*

Proprietary Models

Closed-weight models that are only accessible through a company's official product or API, with the underlying weights kept private.

Example: *Access to Claude or GPT-4 is only available through Anthropic's or OpenAI's official channels, not by downloading the model itself.*

Random Forests

An ensemble model that builds many decision trees on different subsets of the data and combines their answers, usually by majority vote, to make a more accurate and stable prediction than any single tree.

Example: *A random forest predicting whether an email is spam combines the votes of hundreds of individual decision trees to reach a final verdict.*

Reinforcement Learning

A learning approach where a model (called an agent) learns by trial and error, receiving rewards or penalties for its actions, and gradually improving its strategy to maximise reward over time. It is often implemented using neural networks.

Example: *A reinforcement learning agent learns to play chess by playing millions of games against itself, earning a reward for winning.*

Retrieval-Augmented Generation (RAG)

A technique that pairs a language model with a search step: before answering, the system retrieves relevant documents or data from an external source and uses them to ground its response in up-to-date or specific information.

Example: *A company chatbot that searches internal policy documents before answering an employee's question is using RAG, rather than relying only on what the model memorised during training.*

Semi-Supervised Learning

A learning approach that combines a small amount of labelled data with a much larger amount of unlabelled data, blending the strengths of supervised and unsupervised learning to improve accuracy without needing everything labelled by hand.

Example: *A medical imaging model trained on a handful of scans labelled by doctors, plus thousands of unlabelled scans, to improve its diagnostic accuracy.*

Small Language Models (SLMs)

Language models with far fewer parameters than large language models, making them faster, cheaper to run, and often small enough to work directly on a phone or laptop, at the cost of some general capability.

Example: *A small language model built into a keyboard app that suggests the next word as you type is an SLM running entirely on your device.*

Small Models

A general term for models with fewer parameters, designed for speed, low cost, and the ability to run on smaller devices rather than large data centres.

Example: *A small model running offline inside a translation app on a phone, with no need for an internet connection.*

Supervised Learning

A learning approach where a model is trained on data that already has the correct answers attached (called labels), so it can learn to predict those labels for new, unseen inputs.

Example: *Training a model to tell cats from dogs using thousands of photos that are already labelled "cat" or "dog" is supervised learning.*

Time-Series Models

Models specifically designed to analyse data points collected over time, in order to detect trends, seasonality, or patterns and forecast future values.

Example: *A time-series model forecasts next month's electricity demand based on years of past hourly usage data.*

Transformers

A neural network architecture that processes entire sequences of data at once and uses a mechanism called "attention" to weigh how much each part of the input matters to every other part. Transformers are the foundation of modern LLMs.

Example: *The architecture behind GPT and Claude is a transformer, which is why these models are good at understanding context across a long passage of text.*

Unsupervised Learning

A learning approach where a model looks for patterns or structure in data that has no labels or predefined answers, discovering groupings or relationships on its own.

Example: *Grouping customers into segments without being told in advance what those segments should be, as in clustering models, is unsupervised learning.*