

AI Foundations Glossary

Use this printable resource with AI Foundations for Everyone. Write on a printed copy or type into a separate working document.

These definitions control learner-facing usage in the course. Module lessons may add context, but must not contradict them.

Algorithm

A defined procedure for transforming inputs into outputs. An algorithm may be written as explicit rules or used as part of a learning system.

Action

A change a system makes or initiates in a physical or virtual environment, such as sending a message, updating a record, running code, making a purchase request, or controlling a device. Producing a recommendation for a person to consider is different from taking the action itself.

AI agent

An AI system that receives information about its environment and can take self-directed actions, sometimes using external tools, in pursuit of an externally specified goal. Agents vary in capability and autonomy; the label alone does not show what the system can access, change, or do without approval.

Agentic AI

An emerging term whose usage is not fully consistent. The OECD's 2026 analysis uses it mainly for systems of co-ordinated AI agents that can divide work, collaborate, and pursue complex goals over time with limited supervision. This course describes the actual agents, tools, actions, duration, and controls instead of relying on the label alone.

Agentic system

A broad course term for an AI system that pursues a goal across multiple steps and can choose or take actions with some degree of autonomy. It may contain one agent or several co-ordinated agents. Calling a system agentic does not establish that it is reliable, safe, fully autonomous, or suitable for a task.

Artificial intelligence (AI)

An umbrella term for computer systems designed to perform tasks associated with capabilities such as perception, prediction, language processing, planning, or decision support. The term covers different techniques and does not imply human-like intelligence.

Automation

Using technology to carry out part or all of a process. Automation can use fixed rules, AI, or both.

Autonomy

The degree to which a system can learn or act without direct human involvement after people have delegated part of a process to it. Autonomy is a spectrum: a system may only recommend, may require approval before acting, may act unless stopped, or may act without operational human approval. Greater autonomy does not remove human accountability.

Bias

A systematic tendency in data, design, measurement, or system behaviour that can produce skewed or unfair results. Not every statistical difference is unfair, and fairness requires context about affected people and decisions.

Confidential information

Information that a person or organisation is required or expected to protect from unauthorised access or disclosure.

Confirmation

An explicit approval required before a system takes a defined action. Confirmation is meaningful only when the person can understand the proposed action, inspect material information, refuse or change it, and prevent execution. A generic Continue button is not automatically effective oversight.

Context

Information available to a model or application while it processes the current request. Context can include instructions, conversation history, retrieved material, and supplied files, subject to the system's limits.

Context window

The amount of tokenised information a model can process together for a request. A larger window does not guarantee that every detail will be used correctly.

Data

Recorded facts, measurements, text, images, audio, labels, or other representations used as input, evidence, or examples in a system.

Dataset

An organised collection of data used for training, evaluation, analysis, or another defined purpose.

Deep learning

A family of machine-learning methods based on neural networks with multiple processing layers. It is used in many modern language, image, audio, and prediction systems.

Deterministic process

A process designed to produce the same output from the same input and state. Real software can still fail, but deterministic tools are generally a better fit for exact repeatable rules and calculations.

Evaluation

Testing a system or output against defined criteria, evidence, examples, or metrics. Evaluation quality depends on whether those checks represent the real use and consequences.

Fairness

A context-dependent judgement about whether a system's process and effects treat people or groups appropriately. Fairness cannot be established by asking a model to declare itself fair.

Fabricated output

Content produced without adequate support, including invented facts, quotations, links, or citations. Hallucination is common shorthand for this behaviour but does not explain its cause.

Generative AI

AI systems that produce new content such as text, images, audio, video, or code in response to input and context.

Human oversight

Review by a person who has the relevant authority, information, time, criteria, and ability to change, reject, escalate, or stop a result.

Inference

Using a trained model to produce a prediction, classification, score, or generated output for an input.

Language model

A model trained to process and generate sequences of language tokens. A language-model application may add retrieval, tools, memory features, safety controls, and an interface around the model.

Machine learning

A way of building systems whose behaviour is fitted from data or examples rather than specified entirely as fixed task rules.

Model

A learned or designed representation used to transform inputs into predictions or outputs. A model is not the same as the complete product or service around it.

Neural network

A model structure made of connected computational units whose parameters are adjusted during training. Neural networks are used in deep learning.

Output

The text, image, score, prediction, classification, action, or other result produced by a system.

Parameter

A numerical value adjusted during model training that influences how the model transforms inputs into outputs.

Personal data

Information relating to an identified or identifiable person. The exact legal definition varies by jurisdiction, so course examples use a cautious practical interpretation and avoid real personal information.

Probabilistic system

A system that represents or produces results using probabilities or learned statistical patterns. Its outputs can vary and require evaluation appropriate to the task.

Prompt

Input that instructs or provides context to a generative-AI system. A prompt can be part of a broader task specification.

Retrieval

Fetching material from a defined source, index, database, or search system. Retrieval can provide evidence or context, but the retrieved material and any generated interpretation still require review.

Rule-based system

A system whose key behaviour is specified through explicit conditions and actions rather than learned from examples.

Source

The document, dataset, record, person, system, or other origin from which a claim or piece of information comes.

Task specification

A structured description of the intended work, including outcome, audience, context, source material, constraints, output format, and review criteria.

Tool

An external capability a system can call or operate, such as search, retrieval, a calculator, code execution, a database, an application programming interface, or an email service. Access to a tool does not prove that the system will choose or use it correctly.

Tool use

Selecting, calling, and using an external capability as part of a task. Tool use can allow a system to retrieve information or change its environment, so permissions, input validation, output checking, confirmation, logging, and limits may be required.

Token

A unit into which a model processes text or other content. A token may be a whole word, part of a word, punctuation, or another encoded unit depending on the model.

Training

The process of adjusting a model using data and an objective so that it develops behaviour useful for a task. Training is distinct from using the trained model during inference.

Verification

Checking a claim or result against suitable independent evidence, authoritative sources, reproducible calculations, or qualified review.

Workflow

A defined sequence of tasks, inputs, tools, decisions, checks, outputs, owners, and stop or escalation points used to achieve an outcome.

SAFE USE NOTE

Use supplied, public, fictional, or safely prepared synthetic material. Do not add personal, confidential, client, health, financial, legal, authentication, or security-sensitive information.