Understanding AI Terms

(AIGuide 1)

A Global Reference Guide to Artificial Intelligence Terms

AI Sourced Facts (AISF) Pte. Ltd.

Edition: Version 1.0
Publication Date: January 2026

PURPOSE OF THIS GUIDE

Artificial intelligence is becoming part of everyday life — in how people work, learn, communicate, and make decisions. Yet the language used to describe AI is often technical, inconsistent, or difficult for non-specialists to understand.

This booklet exists to provide a clear, neutral, and accessible reference to the most important terms used in discussions about artificial intelligence today.

It is written for:

- professionals who want clarity without jargon
- students and educators seeking reliable explanations
- leaders and policymakers navigating AI decisions
- everyday users who want to understand the tools they encounter
-

Our aim is simple:
to make the language of AI understandable, trustworthy, and useful for everyone.

GLOBAL PERSPECTIVE

A note on neutrality and scope

Artificial intelligence is a global technology, affecting people across cultures, industries, and nations.

This guide is written from a **global and neutral perspective**.
It avoids promoting any single company, country, or commercial viewpoint.

Instead, it reflects:

- commonly used international terminology
- shared public concerns about safety and trust
- widely recognised developments in AI practice

Our intention is to provide a reference that remains useful whether you are reading it in a classroom, a workplace, a policy setting, or at home.

EDITORIAL STANDARD

Our commitment to clarity

Every explanation in this booklet follows one simple rule:

Never assume the reader already understands artificial intelligence.

All entries are written to:

- avoid unnecessary jargon
- explain ideas in everyday language
- connect technical concepts to real-world consequences
- use examples people can recognise from daily life

This guide is not written for engineers alone.
It is written for **people**.

HOW TO USE THIS GUIDE

This booklet is designed to be easy to use, whether you are reading from start to finish or looking up a specific term.

Each term is presented in three clear parts:

What it means
Explains the term in plain language, without assuming technical knowledge.

Why it matters
Shows why the term is important in real discussions about work, safety, policy, or society.

In everyday life
Gives a simple, real-world example of where people might encounter this concept.

All terms are listed in **alphabetical order** for quick reference.
Each term is also tagged with a category code to show where it fits in the wider AI landscape.

VERSION & GOVERNANCE

Edition and maintenance

This booklet is published as part of the public-education work of **AI Sourced Facts (AISF)** Pte. Ltd.

Current edition: Version 1.0
Release date: January 2026

This is a living reference.

As the language of artificial intelligence evolves, future editions will add new terms and update explanations while preserving earlier editions as historical snapshots.

AISF UNIVERSAL DISCLAIMER

This publication is provided for informational and educational purposes only.

It presents a structured and time-bound analysis based on information available at the time of writing. The content reflects observed patterns, interpretations, and perspectives and does not constitute definitive, authoritative, or universally accepted conclusions.

This publication does not provide:

- legal, regulatory, or compliance advice
- financial, investment, or business advice
- technical implementation guidance
- operational or strategic recommendations

No representation or warranty is made as to the accuracy, completeness, or suitability of the information for any specific purpose.

Artificial intelligence systems, technologies, and related terminology evolve rapidly. As such, the observations and interpretations contained in this publication may change over time and may not reflect subsequent developments.

AI Sourced Facts (AISF) does not endorse, promote, or recommend any specific technologies, systems, vendors, products, or approaches.

The use of any information, concepts, or interpretations from this publication is at the reader's own discretion and risk.

AISF publications are intended to support informed human judgment. Final responsibility for decisions, actions, and outcomes remains with the user.

CATEGORY FRAMEWORK

To help readers see how different ideas relate to one another, each term is tagged with one or more category letters:

A — Core AI Identity
Fundamental concepts that define what AI is and how it is commonly understood.

B — Models & Learning
How AI systems are built, trained, and improved.

C — Language & Interaction
How people communicate with AI and how AI understands human language.

D — Generative Capabilities
AI's ability to create new content such as text, images, audio, and video.

E — Model Behaviour & Limits
How AI behaves in practice and where its limitations appear.

F — Ethics, Safety & Trust
Fairness, responsibility, and protection from harm.

G — Governance & Society
Rules, policies, and social impact of AI.

H — System Types & Architectures
Different kinds of AI systems and how they are structured.

I — Deployment & Use
How AI is implemented in real workplaces and services.

J — Future-Facing Language
Terms that look ahead to what AI may become.

K — Emerging & Public Discourse
Language used in media, culture, and public debate about AI.

These categories are not technical barriers — they are **navigation aids** to help readers see the bigger picture.

CATEGORY EXPLANATIONS

A — Core AI Identity

This category covers the fundamental ideas that define what artificial intelligence is and how it is commonly understood. It includes concepts that shape the basic identity of AI, such as different types of intelligence, the nature of AI systems, and the core terminology people use when they first encounter the field. These terms help readers answer the question: *"What is AI, really?"* and form the foundation for understanding everything else in the glossary.

B — Models & Learning

This category focuses on how AI systems are built and how they learn. It includes terms related to training, data use, learning methods, and model development. These are the concepts that explain *how AI improves, adapts, and becomes capable over time*. Readers who want to understand how machines "learn" and why data quality matters will find that this category provides the technical backbone of AI in accessible language.

C — Language & Interaction

This category covers how people communicate with AI and how AI understands human language. It includes terms related to prompts, conversations, instructions, and the way AI processes and responds to text or speech. These concepts explain *how humans and AI interact in practice*, making this category especially relevant for everyday users of chatbots, virtual assistants, and workplace AI tools.

D — Generative Capabilities

This category focuses on AI's ability to create new content rather than simply analyse existing information. It includes terms related to generating text, images, audio, video, and other media. These concepts help readers understand *what AI can produce*, how creative tools work, and why generative AI has become one of the most visible and talked-about applications of artificial intelligence.

E — Model Behaviour & Limits

This category explains how AI systems behave in real-world conditions and where their limitations lie. It includes concepts such as errors, uncertainty, reliability, and unexpected behaviour. These terms help readers understand *what AI can and cannot be trusted to do*, why mistakes happen, and why human judgment remains essential even when AI systems appear confident or intelligent.

F — Ethics, Safety & Trust

This category addresses the moral and safety dimensions of using AI. It includes terms related to fairness, bias, accountability, protection from harm, and building trust in AI systems. These concepts help readers understand *how AI should be used responsibly*, why safeguards are necessary, and how societies and organisations try to ensure that technology benefits people rather than harms them.

G — Governance & Society

This category focuses on how AI is managed at organisational, national, and global levels. It includes ideas about regulation, policy, oversight, social impact, and public responsibility. These terms explain *how societies make decisions about AI*, how rules are formed, and how AI affects work, rights, and institutions beyond just technical use.

H — System Types & Architectures

This category describes the different kinds of AI systems and how they are structured internally. It includes terms related to system design, model architecture, and the technical forms AI can take. These concepts help readers understand *what types of AI exist*, how they are organised, and why different architectures are used for different purposes.

I — Deployment & Use

This category focuses on how AI is applied in real settings. It includes terms about implementation, integration, performance, cost, and everyday operation. These concepts explain *how AI moves from theory to practice*, how it fits into workplaces and services, and what practical issues arise when AI is used at scale.

J — Future-Facing Language

This category covers ideas that look ahead to what AI may become and how it could shape the future. It includes terms related to long-term development, advanced capabilities, and evolving roles of AI in society. These concepts help readers explore *where AI might be heading*, without treating speculation as certainty.

K — Emerging & Public Discourse

This category captures the language that appears in public debates, media, and cultural conversations about AI. It includes terms that reflect hopes, fears, controversies, and social narratives around artificial intelligence. These concepts help readers understand *how AI is talked about in society*, not just how it works technically.

LIST OF AI TERMS

A

AGI (Artificial General Intelligence) — A, J, K
AI (Artificial Intelligence) — A
AI Act (EU AI Act) — G
AI Assistant — I, C
AI Capability — A, J
AI Literacy — G, I
AI Readiness — I, G
AI Safety — F, J
AI Security — F, I
AI Transparency — F, G
AIGC (AI-Generated Content) — D, K
Alignment in AI — F, J
Algorithm — A, G
Algorithmic Accountability in AI — G, F
Anomaly Detection — B, I
ASI (Artificial Superintelligence) — A, J, K
Assistive AI — I
Attention Mechanism — H, B
Audit in AI — G, F
Automated Decision-Making in AI — G, I
Automation with AI — I, G
Autonomous Agent — H, I, F
Autonomous Vehicle AI — I, H

B

Backpropagation — B, H
Base Model — A, B
Batch Size — B
Benchmarking in AI — B, I
Bias in AI — F, G
Bias Mitigation in AI — F, G
Big Data for AI — B, I
Black-Box Model — E, F
Boosting (Ensemble Learning) — B
Business Process Automation with AI — I

C

Chatbot — C, I
Chain-of-Thought Reasoning — C, E
Classification — B, I
Context Window — C, E
Conversational AI — C, I
Copilot Systems — I, C
Cross-Lingual Models — C, B
Code Generation — D, I
Content Moderation with AI — F, I
Creative AI — D

D

Data Augmentation — B
Data Labeling for AI — B, G
Data Mining with AI — B, I
Dataset — B
Dataset Shift — E
Deep Learning — A, B
Deepfake — D, F, K
Deployment Pipeline — I
Diffusion Model — H, D
Distillation — B, I
Document Summarisation — D, C
Domain Adaptation — B
Dual-Use Risk in AI — F, G
Dynamic Prompting — C

E

Edge AI — I
Edge Case — E
Emergent Behaviour — E, J, K
Emergent Capabilities — J, K
Embedding — C, B, I
Energy Use of AI (Compute Footprint) — G, I
Ensemble Models — B
Epistemic Uncertainty — E
Error Rate — E
Ethical Design in AI — F, G
Ethics in AI — F, G
Evaluation of AI (Evals) — E, F, I
Explainability in AI — F, G
Exposure Bias — E

F

Fairness in AI — F, G
Factuality in AI Outputs — E, F
Federated Learning — B, F, I
Feedback Loops in AI Systems — F, B
Few-Shot Learning — C, B
Fine-Tuning — B, I
Foundation Model — A, B
Frontier Models — J, K
Fraud Detection with AI — I
Function Calling / Tool Use — I, C, F

G

Generalisation — E, B
Generative AI — A, D
Governance in AI — G
Green AI — G, I
Ground Truth — B
Grounding (in Sources) — E, I
Guardrails in AI — F, I

H

Hallucination — E, F
Hard vs Soft Prompts — C, B
Heuristics in AI — H
Hidden Prompt / System Prompt — C, F
Homomorphic Encryption for AI — F, I
Human-in-the-Loop — F, I
Human Oversight in AI — F, G
Hybrid AI (Rules + Learning) — H, I
Hyperparameters — B

I

Image Generation — D
Image Recognition — I, H
Incident Response in AI Safety — F, G
In-Context Learning — C, B
Inference — B, I
Inference Cost — I
Inference Latency — I
Information Retrieval — I
Instruction Tuning — B, C
Integration of AI in Workflows — I

Intellectual Property in AI Outputs — G, F
Interpretability in AI — F, E

J

Jailbreak — F, K
Job Displacement from AI — G, K
Job Redesign with AI — G, I
Justification of AI Decisions — F, E

K

Knowledge Cutoff — E, C
Keyword vs Semantic Search — I, C

L

Large Language Model (LLM) — A, C
Latent Space — H, D
Legal Liability in AI Use — G, F
Local AI Models — I
LoRA (Low-Rank Adaptation) — B, I
Long-Context Models — C, E
Low vs High Temperature — C, E

M

Machine Learning — A, B
Misalignment in AI — F, J
Misinformation with AI — F, G, K
Model — A
Model Cards — G, F
Model Collapse — E, K
Model Monitoring — I, E
Model Scaling Laws — J, K
Multimodal AI — A, D, C
Multi-Agent Systems — H, I

N

Narrow AI (Weak AI) — A
Natural Language Processing (NLP) — C, A
Neural Network — H, B
Non-Determinism in AI — E
Novelty in AI Content — D, G

O

Open-Source AI Models — I, G
Open vs Closed AI Models — G, K
Optical Character Recognition (OCR) — I
Out-of-Distribution Data (OOD) — E
Output Filtering in AI — F, I
Overfitting — E, B
Oversight Boards for AI — G

P

Parameters (Model Parameters) — B
Personalisation with AI — I, F
Policy in AI — G
Post-Training — B
Pre-Training — B
Precision & Recall — E
Privacy in AI — F, G
Prompt — C
Prompt Engineering — C
Prompt Injection — F, C
Prompt Libraries — C, I
Proprietary AI Models — I, G
Public Interest AI — G

Q

Quantisation — I, B
Quality Assurance in AI — I, F

R

Rate Limits (APIs) — I
Real-Time AI — I
Recommendation Systems — H, I
Reinforcement Learning — B
Reliability in AI — E, F
Responsible AI — F, G
Responsible Deployment of AI — I, G
Retrieval-Augmented Generation (RAG) — I, C, E
Risk Assessment in AI — G, F
Robustness in AI — E, F
Rule-Based AI (Expert Systems) — H
Runtime Safety — F, I

S

Safety by Design in AI — F, G
Safety Cases for AI — G, F
Self-Supervised Learning — B
Semi-Supervised Learning — B
Sentiment Analysis — I, C
Shielding / Sandboxing of AI Tools — F, I
Small Language Models (SLMs) — A, I
Social Impact of AI — G
Speech Recognition — D, I
Speech Synthesis — D, I
Synthetic Data — B, F
Synthetic Media — D, F
System Prompt — C, F
Scenario Testing for AI — F, I

T

Temperature (Model Setting) — C, E
Text Generation — D
Token — C
Tool Use by AI Agents — I, F
Training — B
Training Data — B, F
Transfer Learning — B
Transformer Models — H, B
Transparency in AI — F, G
Turing Test — K, J

U

Underfitting — E
Uncertainty Estimation — E, F
Unsupervised Learning — B
Use Policies for AI — G, F
User Consent in AI Data Use — G, F

V

Vector Databases — I, C
Verification & Validation in AI — I, F
Video Generation — D
Virtual Agents — I, C

W

Weak AI (Narrow AI) — A
Workforce Impact of AI — G, K
Workflow Automation with AI — I
Watermarking of AI Content — F, G

Z

Zero-Shot Learning — C, B

EXPLANATION OF AI TERMS

A

AGI (Artificial General Intelligence) — *A, J, K*

What it means
Artificial General Intelligence refers to a future type of AI that could understand, learn, and perform a wide range of tasks at a level similar to a human being, rather than being limited to one specific job like translation or image recognition.

Why it matters
AGI is important because it represents a major shift in what AI could become. If such systems ever exist, they would raise profound questions about safety, responsibility, work, and how humans and machines should coexist.

In everyday life
You might hear AGI discussed in documentaries or opinion pieces that ask whether machines could one day "think like people" or replace many human roles.

AI (Artificial Intelligence) — *A*

What it means
Artificial Intelligence refers to computer systems that can perform tasks usually associated with human thinking, such as recognising patterns, understanding language, learning from experience, or making decisions.

Why it matters
AI is becoming part of everyday life, influencing how we work, communicate, learn, and receive services. Understanding what AI is helps people use it more wisely and with realistic expectations.

In everyday life
When a phone suggests the fastest route home or a streaming service recommends a movie, AI is helping make those decisions.

AI Act (EU AI Act) — *G*

What it means
The EU AI Act is a European law designed to regulate how artificial intelligence is developed and used, especially in areas where AI could affect people's rights, safety, or opportunities.

Why it matters
It is one of the first major legal frameworks focused specifically on AI, and it is influencing how companies around the world think about responsible AI design and use.

In everyday life
If you use an AI-powered service in Europe, that company must now follow stricter rules about transparency, safety, and accountability.

AI Assistant — *I, C*

What it means
An AI assistant is a digital tool designed to help people with everyday tasks such as writing emails, organising schedules, answering questions, or summarising information.

Why it matters
AI assistants are often the first direct experience people have with AI, shaping how they understand and trust the technology.

In everyday life
A tool that drafts meeting notes or helps plan a trip is acting as an AI assistant.

AI Capability — *A, J*

What it means
AI capability refers to what an AI system can actually do and how well it performs those tasks, such as understanding language, recognising images, or solving problems.

Why it matters
Knowing an AI's real capabilities helps avoid both over-trust ("it can do everything") and under-use ("it can't help at all"). Balanced understanding leads to safer and more effective use.

In everyday life
Realising that a chatbot can write text but cannot verify facts helps you use it responsibly.

AI Literacy — *G, I*

What it means
AI literacy means having a basic understanding of how AI works, what it can and cannot do, and how to use it wisely and safely.

Why it matters
As AI becomes more common, people who understand its limits are better able to avoid mistakes, misinformation, and misuse.

In everyday life
Double-checking important information from an AI tool before sharing it shows good AI literacy.

AI Readiness — *I, G*

What it means
AI readiness describes how prepared an organisation, school, or society is to adopt AI in a thoughtful and responsible way.

Why it matters
Being ready for AI means having the right skills, policies, and awareness to use the technology without causing confusion, harm, or unrealistic expectations.

In everyday life
A company that trains staff and sets clear rules before introducing AI tools is improving its AI readiness.

AI Safety — *F, J*

What it means
AI safety focuses on preventing AI systems from causing harm, whether through errors, misuse, or unintended consequences.

Why it matters
Without strong safety measures, AI systems could make serious mistakes in areas like healthcare, finance, or public information.

In everyday life
Restrictions that stop chatbots from giving medical diagnoses are an example of AI safety in action.

AI Security — *F, I*

What it means
AI security is about protecting AI systems from being hacked, manipulated, or used in harmful ways.

Why it matters
If AI systems are not secure, they can be exploited to spread misinformation, steal data, or cause real-world harm.

In everyday life
Preventing attackers from tricking a chatbot into revealing private customer information is part of AI security.

AI Transparency — *F, G*

What it means
AI transparency means being open about how an AI system works, what data it uses, and how it reaches decisions.

Why it matters
Transparency helps people trust AI systems and makes it easier to hold organisations accountable for how AI is used.

In everyday life
When a company explains why its AI rejected a loan application, it is practising transparency.

AIGC (AI-Generated Content) — *D, K*

What it means
AI-generated content refers to text, images, audio, or video created by an AI system instead of a human.

Why it matters
This type of content is changing how media, education, and business operate, while also raising questions about originality and authenticity.

In everyday life
An AI-created image used in an advertisement is an example of AI-generated content.

Alignment in AI — *F, J*

What it means
Alignment in AI means making sure that AI systems behave in ways that match human values, intentions, and safety expectations.

Why it matters
If AI systems are not aligned, they may act in ways that are technically correct but socially harmful or unsafe.

In everyday life
Rules that prevent a chatbot from giving dangerous advice are part of keeping AI aligned with human values.

Algorithm — *A, G*

What it means
An algorithm is a set of step-by-step instructions a computer follows to solve a problem or complete a task.

Why it matters
Algorithms form the foundation of many digital systems, including AI, and shape how decisions are made automatically.

In everyday life
Sorting your photos by date or relevance relies on algorithms.

Algorithmic Accountability in AI — *G, F*

What it means
Algorithmic accountability means ensuring that organisations remain responsible for the outcomes of AI-driven decisions.

Why it matters
Without accountability, people can be harmed by automated decisions without knowing who is responsible or how to challenge them.

In everyday life
Being able to appeal a decision made by an automated credit system shows accountability in practice.

Anomaly Detection — *B, I*

What it means
Anomaly detection uses AI to spot unusual patterns that may indicate errors, fraud, or emerging problems.

Why it matters
Detecting anomalies early can prevent financial loss, security breaches, or system failures.

In everyday life
Banks use anomaly detection to flag unusual spending on your credit card.

API (Application Programming Interface) — *I*

What it means
An API is a way for different software systems to communicate, allowing apps to use AI features without building them from scratch.

Why it matters
APIs make it easier for AI to appear in many everyday tools and services.

In everyday life
A language-learning app that uses an external AI translation service is using an API.

ASI (Artificial Superintelligence) — *A, J, K*

What it means
Artificial Superintelligence refers to a hypothetical future stage where AI would exceed human intelligence in most areas.

Why it matters
Although speculative, ASI is central to debates about long-term risks and responsibilities surrounding AI development.

In everyday life
You may hear ASI discussed in science-fiction films or future-focused opinion pieces.

Assistive AI — *I*

What it means
Assistive AI refers to systems designed to support people rather than replace them, helping with tasks while keeping humans in control.

Why it matters
This approach emphasises collaboration between humans and machines, rather than full automation.

In everyday life
A writing tool that suggests improvements but leaves final decisions to the user is assistive AI.

Attention Mechanism — *H, B*

What it means
An attention mechanism helps an AI system focus on the most relevant parts of information, much like how humans focus on key details in a conversation.

Why it matters
It greatly improves how AI understands language, images, and complex data.

In everyday life
When a chatbot responds accurately to the main point of a long message, attention mechanisms are at work.

Audit in AI — *G, F*

What it means
Audit in AI means carefully reviewing AI systems to check whether they are accurate, fair, safe, and compliant with rules.

Why it matters
Audits help organisations detect hidden risks and prevent harm before it affects people.

In everyday life
A company reviewing its hiring AI to ensure it does not discriminate is carrying out an AI audit.

Automated Decision-Making in AI — *G, I*

What it means
This refers to situations where AI systems make or strongly influence decisions with little human involvement.

Why it matters
Automation can improve efficiency, but it also raises concerns about fairness, accountability, and transparency.

In everyday life
Systems that automatically decide loan approvals use automated decision-making.

Automation with AI — *I, G*

What it means
Automation with AI means using intelligent systems to perform tasks that people previously did manually.

Why it matters
It can save time and money, but it also changes how jobs are designed and valued.

In everyday life
Automated customer service chat systems are a common example.

Autonomous Agent — *H, I, F*

What it means
An autonomous agent is an AI system that can plan actions and take steps toward goals with minimal human guidance.

Why it matters
Such systems can be powerful, but they also require strong safeguards to prevent unintended actions.

In everyday life
An AI that books meetings, sends reminders, and follows up automatically behaves like an autonomous agent.

Autonomous Vehicle AI — *I, H*

What it means
These are AI systems that allow vehicles to sense their surroundings and drive with little or no human input.

Why it matters
They promise safer transport but also raise major questions about responsibility when accidents occur.

In everyday life
Self-driving features in modern cars rely on autonomous vehicle AI.

B

Backpropagation — *B, H*

What it means
Backpropagation is a training method that helps AI systems learn by correcting mistakes step by step.

Why it matters
It is one of the key techniques that made modern deep learning possible.

In everyday life
When an AI becomes better at recognising faces after training, backpropagation played a role.

Base Model — *A, B*

What it means
A base model is a general-purpose AI system trained broadly before being adapted to specific tasks.

Why it matters
It allows developers to build many tools on top of one strong foundation.

In everyday life
A general chatbot later customised for healthcare support uses a base model.

Batch Size — *B*

What it means
Batch size refers to how much data an AI system processes at one time during training.

Why it matters
It affects how quickly and effectively an AI learns.

In everyday life
Larger batches can speed up training but require more computing power.

Benchmarking in AI — *B, I*

What it means
Benchmarking in AI means comparing different AI systems using standard tests.

Why it matters
These comparisons influence public perception and business decisions.

In everyday life
Articles claiming one chatbot is “better” than another often rely on benchmark results.

Bias in AI — *F, G*

What it means
When an AI system produces unfair or unequal results, often because it was trained on data that reflects human prejudices, historical inequalities, or gaps in representation.

Why it matters
Biased AI can cause real harm — such as unfair hiring decisions, unequal access to loans, or higher error rates in facial recognition for certain communities. These outcomes can reinforce existing inequalities instead of reducing them.

In everyday life
A news report about an automated hiring tool that consistently favors male candidates over equally qualified women is an example of bias in AI.

Bias Mitigation in AI — *F, G*

What it means
Bias mitigation in AI refers to the steps taken to reduce unfair outcomes in AI systems.

Why it matters
Reducing bias helps ensure that AI benefits everyone rather than disadvantaging certain groups.

In everyday life
A company reviewing its training data to ensure diverse representation is working on bias mitigation.

Big Data for AI — *B, I*

What it means
Big data for AI refers to the large amounts of information used to train and improve AI systems.

Why it matters
Having more data can improve AI, but poor-quality data can also amplify mistakes.

In everyday life
Streaming services collect viewing data to improve their recommendation systems.

Black-Box Model — *E, F*

What it means
A black-box model is an AI system whose internal decision process is difficult for humans to understand.

Why it matters
Lack of clarity can make it hard to trust or challenge AI-driven decisions.

In everyday life
If a system denies a benefit without explanation, it is acting like a black box.

Boosting (Ensemble Learning) — *B*

What it means
Boosting is a technique that combines several simpler models to create a stronger and more accurate AI system.

Why it matters
It improves reliability in tasks like prediction and classification.

In everyday life
Spam filters often rely on ensemble techniques like boosting.

Business Process Automation with AI — *I*

What it means
This refers to using AI to automate routine business tasks such as data entry, scheduling, or invoice processing.

Why it matters
It increases efficiency but also changes how people work and what skills are needed.

In everyday life
An automated system that processes expense claims is using AI-based automation.

C

Chatbot — *C, I*

What it means
A chatbot is an AI system designed to interact with people through text or voice conversations.

Why it matters
Chatbots are often the most visible way the public encounters AI.

In everyday life
Customer service chat windows on websites commonly use chatbots.

Chain-of-Thought Reasoning — *C, E*

What it means
This refers to an AI showing the steps it takes to reach an answer, rather than just giving the final result.

Why it matters
Seeing the reasoning can help people understand and check AI outputs, but it does not guarantee correctness.

In everyday life
When an AI explains how it solved a maths problem step by step, it is using chain-of-thought reasoning.

Classification — *B, I*

What it means
Classification is when an AI sorts information into categories based on patterns it has learned.

Why it matters
Many everyday AI tools rely on classification to function.

In everyday life
Email systems classifying messages as “spam” or “not spam” use this method.

Context Window — *C, E*

What it means
The context window is the amount of information an AI can keep in mind at one time.

Why it matters
It limits how well AI can follow long conversations or analyse lengthy documents.

In everyday life
If a chatbot forgets what you said earlier, it has likely reached its context limit.

Conversational AI — *C, I*

What it means
Conversational AI refers to systems designed to hold natural-sounding conversations with people.

Why it matters
These systems shape how comfortable and confident people feel using AI.

In everyday life
Voice assistants that answer questions and hold short dialogues use conversational AI.

Copilot Systems — *I, C*

What it means
Copilot systems are AI tools that work alongside people, offering suggestions rather than taking full control.

Why it matters
They change work from doing everything manually to supervising and refining AI output.

In everyday life
A coding assistant that suggests lines of code is a copilot system.

Cross-Lingual Models — *C, B*

What it means
These are AI systems that understand and work across different languages.

Why it matters
They support global communication and make AI more inclusive.

In everyday life
Translation apps rely on cross-lingual models to convert text between languages.

Code Generation — *D, I*

What it means
Code generation is when an AI system writes computer programs or parts of programs based on instructions in plain language.

Why it matters
It can speed up software development and help non-programmers create digital tools, but it also raises questions about accuracy and responsibility when mistakes happen.

In everyday life
A developer asking an AI to create a simple website form is using code generation.

Content Moderation with AI — *F, I*

What it means
This is the use of AI to detect and manage harmful or inappropriate content such as hate speech, scams, or explicit material.

Why it matters
AI moderation helps platforms scale, but mistakes can silence legitimate voices or miss harmful content if not carefully supervised.

In everyday life
Social media platforms using AI to flag abusive comments are applying content moderation with AI.

Creative AI — *D*

What it means
Creative AI refers to systems that help generate art, music, writing, or design ideas.

Why it matters
It is changing how creativity is expressed and who can participate in creative work.

In everyday life
An AI tool that helps design a logo or compose background music is creative AI.

D

Data Augmentation — *B*

What it means
Data augmentation means creating additional training examples by slightly changing existing data, such as rotating images or rephrasing text.

Why it matters
It helps AI systems learn better when real-world data is limited.

In everyday life
An image-recognition system trained on slightly altered photos uses data augmentation.

Data Labeling for AI — *B, G*

What it means
Data labeling is the process of tagging information so AI systems can learn from it, such as marking images with “cat” or “dog.”

Why it matters
Labels shape what AI learns, and poor labeling can lead to errors or bias.

In everyday life
Human reviewers tagging customer emails to train a support chatbot are doing data labeling.

Data Mining with AI — *B, I*

What it means
Data mining with AI involves analysing large amounts of data to discover patterns or trends.

Why it matters
It helps organisations make better decisions but can raise privacy concerns if misused.

In everyday life
Retailers analysing purchase data to predict demand are using AI-driven data mining.

Dataset — *B*

What it means
A dataset is a collection of information used to train or test AI systems.

Why it matters
The quality and diversity of a dataset strongly affect how fair and accurate an AI system becomes.

In everyday life
Photos used to train a face-recognition system form its dataset.

Dataset Shift — *E*

What it means
Dataset shift happens when the data an AI sees in real life is different from the data it was trained on.

Why it matters
It can quietly reduce accuracy and reliability over time.

In everyday life
A spam filter trained on old email patterns may struggle with new types of scams.

Deep Learning — *A, B*

What it means
Deep learning is a type of machine learning that uses layered neural networks to recognise complex patterns.

Why it matters
It powers many of today's major AI breakthroughs in speech, vision, and language.

In everyday life
Voice assistants that understand spoken commands rely on deep learning.

Deepfake — *D, F, K*

What it means
A deepfake is highly realistic fake audio, video, or images created using AI.

Why it matters
Deepfakes can be used for entertainment, but also for fraud, misinformation, and reputational harm.

In everyday life
A fake video showing a public figure saying something they never said is a deepfake.

Deployment Pipeline — *I*

What it means
A deployment pipeline is the process used to move an AI system from development into real-world use.

Why it matters
Poor deployment can introduce errors, security gaps, or unexpected behaviour.

In everyday life
Updating a chatbot and rolling it out to all users involves a deployment pipeline.

Diffusion Model — *H, D*

What it means
A diffusion model is a type of AI that creates images or media by gradually refining random noise into a clear result.

Why it matters
It has become a leading method behind modern image-generation tools.

In everyday life
Text-to-image tools that create artwork from descriptions often use diffusion models.

Distillation — *B, I*

What it means
Distillation is a way to transfer knowledge from a large AI model to a smaller, more efficient one.

Why it matters
It allows powerful AI to run on smaller devices with lower costs.

In everyday life
A slimmed-down chatbot on a phone may be distilled from a larger system.

Document Summarisation — *D, C*

What it means
This is when AI shortens long texts into key points while keeping the main meaning.

Why it matters
It saves time but can miss nuance if used without review.

In everyday life
An AI tool that condenses a long report into bullet points is doing document summarisation.

Domain Adaptation — *B*

What it means
Domain adaptation means adjusting an AI system trained in one area to work well in another.

Why it matters
It helps AI stay useful when applied to new industries or contexts.

In everyday life
A general language model adapted for legal writing is using domain adaptation.

Dual-Use Risk in AI — *F, G*

What it means
Dual-use risk refers to AI tools that can be used for both helpful and harmful purposes.

Why it matters
It makes safety planning more complex because good technology can be misused.

In everyday life
Image-editing AI that helps artists but can also create fake IDs shows dual-use risk.

Dynamic Prompting — *C*

What it means
Dynamic prompting means changing the instructions given to an AI while it is being used, based on context or user needs.

Why it matters
It makes AI more flexible but can introduce inconsistency if not managed carefully.

In everyday life
A chatbot that adjusts its tone depending on whether you ask for help or advice uses dynamic prompting.

E

Edge AI — *I*

What it means
Edge AI refers to running AI directly on devices like phones or cameras instead of in distant data centres.

Why it matters
It improves speed and privacy by keeping data closer to the user.

In everyday life
A smartphone that recognises faces without sending images to the cloud uses edge AI.

Edge Case — *E*

What it means
An edge case is a rare or unusual situation that an AI system may not handle well.

Why it matters
These cases often reveal hidden weaknesses in AI systems.

In everyday life
A self-driving car struggling with unusual road signs is facing an edge case.

Emergent Behaviour — *E, J, K*

What it means
Emergent behaviour occurs when an AI system shows abilities that were not specifically planned by its designers.

Why it matters
It can lead to innovation but also unexpected risks.

In everyday life
A chatbot unexpectedly learning to solve puzzles without being trained to do so shows emergent behaviour.

Emergent Capabilities — *J, K*

What it means
These are new skills that appear as AI systems become larger and more complex.

Why it matters
They challenge our ability to predict what advanced AI might do.

In everyday life
An AI suddenly becoming good at translation without special training shows emergent capability.

Embedding — *C, B, I*

What it means
Embeddings are numerical representations that capture the meaning of words or images so AI can compare them.

Why it matters
They power search, recommendations, and question-answering systems.

In everyday life
A tool that finds “documents similar to this one” uses embeddings.

Energy Use of AI (Compute Footprint) — *G, I*

What it means
This refers to the electricity and computing power needed to train and run AI systems.

Why it matters
High energy use raises cost and environmental concerns.

In everyday life
Large data centres running AI models consume significant power.

Ensemble Models — *B*

What it means
Ensemble models combine several AI systems to make more reliable predictions.

Why it matters
They reduce the chance of a single model's mistake dominating results.

In everyday life
Weather forecasts often combine multiple models to improve accuracy.

Epistemic Uncertainty — *E*

What it means
This is uncertainty caused by a lack of knowledge, not randomness.

Why it matters
Recognising uncertainty helps avoid overconfidence in AI outputs.

In everyday life
A medical AI saying "I'm not sure" about a rare condition reflects epistemic uncertainty.

Error Rate — *E*

What it means
The frequency with which an AI system makes mistakes.

Why it matters
Understanding error rates helps people judge whether an AI tool is safe for important decisions.

In everyday life
A speech-to-text system mishearing words shows its error rate.

Ethical Design in AI — *F, G*

What it means
Ethical design in AI means building systems with fairness, safety, and respect for people in mind from the start.

Why it matters
Design choices can prevent harm before it happens.

In everyday life
Including diverse voices when designing a hiring AI tool is ethical design.

Ethics in AI — *F, G*

What it means
Ethics in AI focuses on the moral questions raised by using intelligent systems, such as fairness, privacy, and responsibility.

Why it matters
Ethical failures can damage lives and trust in technology.

In everyday life
Debates about whether AI should be used in surveillance are ethical questions.

Evaluation of AI (Evals) — *E, F, I*

What it means
Evals are structured tests that measure how well an AI system performs and behaves.

Why it matters
They help organisations detect safety risks and performance drops.

In everyday life
Testing a chatbot to see whether it gives unsafe advice is part of evaluation.

Explainability in AI — *F, G*

What it means
Explainability means how clearly an AI system's decisions can be understood by people — so users are not left guessing why something happened.

Why it matters
It builds trust and is essential in sensitive areas like healthcare or finance.

In everyday life
A system that explains why it denied a loan shows explainability.

Exposure Bias — *E*

What it means
Exposure bias happens when an AI learns from ideal examples but struggles in real-world situations.

Why it matters
It can make systems seem better in testing than in practice.

In everyday life
A translation tool working well in demos but poorly with slang shows exposure bias.

F

Fairness in AI — *F, G*

What it means
Fairness in AI means ensuring systems treat people equally and do not disadvantage certain groups.

Why it matters
Unfair systems can deepen social inequalities.

In everyday life
A hiring AI designed to give equal opportunity across genders reflects fairness.

Factuality in AI Outputs — *E, F*

What it means
Factuality refers to how often an AI provides correct, reliable information.

Why it matters
Incorrect information can mislead people and cause harm.

In everyday life
Checking an AI's answer against a trusted source helps confirm factuality.

Federated Learning — *B, F, I*

What it means
Federated learning trains AI systems using data stored on many devices without moving the data to a central server.

Why it matters
It improves privacy while still allowing AI to learn.

In everyday life
Smartphones improving predictive typing without sending your texts to a central database use federated learning.

Feedback Loops in AI Systems — *F, B*

What it means
Feedback loops occur when AI systems learn from the results of their own decisions.

Why it matters
They can improve systems but also reinforce mistakes or bias if not monitored.

In everyday life
A recommendation system that keeps showing the same type of content because you once clicked on it is part of a feedback loop.

Few-Shot Learning — *C, B*

What it means
Few-shot learning is when an AI learns a new task from just a few examples.

Why it matters
It makes AI more flexible and easier to adapt.

In everyday life
Showing an AI three sample emails and asking it to write similar ones uses few-shot learning.

Fine-Tuning — *B, I*

What it means
Fine-tuning is adjusting a general AI system to perform better on a specific task or topic.

Why it matters
It improves relevance but must be done carefully to avoid introducing new risks.

In everyday life
Customising a chatbot for legal advice is an example of fine-tuning.

Foundation Model — *A, B*

What it means
A foundation model is a large AI system trained broadly and then adapted for many uses.

Why it matters
It reduces the need to build separate models for every task.

In everyday life
One AI system powering writing, translation, and summarisation is a foundation model.

Frontier Models — *J, K*

What it means
Frontier models are the most advanced AI systems currently available.

Why it matters
They shape future possibilities and risks in AI.

In everyday life
The latest generation of chatbots often uses frontier models.

Fraud Detection with AI — *I*

What it means
This is the use of AI to identify suspicious activity such as scams or identity theft.

Why it matters
It protects people and organisations from financial harm.

In everyday life
Banks flagging unusual transactions rely on AI-based fraud detection.

Function Calling / Tool Use — *I, C, F*

What it means
This is when an AI system can use other digital tools, such as calendars, databases, or software applications, to complete tasks.

Why it matters
It makes AI more useful but also increases safety and security responsibilities.

In everyday life
An AI that books a meeting for you is using tool-calling capabilities.

G

Generalisation — *E, B*

What it means
Generalisation is an AI system's ability to apply what it has learned to new situations it has not seen before.

Why it matters
An AI that only works well on familiar examples but fails in new situations can be unreliable and risky.

In everyday life
A photo app that recognises your dog even in new settings is showing good generalisation.

Generative AI — *A, D*

What it means
Generative AI refers to systems that create new content such as text, images, music, or video instead of just analysing information.

Why it matters
It is transforming how people work, learn, and express creativity — while also raising questions about authenticity and ownership.

In everyday life
An AI tool that writes a draft email or creates an image from a description is using generative AI.

Governance in AI — *G*

What it means
Governance in AI refers to the rules, policies, and oversight used to guide how AI is developed and used responsibly.

Why it matters
Good governance helps ensure AI benefits society rather than causing harm or unfairness.

In everyday life
A company setting clear rules about when staff can use AI tools is practising AI governance.

Green AI — *G, I*

What it means
Green AI focuses on developing AI systems that use less energy and have a smaller environmental impact.

Why it matters
As AI grows, its energy use affects costs and climate goals.

In everyday life
Choosing efficient AI models that run on local devices supports Green AI.

Ground Truth — *B*

What it means
Ground truth is the correct, real-world information used to train or test AI systems.

Why it matters
Without accurate ground truth, AI systems learn the wrong lessons.

In everyday life
Correctly labelled medical images provide ground truth for training diagnostic tools.

Grounding (in Sources) — *E, I*

What it means
Grounding means connecting AI responses to reliable information sources instead of letting the system guess.

Why it matters
It reduces errors and increases trust in AI outputs.

In everyday life
A chatbot that cites official health guidelines instead of inventing advice is using grounding.

Guardrails in AI — *F, I*

What it means
Guardrails are safety limits that prevent AI systems from producing harmful or inappropriate outputs.

Why it matters
They help keep AI useful without putting people at risk.

In everyday life
A chatbot refusing to give instructions for illegal activities is following guardrails.

H

Hallucination — *E, F*

What it means
Hallucination occurs when an AI confidently gives false or made-up information.

Why it matters
People may trust incorrect answers and make poor decisions as a result.

In everyday life
A chatbot inventing a fake legal case is an example of hallucination.

Hard vs Soft Prompts — *C, B*

What it means
Hard prompts are fixed instructions built into an AI system, while soft prompts are flexible user instructions given during use.

Why it matters
Understanding this difference helps explain why some AI behaviour can be changed and some cannot.

In everyday life
System rules that always apply are hard prompts; your typed instructions are soft prompts.

Heuristics in AI — *H*

What it means
Heuristics are simple rules or shortcuts AI uses to make decisions when perfect answers are not possible.

Why it matters
They make systems faster but can introduce bias or errors.

In everyday life
A navigation app choosing a route based on typical traffic patterns uses heuristics.

Hidden Prompt / System Prompt — *C, F*

What it means
A system prompt is a hidden set of instructions that guides how an AI behaves behind the scenes.

Why it matters
These instructions shape tone, safety, and limits — often without users realising it.

In everyday life
If a chatbot always stays polite and avoids certain topics, its system prompt is guiding that behaviour.

Homomorphic Encryption for AI — *F, I*

What it means
This is a way for AI to work with encrypted data without seeing the raw information.

Why it matters
It allows AI to help analyse sensitive data while protecting privacy.

In everyday life
A hospital analysing patient data without exposing identities benefits from this technology.

Human-in-the-Loop — *F, I*

What it means
Human-in-the-loop means keeping people involved in AI decision-making rather than letting systems act alone.

Why it matters
It reduces risk and ensures accountability.

In everyday life
A doctor reviewing AI medical suggestions before acting is part of a human-in-the-loop system.

Human Oversight in AI — *F, G*

What it means
Human oversight refers to people monitoring AI systems and stepping in when needed.

Why it matters
It ensures that AI remains a tool, not an unchecked authority — so that important decisions always stay with people, not machines.

In everyday life
Moderators reviewing flagged content after AI detection provide human oversight.

Hybrid AI (Rules + Learning) — *H, I*

What it means
Hybrid AI combines traditional rule-based systems with modern machine learning.

Why it matters
This approach blends reliability with flexibility.

In everyday life
A fraud system using fixed rules plus AI predictions is hybrid AI.

Hyperparameters — *B*

What it means
Hyperparameters are settings that control how an AI model learns — similar to adjusting the speed or intensity of a training programme.

Why it matters
They affect performance, accuracy, and reliability.

In everyday life
Adjusting how fast an AI learns is changing a hyperparameter.

I

Image Generation — *D*

What it means
Image generation is when AI creates pictures from descriptions or ideas.

Why it matters
It changes how visual content is produced and who can create it.

In everyday life
An AI tool that designs posters from text instructions uses image generation.

Image Recognition — *I, H*

What it means
Image recognition allows AI to identify objects, people, or scenes in pictures.

Why it matters
It powers many safety, medical, and consumer applications.

In everyday life
Your phone unlocking by recognising your face uses image recognition.

Incident Response in AI Safety — *F, G*

What it means
This refers to how organisations respond when an AI system causes harm or behaves unexpectedly.

Why it matters
Quick and transparent responses help prevent further damage and rebuild trust.

In everyday life
A company disabling a faulty chatbot after it spreads misinformation is handling an AI incident.

In-Context Learning — *C, B*

What it means
In-context learning means AI adapts based on examples given during a conversation, without permanent retraining.

Why it matters
It makes AI more flexible and responsive.

In everyday life
Showing an AI a few examples of how you like emails written and then asking it to continue is in-context learning.

Inference — *B, I*

What it means
Inference is when an AI system uses what it has learned to make predictions or generate responses.

Why it matters
It is the stage where AI delivers real value to users.

In everyday life
When a chatbot answers your question, it is performing inference.

Inference Cost — *I*

What it means
Inference cost refers to the resources needed each time an AI produces an output.

Why it matters
High costs can limit how widely AI can be used.

In everyday life
A company limiting chatbot use due to expense is managing inference cost.

Inference Latency — *I*

What it means
Inference latency is the delay between asking an AI for something and receiving the result.

Why it matters
Long delays reduce usability and trust because people expect AI tools to respond quickly — especially in time-sensitive situations.

In everyday life
Waiting several seconds for voice assistants to respond shows high latency.

Information Retrieval — *I*

What it means
Information retrieval is finding relevant data from large collections.

Why it matters
It powers search engines and knowledge tools.

In everyday life
Typing a question into a search engine uses information retrieval.

Instruction Tuning — *B, C*

What it means
Instruction tuning is training AI to follow human instructions more accurately.

Why it matters
It makes AI more helpful and safer to use.

In everyday life
A chatbot becoming better at following clear directions after updates reflects instruction tuning.

Integration of AI in Workflows — *I*

What it means
This refers to embedding AI tools into everyday work processes.

Why it matters
Good integration improves efficiency; poor integration causes frustration.

In everyday life
Adding AI to automate meeting notes is workflow integration.

Intellectual Property in AI Outputs — *G, F*

What it means
This concerns who owns content created by AI systems.

Why it matters
Ownership affects artists, businesses, and legal rights.

In everyday life
Questions about who owns an AI-generated logo involve intellectual property issues.

Interpretability in AI — *F, E*

What it means
Interpretability refers to how easily humans can understand how an AI reaches decisions.

Why it matters
Clear understanding builds trust and accountability.

In everyday life
A system that shows which factors influenced a decision improves interpretability.

J

Jailbreak — *F, K*

What it means
A jailbreak is when someone deliberately tries to bypass an AI system's safety rules to make it produce harmful, illegal, or restricted content.

Why it matters
Jailbreaking exposes weaknesses in AI safety systems and can lead to real-world harm, such as spreading misinformation, enabling scams, or encouraging dangerous behaviour.

In everyday life
When people share online tricks to make a chatbot give forbidden advice, they are attempting a jailbreak.

Job Displacement from AI — *G, K*

What it means
Job displacement from AI refers to situations where tasks or roles once done by people are taken over by automated systems.

Why it matters
While AI can increase productivity, it can also create uncertainty for workers and requires societies to think seriously about retraining and new career paths.

In everyday life
Self-checkout machines replacing some cashier roles is an early example of job displacement through automation.

Job Redesign with AI — *G, I*

What it means
Job redesign with AI means reshaping roles so that people work alongside AI rather than being replaced by it.

Why it matters
This approach helps preserve human value while benefiting from automation.

In everyday life
A journalist using AI to draft outlines but focusing more on investigation and storytelling is an example of job redesign.

Justification of AI Decisions — *F, E*

What it means
This refers to explaining why an AI system made a particular choice or recommendation.

Why it matters
Without justification, people may feel powerless or unfairly treated by automated systems.

In everyday life
A bank explaining why an AI rejected a loan application is providing justification of its AI decision.

K

Knowledge Cutoff — *E, C*

What it means
A knowledge cutoff is the point in time after which an AI system no longer has updated information.

Why it matters
Users may wrongly assume AI always has the latest facts, leading to outdated or incorrect decisions.

In everyday life
If a chatbot cannot answer about very recent events, it may be limited by its knowledge cutoff.

Keyword vs Semantic Search — *I, C*

What it means
Keyword search looks for exact words, while semantic search tries to understand the meaning behind a question.

Why it matters
Semantic search gives more helpful results when people don't know the exact terms to use.

In everyday life
Searching "how to fix slow internet" and getting useful advice even without technical wording shows semantic search in action.

L

Large Language Model (LLM) — *A, C*

What it means
A large language model is an AI system trained on vast amounts of text so it can understand and generate human-like language.

Why it matters
LLMs power many modern tools, from chatbots to writing assistants, shaping how people interact with technology.

In everyday life
The AI behind a customer-support chatbot is often a large language model.

Latent Space — *H, D*

What it means
Latent space is the hidden way an AI represents ideas, patterns, or images internally.

Why it matters
It helps explain how AI can create new content that feels original rather than copied.

In everyday life
When an AI blends different art styles into a new image, it is working within its latent space.

Legal Liability in AI Use — *G, F*

What it means
Legal liability in AI use concerns who is responsible when an AI system causes harm or makes a serious mistake.

Why it matters
Clear responsibility is essential for protecting people's rights and ensuring justice.

In everyday life
If an automated system gives harmful financial advice, questions about liability arise immediately.

Local AI Models — *I*

What it means
Local AI models run directly on a person's device instead of relying on cloud servers.

Why it matters
They offer better privacy and faster response times, especially for sensitive data.

In everyday life
A phone that performs voice recognition without sending audio to the internet uses a local AI model.

LoRA (Low-Rank Adaptation) — *B, I*

What it means
LoRA is a method that allows large AI models to be customised efficiently without retraining everything from scratch.

Why it matters
It makes AI adaptation cheaper and more accessible for smaller organisations.

In everyday life
A small business customising a chatbot for customer service using limited resources benefits from LoRA-style techniques.

Long-Context Models — *C, E*

What it means
Long-context models are AI systems that can handle very long conversations or documents without losing track of earlier information.

Why it matters
They improve reliability when AI is used for complex tasks like legal reviews or research.

In everyday life
An AI that remembers details from a long meeting transcript is using long-context capabilities.

Low vs High Temperature — *C, E*

What it means
Temperature is a setting that controls how creative or predictable an AI's responses are.

Why it matters
Understanding this helps users choose between safe, factual answers and more imaginative ones.

In everyday life
Using a low-temperature setting for legal text and a higher one for creative writing shows this difference.

M

Machine Learning — *A, B*

What it means
Machine learning is a way of building computer systems that improve by learning from examples rather than being programmed with fixed rules. Instead of being told exactly what to do, the system finds patterns in data and uses them to make predictions or decisions.

Why it matters
Most modern AI systems rely on machine learning. Understanding this helps people see that AI is shaped by data and experience, not magic or independent thinking.

In everyday life
When your email app learns which messages you usually mark as spam and starts filtering them automatically, it is using machine learning.

Misalignment in AI — *F, J*

What it means
Misalignment happens when an AI system's behaviour does not match human values, intentions, or expectations — even if the system is technically working as designed.

Why it matters
A misaligned AI can produce harmful or confusing outcomes, not because it is broken, but because it optimises for the wrong goals.

In everyday life
An AI customer-service bot that focuses only on ending conversations quickly instead of actually helping people is showing signs of misalignment.

Misinformation with AI — *F, G, K*

What it means
This refers to false or misleading information that is created or spread using AI tools, sometimes unintentionally and sometimes on purpose.

Why it matters
AI can generate content at scale, which makes it easier for misinformation to spread quickly and widely, affecting public trust and decision-making.

In everyday life
An AI-generated news article that looks real but contains false claims is an example of AI-driven misinformation.

Model — *A*

What it means
In AI, a model is the trained system that has learned patterns from data and can now make predictions or generate responses.

Why it matters
The quality of an AI system depends heavily on the model behind it — how it was trained, what data it used, and how it is updated.

In everyday life
The "brain" behind a chatbot or image generator is its AI model.

Model Cards — *G, F*

What it means
Model cards are documents that explain how an AI model was built, what it is good at, and where its limits are.

Why it matters
They help users and organisations understand risks and use AI responsibly instead of blindly trusting it.

In everyday life
A company publishing details about what its AI can and cannot do is using a model card approach.

Model Collapse — *E, K*

What it means
Model collapse happens when AI systems are trained too much on content created by other AI systems, causing quality and originality to decline over time.

Why it matters
If unchecked, it could make future AI less reliable and less creative.

In everyday life
A writing tool that slowly starts producing repetitive or generic text because it learns mostly from AI-generated material shows signs of model collapse.

Model Monitoring — *I, E*

What it means
Model monitoring means regularly checking how an AI system performs after it has been deployed.

Why it matters
AI behaviour can change over time as data and conditions shift, so ongoing monitoring helps catch problems early.

In everyday life
A company tracking whether its recommendation system is still fair and accurate months after launch is doing model monitoring.

Model Scaling Laws — *J, K*

What it means
These describe how AI performance tends to improve as models get larger and are trained with more data and computing power.

Why it matters
They explain why bigger AI systems often become more capable — but also why they become more expensive and harder to manage safely.

In everyday life
When newer AI tools feel noticeably more capable than older versions, scaling laws help explain why.

Multimodal AI — *A, D, C*

What it means
Multimodal AI can understand and work with different types of information at the same time, such as text, images, and audio.

Why it matters
It brings AI closer to how humans experience the world, making systems more useful and intuitive.

In everyday life
An app that lets you upload a photo and ask questions about it is using multimodal AI.

Multi-Agent Systems — *H, I*

What it means
These are setups where several AI systems work together, each handling part of a task or making its own decisions.

Why it matters
They can solve complex problems but also make accountability more challenging — because it can be harder to tell which system made which decision.

In everyday life
Multiple AI bots coordinating to manage customer support across different channels form a multi-agent system.

N

Narrow AI (Weak AI) — *A*

What it means
Narrow AI refers to systems designed to perform one specific task very well, such as translation or image recognition.

Why it matters
Most AI today is narrow AI, not human-like intelligence — and understanding this prevents unrealistic expectations.

In everyday life
A navigation app that gives directions but cannot hold a conversation is narrow AI.

Natural Language Processing (NLP) — *C, A*

What it means
NLP is the field of AI that focuses on understanding and generating human language.

Why it matters
It enables chatbots, translation tools, and voice assistants to work.

In everyday life
When your phone turns speech into text, NLP is behind it.

Neural Network — *H, B*

What it means
A neural network is a type of AI model inspired by how the human brain processes information, using layers of connected units to learn patterns.

Why it matters
Neural networks are the foundation of many powerful AI systems today.

In everyday life
Image recognition tools that identify faces or objects rely on neural networks.

Non-Determinism in AI — *E*

What it means
Non-determinism means an AI system may give different answers to the same question at different times.

Why it matters
It explains why AI responses are not always perfectly predictable — and why people should treat AI answers as helpful suggestions, not absolute truth.

In everyday life
Asking the same chatbot question twice and getting slightly different answers shows non-determinism.

Novelty in AI Content — *D, G*

What it means
This refers to how original or new AI-generated content appears.

Why it matters
It affects how people judge creativity, copyright, and value.

In everyday life
An AI artwork that feels genuinely new rather than copied raises questions about novelty.

O

Open-Source AI Models — *I, G*

What it means
Open-source AI models are systems whose code and design are shared publicly so others can study, improve, or adapt them.

Why it matters
They promote transparency and innovation, but also raise concerns about misuse.

In everyday life
Developers using a free, publicly available AI model to build new apps are relying on open-source AI.

Open vs Closed AI Models — *G, K*

What it means
This contrasts AI systems that are openly shared with those controlled by a single company.

Why it matters
It shapes who controls AI power and how widely benefits and risks are distributed.

In everyday life
A public AI tool anyone can customise versus a proprietary one locked behind a company platform illustrates this difference.

Optical Character Recognition (OCR) — *I*

What it means
OCR is a technology that allows computers to read text from images or scanned documents.

Why it matters
It makes paper information searchable and usable in digital systems.

In everyday life
Scanning a receipt and turning it into editable text uses OCR.

Out-of-Distribution Data (OOD) — *E*

What it means
This refers to data that looks very different from what an AI system was trained on.

Why it matters
AI systems often perform poorly on unfamiliar data, which can lead to unexpected errors.

In everyday life
A voice assistant struggling with strong accents it was not trained on is facing out-of-distribution data.

Output Filtering in AI — *F, I*

What it means
Output filtering means checking and adjusting what an AI produces to remove harmful or inappropriate content.

Why it matters
It protects users and organisations from unintended harm.

In everyday life
A chatbot that blocks offensive language before sending a response is using output filtering.

Overfitting — *E, B*

What it means
Overfitting happens when an AI system learns its training data too well and struggles with new situations.

Why it matters
It can make systems appear accurate in testing but unreliable in real life.

In everyday life
A handwriting app that works perfectly for one person but fails with others shows overfitting.

Oversight Boards for AI — *G*

What it means
These are groups that review how AI systems are developed and used to ensure they meet ethical and safety standards.

Why it matters
They provide independent judgment and accountability.

In everyday life
A hospital setting up a committee to review AI use in patient care is forming an oversight board.

P

Parameters (Model Parameters) — *B*

What it means
Parameters are the internal settings inside an AI model that determine how it behaves. They

are adjusted during training so the system learns patterns from data, much like tuning thousands or millions of tiny knobs.

Why it matters
The number and quality of parameters strongly influence how capable an AI system becomes. Larger models with more parameters can often handle more complex tasks, but they also require more computing power and careful control.

In everyday life
When a new version of an AI tool feels "smarter" than the old one, it often has more or better-trained parameters behind it.

Personalisation with AI — *I, F*

What it means
Personalisation with AI means tailoring content, services, or experiences to individual users based on their behaviour, preferences, or history.

Why it matters
Personalisation can improve convenience, but it also raises concerns about privacy and fairness if people feel they are being watched or treated differently without knowing why.

In everyday life
A shopping app suggesting products based on what you have viewed before is using AI-driven personalisation.

Policy in AI — *G*

What it means
Policy in AI refers to the rules, guidelines, and laws that shape how AI should be developed and used in society.

Why it matters
Good policies help protect people's rights, encourage responsible innovation, and prevent misuse of powerful technologies.

In everyday life
Government rules about how AI can be used in schools or workplaces are examples of AI policy in action.

Post-Training — *B*

What it means
Post-training is the stage where an AI system is refined after its initial learning phase, often to improve safety, accuracy, or usefulness.

Why it matters
Many important improvements — such as better behaviour and clearer responses — happen during post-training, not just during initial development.

In everyday life
When a chatbot becomes more polite or helpful after an update, it has likely gone through post-training.

Pre-Training — *B*

What it means
Pre-training is the early phase where an AI system learns general patterns from large amounts of data before being specialised.

Why it matters
This stage gives AI a broad foundation of knowledge that later makes it adaptable to many tasks.

In everyday life
An AI that already understands language before being trained for customer support relies on pre-training.

Precision & Recall — *E*

What it means
Precision and recall are two ways of measuring how well an AI system makes correct decisions. Precision looks at how often its positive results are right, while recall looks at how many true cases it successfully finds.

Why it matters
Balancing these measures is crucial in sensitive areas like healthcare or fraud detection, where missing a case or flagging too many false ones can both cause harm.

In everyday life
A spam filter that blocks most junk mail but occasionally blocks real emails shows the trade-off between precision and recall.

Privacy in AI — *F, G*

What it means
Privacy in AI concerns how personal data is collected, stored, and used when training and operating AI systems.

Why it matters
Poor privacy practices can lead to data misuse, loss of trust, and harm to individuals.

In everyday life
Being able to opt out of data collection when using an AI service reflects attention to privacy.

Prompt — *C*

What it means
A prompt is the instruction or question you give an AI system to guide what it produces.

Why it matters
The way a prompt is written can strongly affect the quality and safety of the AI's response.

In everyday life
Asking "Explain this simply" instead of "Explain this" is changing the prompt to get a better result.

Prompt Engineering — *C*

What it means
Prompt engineering is the skill of writing clear and effective instructions so AI systems give better answers.

Why it matters
It helps people get more useful, accurate, and safe outputs from AI tools.

In everyday life
Learning how to phrase questions to get clearer answers from a chatbot is practising prompt engineering.

Prompt Injection — *F, C*

What it means
Prompt injection is when someone tries to trick an AI system into ignoring its safety rules by hiding harmful instructions inside normal-looking input.

Why it matters
It exposes security weaknesses and can lead to dangerous or misleading outputs.

In everyday life
A user sneaking hidden commands into a form to manipulate an AI response is attempting prompt injection.

Prompt Libraries — *C, I*

What it means
Prompt libraries are collections of ready-made prompts that help users interact more effectively with AI tools.

Why it matters
They lower the learning curve and help people use AI more productively.

In everyday life
Using a saved template to ask an AI to summarise documents is relying on a prompt library.

Proprietary AI Models — *I, G*

What it means
Proprietary AI models are systems owned and controlled by a company, rather than shared openly.

Why it matters
They often offer strong performance but raise questions about transparency and public oversight.

In everyday life
A company using its own private AI system that others cannot access is using a proprietary model.

Public Interest AI — *G*

What it means
Public interest AI refers to using artificial intelligence to benefit society as a whole rather than only private profit.

Why it matters
It encourages the development of AI that supports education, health, safety, and social good.

In everyday life
AI tools that help predict natural disasters or improve access to public services are examples of public interest AI.

Q

Quantisation — *I, B*

What it means
Quantisation is a technique that makes AI models smaller and faster by simplifying how they store numbers.

Why it matters
It allows powerful AI to run on everyday devices instead of only on expensive servers.

In everyday life
An AI app that works smoothly on your phone instead of needing the cloud may rely on quantisation.

Quality Assurance in AI — *I, F*

What it means
Quality assurance in AI involves checking that systems behave reliably, safely, and as intended.

Why it matters
It helps prevent harmful mistakes and builds confidence in AI tools.

In everyday life
Testing a chatbot before releasing it to customers is part of AI quality assurance.

R

Rate Limits (APIs) — *I*

What it means
Rate limits control how often a user or system can access an AI service within a certain time.

Why it matters
They protect systems from overload and misuse.

In everyday life
If an app temporarily blocks too many rapid requests, it is enforcing rate limits.

Real-Time AI — *I*

What it means
Real-time AI refers to systems that respond instantly or nearly instantly to changing information.

Why it matters
Fast responses are critical in areas like navigation, safety, and customer support.

In everyday life
Traffic apps that update routes as conditions change use real-time AI.

Recommendation Systems — *H, I*

What it means
These are AI systems that suggest products, content, or actions based on patterns in user behaviour.

Why it matters
They shape what people see and buy, influencing culture and choices.

In everyday life
Streaming services recommending movies are using recommendation systems.

Reinforcement Learning — *B*

What it means
Reinforcement learning is a method where AI learns by trying actions and receiving feedback about what works best.

Why it matters
It enables AI to improve through experience, especially in complex environments.

In everyday life
Game-playing AIs that get better after each match use reinforcement learning.

Reliability in AI — *E, F*

What it means
Reliability in AI means that systems behave consistently and as expected over time.

Why it matters
Unreliable AI can undermine trust and lead to risky decisions.

In everyday life
A navigation app that gives dependable directions every day shows reliability.

Responsible AI — *F, G*

What it means
Responsible AI refers to developing and using AI in ways that respect people, society, and the law.

Why it matters
It ensures that progress does not come at the cost of harm or injustice.

In everyday life
A company openly explaining how it protects user data in its AI tools is practising responsible AI.

Responsible Deployment of AI — *I, G*

What it means
This means introducing AI systems carefully, with testing, oversight, and safeguards in place.

Why it matters
Rushed deployment can create serious risks.

In everyday life
Piloting an AI system with a small group before a full rollout shows responsible deployment.

Retrieval-Augmented Generation (RAG) — *I, C, E*

What it means
RAG combines AI generation with real information sources, so responses are based on facts rather than guesses.

Why it matters
It improves accuracy and reduces hallucinations.

In everyday life
A chatbot that pulls answers from company documents instead of inventing them is using RAG.

Risk Assessment in AI — *G, F*

What it means
Risk assessment in AI means identifying possible harms before deploying a system.

Why it matters
It helps prevent problems rather than reacting after damage is done.

In everyday life
A school evaluating privacy risks before using AI tools for students is doing risk assessment.

Robustness in AI — *E, F*

What it means
Robustness refers to how well an AI system performs under unexpected conditions.

Why it matters
Robust systems are safer and more trustworthy.

In everyday life
A voice assistant that still works in noisy environments shows robustness.

Rule-Based AI (Expert Systems) — *H*

What it means
Rule-based AI uses fixed instructions created by humans rather than learning from data.

Why it matters
It offers transparency but less flexibility than modern learning-based systems.

In everyday life
A tax program following strict legal rules is an example of rule-based AI.

Runtime Safety — *F, I*

What it means
Runtime safety means protecting people from harm while an AI system is actively running.

Why it matters
Some risks only appear during real use, not in testing.

In everyday life
A chatbot blocking harmful content in real time demonstrates runtime safety.

S

Safety by Design in AI — *F, G*

What it means
Safety by design means building AI systems with protection and risk prevention in mind from the very beginning, rather than adding safety features later as an afterthought.

Why it matters
When safety is built in early, AI systems are less likely to cause harm, spread misinformation, or be misused. This approach reduces risks before they reach the public.

In everyday life
An AI tool that is designed to refuse dangerous medical advice from day one reflects safety by design.

Safety Cases for AI — *G, F*

What it means
A safety case is a clear explanation of why an AI system is considered safe enough to use, based on evidence and testing.

Why it matters
It helps organisations prove they have taken risks seriously instead of simply hoping for the best.

In everyday life
Before introducing AI into a hospital system, leaders may require a safety case showing it will not endanger patients.

Self-Supervised Learning — *B*

What it means
Self-supervised learning is a method where AI learns from data without needing humans to label everything manually.

Why it matters
It allows AI systems to learn from vast amounts of information more efficiently and cheaply.

In everyday life
An AI that learns language patterns just by reading large volumes of text is using self-supervised learning.

Semi-Supervised Learning — *B*

What it means
Semi-supervised learning combines a small amount of labelled data with a large amount of unlabelled data to train AI systems.

Why it matters
It balances accuracy with practicality when labelling data is expensive or slow.

In everyday life
A system learning from a few labelled medical images and many unlabelled ones uses semi-supervised learning.

Sentiment Analysis — *I, C*

What it means
Sentiment analysis is when AI detects emotions or opinions in text, such as whether something sounds positive, negative, or neutral.

Why it matters
It helps organisations understand public opinion, customer satisfaction, and social trends.

In everyday life
Companies analysing social media comments to see how people feel about a product use sentiment analysis.

Shielding / Sandboxing of AI Tools — *F, I*

What it means
Shielding or sandboxing means isolating AI systems so they can be tested or used safely without affecting real systems.

Why it matters
It reduces the risk of accidents, security breaches, or unexpected behaviour.

In everyday life
Testing a new chatbot in a private environment before letting customers use it is sandboxing.

Small Language Models (SLMs) — *A, I*

What it means
Small language models are lighter versions of large AI systems designed to run efficiently on limited hardware.

Why it matters
They make AI more accessible and affordable, especially for small organisations and personal devices.

In everyday life
A voice assistant working entirely on your phone may use a small language model.

Social Impact of AI — *G*

What it means
This refers to how AI affects society, including jobs, education, relationships, and equality.

Why it matters
Technology shapes culture and opportunity, so understanding social impact helps guide responsible choices.

In everyday life
Discussions about AI changing the job market reflect concern about its social impact.

Speech Recognition — *D, I*

What it means
Speech recognition allows AI to turn spoken words into text or commands.

Why it matters
It makes technology more accessible, especially for people with disabilities.

In everyday life
Dictating a message instead of typing uses speech recognition.

Speech Synthesis — *D, I*

What it means
Speech synthesis is when AI produces spoken language from text.

Why it matters
It supports accessibility and natural interaction with technology.

In everyday life
Navigation apps reading directions aloud use speech synthesis.

Synthetic Data — *B, F*

What it means
Synthetic data is artificial data created by AI to mimic real information without using actual personal data.

Why it matters
It helps protect privacy while still allowing AI to learn.

In everyday life
Training a medical AI on simulated patient records instead of real ones uses synthetic data.

Synthetic Media — *D, F*

What it means
Synthetic media refers to audio, video, or images created by AI rather than captured from the real world.

Why it matters
It expands creativity but also raises concerns about authenticity and trust.

In everyday life
An AI-generated voice in a video is an example of synthetic media.

System Prompt — *C, F*

What it means
A system prompt is a hidden set of instructions that guides how an AI behaves, including tone, limits, and safety rules.

Why it matters
These unseen instructions strongly shape what users experience without them realising it.

In everyday life
A chatbot always responding politely and refusing certain requests is following its system prompt.

Scenario Testing for AI — *F, I*

What it means
Scenario testing involves checking how an AI system behaves in difficult or unusual situations.

Why it matters
It helps identify weaknesses before real harm occurs.

In everyday life
Testing how a customer-service bot reacts to angry users is scenario testing.

T

Temperature (Model Setting) — *C, E*

What it means
Temperature is a setting that controls how creative or predictable an AI's responses are.

Why it matters
Understanding this helps users choose between safer, factual outputs and more imaginative ones.

In everyday life
Using a low temperature for legal writing and a higher one for creative stories shows how this setting works.

Text Generation — *D*

What it means
Text generation is when AI creates written content such as emails, stories, or reports.

Why it matters
It changes how people write and work, saving time but also raising questions about originality.

In everyday life
An AI drafting a cover letter for a job application is generating text.

Token — *C*

What it means
A token is a small unit of text that AI uses to process language, such as a word or part of a word.

Why it matters
Tokens affect how much text AI can handle and how costs are calculated.

In everyday life
Longer prompts using more words require more tokens.

Tool Use by AI Agents — *I, F*

What it means
This is when AI systems can interact with other software tools to complete tasks.

Why it matters
It makes AI more powerful but also increases the need for security and oversight, because these systems can now take actions — not just give advice."

In everyday life
An AI scheduling meetings through your calendar is using tools on your behalf.

Training — *B*

What it means
Training is the process of teaching an AI system by exposing it to data so it can learn patterns.

Why it matters
The quality of training determines how reliable and fair an AI system becomes.

In everyday life
An image-recognition system learning from thousands of labelled photos is being trained.

Training Data — *B, F*

What it means
Training data is the information used to teach an AI system how to perform tasks.

Why it matters
Biased or poor-quality data leads to biased or unreliable AI behaviour.

In everyday life
Photos used to train facial recognition systems are training data.

Transfer Learning — *B*

What it means
Transfer learning means reusing knowledge from one AI task to help with another.

Why it matters
It saves time and resources when building new systems.

In everyday life
An AI trained on general language and then adapted for legal writing uses transfer learning.

Transformer Models — *H, B*

What it means
Transformer models are a type of AI design that helps computers understand and work with language by looking at how words relate to each other in context.

Why it matters
They power most modern language tools and chatbots.

In everyday life
The AI behind many chat systems uses transformer models.

Transparency in AI — *F, G*

What it means
Transparency in AI means being open about how systems work and how decisions are made.

Why it matters
It builds trust and helps people challenge unfair outcomes — especially when an AI decision affects jobs, money, or access to services.

In everyday life
A company explaining how its AI evaluates job applications shows transparency.

Turing Test — *K, J*

What it means
The Turing Test is a thought experiment where a machine tries to appear human in conversation.

Why it matters
It shaped early thinking about what it means for machines to seem intelligent.

In everyday life
People joking that a chatbot "sounds human" are referencing the idea behind the Turing Test.

U

Underfitting — *E*

What it means
Underfitting happens when an AI system is too simple to learn patterns in data properly.

Why it matters
It leads to poor performance even on familiar tasks.

In everyday life
A handwriting app that struggles to recognise even clear writing may be underfitting.

Uncertainty Estimation — *E, F*

What it means
Uncertainty estimation is when AI systems express how confident they are in their answers.

Why it matters
It helps users know when to trust results and when to seek human judgment.

In everyday life
A medical AI saying "I am not confident about this diagnosis" is estimating uncertainty.

Unsupervised Learning — *B*

What it means
Unsupervised learning is when AI finds patterns in data without being told what to look for.

Why it matters
It helps discover hidden trends humans might miss.

In everyday life
A music app grouping similar songs without human labels uses unsupervised learning.

Use Policies for AI — *G, F*

What it means
Use policies define how people are allowed to use AI systems responsibly.

Why it matters
They prevent misuse and protect users and organisations.

In everyday life
A company setting rules about not using AI to handle sensitive data is applying use policies.

User Consent in AI Data Use — *G, F*

What it means
This refers to getting clear permission from people before using their data in AI systems.

Why it matters
Consent respects privacy and builds trust.

In everyday life
Being asked to agree before your data is used to train an AI tool is an example of user consent.

V

Vector Databases — *I, C*

What it means
A vector database stores information in a way that allows AI systems to search by meaning rather than just exact words. Instead of matching text literally, it helps AI find ideas that are similar in meaning.

Why it matters
This makes search, recommendations, and question-answering far more accurate, especially when people don't know the exact wording to use.

In everyday life
When you search for "how to stay healthy" and the system finds articles about exercise, sleep, and diet—even if those exact words weren't used—it is likely using a vector database.

Verification & Validation in AI — *I, F*

What it means
Verification and validation mean checking that an AI system works the way it is supposed to and is safe for real-world use. Verification asks, "Did we build the system correctly?" Validation asks, "Did we build the right system?"

Why it matters
Without proper checks, AI tools can quietly cause harm, spread misinformation, or fail in critical situations.

In everyday life
Testing a medical AI thoroughly before using it in hospitals is an example of verification and validation.

Video Generation — *D*

What it means
Video generation is when AI creates moving images or short films based on text descriptions or examples.

Why it matters
It changes how video content is produced and raises new questions about trust, authenticity, and misuse.

In everyday life
An AI tool that creates a short promotional video from a written script is using video generation.

Virtual Agents — *I, C*

What it means
Virtual agents are AI systems that act on behalf of people to perform tasks such as booking appointments, answering questions, or managing schedules.

Why it matters
They reduce routine work but must be carefully controlled so they act responsibly and securely.

In everyday life
A digital assistant that automatically schedules meetings is acting as a virtual agent.

W

Weak AI (Narrow AI) — *A*

What it means
Weak AI refers to systems designed to perform specific tasks very well, without true understanding or general intelligence.

Why it matters
Most AI today is weak AI, and recognising this helps people avoid unrealistic expectations about what machines can actually do.

In everyday life
A translation app that works well for language but cannot reason about topics is an example of weak AI.

Workforce Impact of AI — *G, K*

What it means
This term describes how AI affects jobs, skills, and the way people work.

Why it matters
AI can both create new opportunities and disrupt existing roles, making adaptation and retraining essential.

In everyday life
Employees learning new digital skills because of automation are experiencing workforce impact from AI.

Workflow Automation with AI — *I*

What it means
Workflow automation with AI means using intelligent systems to streamline tasks and processes that once required manual effort.

Why it matters
It saves time and reduces errors but also changes how work is organised.

In everyday life
An AI that automatically routes customer requests to the right department is automating workflows.

Watermarking of AI Content — *F, G*

What it means
Watermarking is a way of marking AI-generated content so people can recognise that it was created by a machine.

Why it matters
It helps fight misinformation and supports transparency in a world where AI content looks increasingly real.

In everyday life
A label showing that an image was created by AI is a form of watermarking.

Z

Zero-Shot Learning — *C, B*

What it means
Zero-shot learning is when an AI system can perform a task it was never specifically trained to do, using general knowledge and reasoning instead.

Why it matters
It makes AI more flexible and useful in new situations without needing constant retraining.

In everyday life
Asking an AI to write a poem in a style it has never practised and getting a reasonable result is an example of zero-shot learning.

Future Edition Updates & User Submissions

Artificial intelligence terminology continues to evolve as new concepts emerge and existing terms are used in different ways across industries and regions.

Future editions of this guide may:

- include newly introduced terms
- refine explanations based on evolving usage
- improve clarity based on reader feedback

AISF welcomes suggestions for additional terms or areas where further clarification may be helpful.

Submissions may be considered for future editions, subject to editorial review and alignment with AISF's commitment to clarity, neutrality, and accessibility.

Acknowledgements

This guide reflects a synthesis of widely used terminology across global AI discourse, informed by publicly available materials, industry usage, academic references, and cross-system AI research.

AISF acknowledges the broader global community — including researchers, practitioners, educators, policymakers, and everyday users — whose collective use of language continues to shape how artificial intelligence is understood.

All content has been editorially structured and presented under AISF governance.

ABOUT

AI SOURCED FACTS (AISF) PTE. LTD.

AISF is a Singapore-headquartered institution dedicated to structured reasoning, responsible AI navigation, and governance-informed adoption of artificial intelligence systems.

AISF operates with a capability-first, vendor-neutral posture. Its publications do not rank platforms, endorse providers, or promote specific technologies. Instead, AISF develops structured frameworks that help individuals, professionals, and institutions reason clearly before integrating AI into operational, strategic, or educational environments.

AISF's work spans whitepapers, applied insight books, education instruments, governance architectures, and structured research initiatives. These outputs are informed by cross-system AI research methodologies and reflect globally observed usage patterns at the time of publication. Human accountability remains central across all AISF frameworks.

AISF does not provide regulatory, legal, financial, investment, or compliance advice. Its publications are designed to support structured thinking, proportionate governance, and disciplined evaluation of AI capabilities prior to deployment or reliance.

As artificial intelligence systems continue to evolve, AISF's focus remains constant: clarity before integration, governance proportionate to capability, and long-term institutional resilience in the age of AI.

BACK COVER

Artificial intelligence is now part of everyday conversation. Terms are used widely — often inconsistently — across media, workplaces, and public discourse.

This guide provides a clear and accessible reference to commonly used AI terms, helping readers understand what they mean, how they are used, and where confusion may arise.

It is designed to help readers:

- understand commonly used AI terminology in plain language
- distinguish between technical meaning and everyday usage
- identify where terms are used loosely or inconsistently
- build confidence in engaging with AI-related discussions

This publication is intended for:

- individuals seeking clarity without technical complexity
- professionals navigating AI-related conversations
- students and educators building foundational understanding
- organisations supporting broader AI awareness

AISF publications are developed using cross-system research methodologies and reflect globally observed AI usage patterns at the time of publication. They are designed to support clarity, responsible integration, and informed human judgment.

Artificial intelligence systems assist.
Responsibility remains human.

www.aisourcedfacts.com