

Ai Glossary Beginner Cheat Sheet

Plain-English AI reference guide | cyberside.ai

Term	Plain-English meaning
Accuracy	How often an AI system gives correct results for a defined task. A model can be accurate on one task and weak on another.
Active learning	A training approach where a model asks for labels or feedback on the examples that would help it learn most.
Agent	An AI system designed to take steps toward a goal, often by using tools, following rules, or working through multiple tasks. Think “intern with a checklist,” not “magic robot employee.”
AGI	Artificial General Intelligence: a hypothetical AI with broad, flexible abilities across many kinds of tasks. It is not a settled benchmark or a product label you should accept uncritically.
AI literacy	The practical ability to understand what AI can and cannot do, use it thoughtfully, and spot when human judgment is needed.
AI model	A mathematical system trained to recognize patterns and produce outputs such as text, images, predictions, or classifications.
Algorithm	A defined set of steps for solving a problem. AI systems use algorithms, but not every algorithm is AI.
Analytics	The process of examining data to understand patterns, trends, or performance. AI can assist with analytics but does not replace checking the underlying numbers.
Annotation	A label or note added to data so humans or models can understand what it represents, such as marking a photo as containing a dog.
Artificial intelligence	A broad field of computing focused on systems that perform tasks associated with learning, perception, language, reasoning, or decision support.
Attribute	A field or characteristic attached to a record, object, person, or item in data.
Backpropagation	A training method that adjusts a neural network after it compares its prediction with the correct answer.
Benchmark	A standardized test used to compare models or systems on specific tasks. A strong benchmark score is useful evidence, not a guarantee of real-world performance.
Bias	A systematic pattern that can produce unfair, distorted, or unrepresentative results. Bias can come from data, model design, prompts, or how people use a system.
Big data	Very large, fast-moving, or complex data that requires specialized storage, processing, or analysis methods.
Binary classification	A prediction task with two possible categories, such as spam versus not spam.
Black box	A system whose internal reasoning is hard to inspect or explain clearly. The phrase does not mean the system is mysterious magic; it means visibility is limited.
Chain of thought	A model’s internal step-by-step reasoning process. It can be useful to ask for a concise explanation or a checked solution, but do not treat a detailed-looking rationale as proof.
Chatbot	A conversational software interface that responds to messages. Some chatbots are AI-powered; some follow fixed scripts.
Classification	Assigning an item to one of several categories, such as labeling a support ticket by topic.

Term	Plain-English meaning
Classifier	A model or rule system that performs classification.
Confidence score	A number expressing how strongly a model favors a prediction. It is not the same as a guarantee that the prediction is correct.
Copyright	A legal right that can protect original creative works. AI use can raise copyright questions, especially when training, uploading, or publishing content; get qualified advice for high-stakes situations.
CRM	Customer Relationship Management software: a system for tracking customer interactions, leads, and follow-ups.
Cross-validation	A way to test a model by training and evaluating it on different slices of the same dataset.
Data augmentation	Creating modified versions of training examples to help a model handle more variation, such as rotating images or rephrasing text.
Data drift	A change over time in the real-world data a model sees, which can reduce performance after deployment.
Data labeling	The process of adding useful tags or answers to data for training or evaluation.
Data leakage	Sensitive, private, or restricted data being exposed to people or systems that should not receive it.
Data lineage	A record of where data came from, how it changed, and where it was used.
Data minimization	A privacy practice of collecting and using only the data genuinely needed for a purpose.
Dataset	An organized collection of data used for analysis, training, testing, or reporting.
Decision tree	A model or rule structure that makes decisions through a series of branching questions.
Deep learning	A branch of machine learning that uses multi-layer neural networks to learn complex patterns from data.
Deployment	Putting a model, application, or workflow into real use. Deployment should include monitoring, access controls, and a rollback plan.
Emergent behavior	A capability or behavior that appears when a system becomes larger or is used in a new way, even if it was not directly programmed as a feature.
Entity extraction	Identifying things such as names, organizations, dates, places, or product codes in text.
False negative	A result that incorrectly says something is absent or safe when it is actually present or risky.
False positive	A result that incorrectly flags something as present, suspicious, or likely when it is not.
Feature	An input variable or measurable property used by a model, such as purchase amount, time of day, or word count.
Feature engineering	Selecting, transforming, or creating useful input features for a model.
Fine-tuning	Additional training that adjusts a base model for a narrower style, task, or domain. It is different from simply giving an AI a detailed prompt.
Foundation model	A large, broadly trained model that can be adapted to many tasks through prompting, fine-tuning, or tools.
Generalization	A model's ability to work on new examples that were not part of its training data.
Generative adversarial network	A model architecture in which one network generates examples and another tries to detect them, historically used for synthetic media.

Educational reference. Verify high-stakes decisions with qualified professionals.