A BEGINNER'S GUIDE

What Is AI?

A plain English guide to the technology behind chatbots, recommendations and creative tools.

The short answer

Artificial intelligence, usually shortened to AI, is technology that performs tasks we associate with human thinking. It can recognise patterns, understand or produce language, make predictions and help choose what to do next. AI does not think or understand exactly as a person does. It uses software and data to produce useful outputs.

You may already use it

A streaming service suggests a programme. Your phone groups photos by face. A map estimates traffic and recommends a faster route.

AI is a broad label

Some AI follows hand-written rules. Other AI learns patterns from examples. Generative AI can create new text, pictures or sound. These are related approaches, not synonyms.

How machines learn from examples

A lot of modern AI uses machine learning. Instead of writing a separate instruction for every possible situation, people give a system examples. The system adjusts its internal settings to find patterns that help it make predictions on new examples.

Machine learning

Machine learning is a way to build AI by learning from data. Imagine showing an email filter many messages marked as unwanted or wanted. It learns clues such as suspicious links or repeated phrases, then estimates whether a new message is spam.

Deep learning

Deep learning is a kind of machine learning that uses many connected layers of computation, called a neural network. It can learn complicated patterns from large collections of examples, such as the shapes and sounds that help a system recognise speech.

A useful way to picture the difference

Think of learning to identify birds. A simple machine-learning system might use a set of measured features, such as size, beak shape and colour. A deep-learning system can work from the pixels in the photographs and discover useful visual features through its layers. Deep learning is powerful for complex data, but often needs more examples and computing.

Keep the relationship straight

Deep learning sits inside machine learning, and machine learning sits inside the wider field of AI. Not every AI tool learns, and not every learning system creates content.

Generative AI and the models behind it

Generative AI is AI designed to produce new material in response to a request. It can draft an email, suggest a recipe, summarise notes, create an illustration or help write computer code. The result is newly assembled, but it is shaped by patterns learned from training examples and by the instructions it receives.

What is a model?

A model is the learned system that turns an input into an output. Think of it as a recipe made from patterns rather than a fixed list of answers. A language model learns patterns in text and uses them to estimate what words or pieces of text are likely to follow a prompt.

One family of tools, many kinds of output

Text

Draft a polite reply, outline a topic or explain a difficult idea in simpler words.

Images

Turn a written description into a picture, concept sketch or visual variation.

Other tasks

Some tools work with speech, music, video or code. Others combine several formats.

A prompt is the instruction or question you give the tool. More context can make the result more relevant: for example, name the audience, the format and any limits. A prompt guides the model, but it does not guarantee a correct answer.

Training and inference, simply explained

Two words help explain what happens before and during an AI tool's use. Training is how a model learns patterns from data. Inference is what happens when a trained model uses those patterns to respond to a new input.

01 Gather examples

For a cat-photo classifier, people provide many photos with labels. Data quality and variety matter: confusing or unbalanced examples can teach poor patterns.

02 Train the model

During training, the system compares its guesses with the examples and adjusts internal values to improve. The result is a trained model that has learned statistical patterns.

03 Use it

Give it a new photo and it estimates whether it shows a cat. Producing this response is inference. The model is applying what it learned, not repeating the training process from scratch.

A familiar classroom analogy

Imagine practising with marked maths questions. Studying the worked examples and learning the methods is like training. Answering a fresh question later is like inference. The analogy is not exact, because a model adjusts mathematical settings rather than learning with human understanding, but it captures the difference between learning and applying.

One important distinction

The examples used to train a model are not necessarily the same as the information it can look up while answering. Some products connect to search or other data sources; others respond only from their learned patterns and the conversation context.

What AI can get wrong, plus a mini glossary

AI can be helpful, but it is not a reliable authority on everything. A fluent answer can still be incorrect, incomplete or out of date. Treat its output as a starting point, not automatic proof.

Why it can fail

A model may produce a confident-sounding false statement, sometimes called a hallucination. It can reflect gaps or bias in the examples it learned from, misunderstand a vague request, or miss a recent change. A generated answer is not the same as a checked fact.

Use it thoughtfully

Check important claims against dependable sources, especially health, legal, financial or safety advice. Avoid sharing private or sensitive information unless you understand how the service handles it. Consider fairness, copyright and who is responsible for a decision made with AI.

Glossary

AI: technology that performs tasks associated with human intelligence.

Machine learning: a way of building AI systems that learn patterns from data.

Deep learning: machine learning that uses layered neural networks to learn complex patterns.

Generative AI: AI that produces new content, such as text, images or audio.

Model: a learned system that maps an input to a prediction or output.

Training: the process of adjusting a model using examples and data.

Inference: using a trained model to produce an answer or prediction for a new input.

Prompt: the question, instruction or context given to a generative AI tool.