



What is a neural network?

KS3

KS4

Ages 11-16 ⌚ 4 min read

Every time your phone unlocks by recognising your face, every time a spam filter catches a suspicious email, every time you ask a voice assistant something — there's a neural network at work. They're the engine behind modern AI. But the name is slightly misleading. They're inspired by brains, but they don't work like one.

What is a neuron, really?

In your brain, neurons are cells that receive signals from other neurons, and fire their own signal if the combined input is strong enough. A simple concept, repeated 86 billion times with trillions of connections, produces thought, memory, and consciousness. Artificial neural networks borrow the basic concept: nodes (like neurons) receive inputs, apply a mathematical operation, and pass the result to the next layer of nodes.

💡 Imagine a chain of people passing a message. The first person hears several sounds and shouts the loudest one to the next person. That person combines what they hear from several people and shouts a summary. By the end of the chain, a complex raw input has been transformed into a simple answer. Each person in the chain is like an artificial neuron — they receive inputs, combine them, and pass something forward. The "learning" happens by adjusting how much each person weighs the different inputs they receive.

How does it learn?

A neural network starts with random weights — essentially random decisions about how important each connection is. It then processes a training example and makes a prediction. That prediction is compared to the correct answer, and the error is used to slightly adjust the weights throughout the network. This process, called backpropagation, is repeated millions of times across millions of examples. Gradually the network gets better at its task — not because anyone programmed the rules, but because the weights evolved to capture the patterns in the data.

Layers and depth

Simple neural networks have three layers: input, hidden, and output. "Deep" neural networks (the "deep" in deep learning) have many hidden layers — sometimes hundreds. Each layer learns to detect increasingly abstract patterns: in an image-recognition network, early layers might detect edges, middle layers might detect shapes, and later layers might detect faces or objects. The depth allows the network to learn complex, hierarchical representations.

What can they do?

Neural networks now exceed human performance at image recognition, many language tasks, protein structure prediction, playing complex games, and medical diagnosis in specific domains. They have weaknesses too — they can be fooled by subtly altered inputs, they require vast amounts of training data, and they can perpetuate biases in that data. But they represent the most significant advance in artificial intelligence in decades.