

Open Weights section I had to remove:

But before looking at the larger competition, I want to explain the term at the center of the article: **open-weight AI**.

Open-weight AI

An AI model contains billions—or sometimes trillions—of numerical settings called weights. During training, those numbers are adjusted as the model learns patterns in language, images, computer code and other information.

The weights are essentially the model's trained machinery.

With a closed model such as the main versions of ChatGPT or Claude, we can use the model, but we cannot download the whole system and operate it ourselves. The company controls the servers, updates, safeguards and price.

With an open-weight model, the developer releases those trained numerical settings. A company, university or government can download the model, run it on its own computers and adapt it for a particular purpose.

A hospital, for example, could operate one internally so confidential patient information never leaves the hospital. A manufacturer could customize one using its own technical manuals and machinery records.

Open-weight, however, does not necessarily mean open-source.

With traditional open-source software, the human-written source code is available for people to inspect and modify.

An open-weight AI company may release the trained model while withholding the training data, the complete training code, the process used to select the data and some of its safety testing.

One simple analogy is this:

Open-source gives you the recipe, the ingredients and the cooking instructions.

Open-weight gives you the finished meal—and perhaps the trained chef—but may not tell you exactly how the meal was created.

Alternate Explanation

You may have heard people say that large language models, or LLMs, are “trained.”

During training, the model learns patterns and relationships from enormous amounts of text. For example, it may learn that:

- cats are related to animals
- Paris is related to France
- the word “doctor” often appears near “hospital”
- “dog” is more closely related to “puppy” than to “airplane”

The model learns relationships like these at an unimaginably complex level.

It stores what it has learned in billions of numerical settings called **weights**. A single weight does not represent one fact or idea. Instead, each weight affects the strength of a connection inside the model. Individually, the weights mean almost nothing. Collectively, they form the model’s learned structure.

Now suppose you ask:

What is the capital of France?

The model processes your question using its billions of weights. It then calculates which word, or more precisely which **token**, is most likely to come next.

In a simplified example, its probabilities might look something like this:

Paris — 99.999%

London — 0.0004%

Berlin — 0.0002%

Madrid — 0.0001%

The probabilities are not the weights. The probabilities are produced by the model using the weights.

It is also important to remember that these are probabilities for the next token, not a guaranteed measure of factual accuracy. A model can sometimes assign a high probability to an answer that sounds convincing but is incorrect.

One way to think about this is to imagine a student who has spent years studying a subject. The student's knowledge is not stored as one fact in one brain cell. It is distributed across countless connections formed through learning.

The weights in an AI model work in a somewhat similar way. They represent the strength of billions of connections inside the model's artificial neural network.

When an AI company releases an **open-weight model**, it makes this trained network available for others to download. A business, university or government may then run the model on its own computers, adapt it and build new systems on top of it.

Open-weight does not necessarily mean open-source.

An open-weight company may share the trained model without sharing:

- all the data used to train it
- the complete training process
- all the software used to create it
- the decisions made about safety and data selection

A simple way to remember the distinction is:

Open-weight means the company shares the trained model. Open-source usually means it also shares much more of the recipe used to create it.

Or, using a human analogy:

Open-weight is like receiving a trained brain. Open-source is closer to receiving the training materials, methods and instructions as well.