

MistralAI on Microsoft Azure



Azure



MISTRAL
AI_



MISTRAL AI

OVERVIEW

Mistral AI, established in 2023, quickly emerged as a significant player in the AI landscape, dedicated to developing cutting-edge language models and technologies.

Partnering with leading technology platforms like Microsoft Azure, Mistral AI ensures that our state-of-the-art models are **accessible to developers, businesses, and researchers worldwide**. This helps integrate AI into diverse applications, from enterprise solutions to creative endeavors.

With a commitment to continuous improvement and customer-focused innovation, Mistral AI is dedicated to enhancing the capabilities of AI to **solve real-world problems** effectively and ethically.

*Data as of July 20th, 2024

MODELS

Mistral 7B, Release date: 27/09/2023

Mixtral 8x7B, Release date: 9/12/2023

Mixtral 8x22B, Release date: 10/04/2024

Codestral 22B, Release date: 29/05/2024


OPEN WIEGHT

Mistral Small, Release date: 26/02/2024

Mistral Large, Release date: 26/02/2024

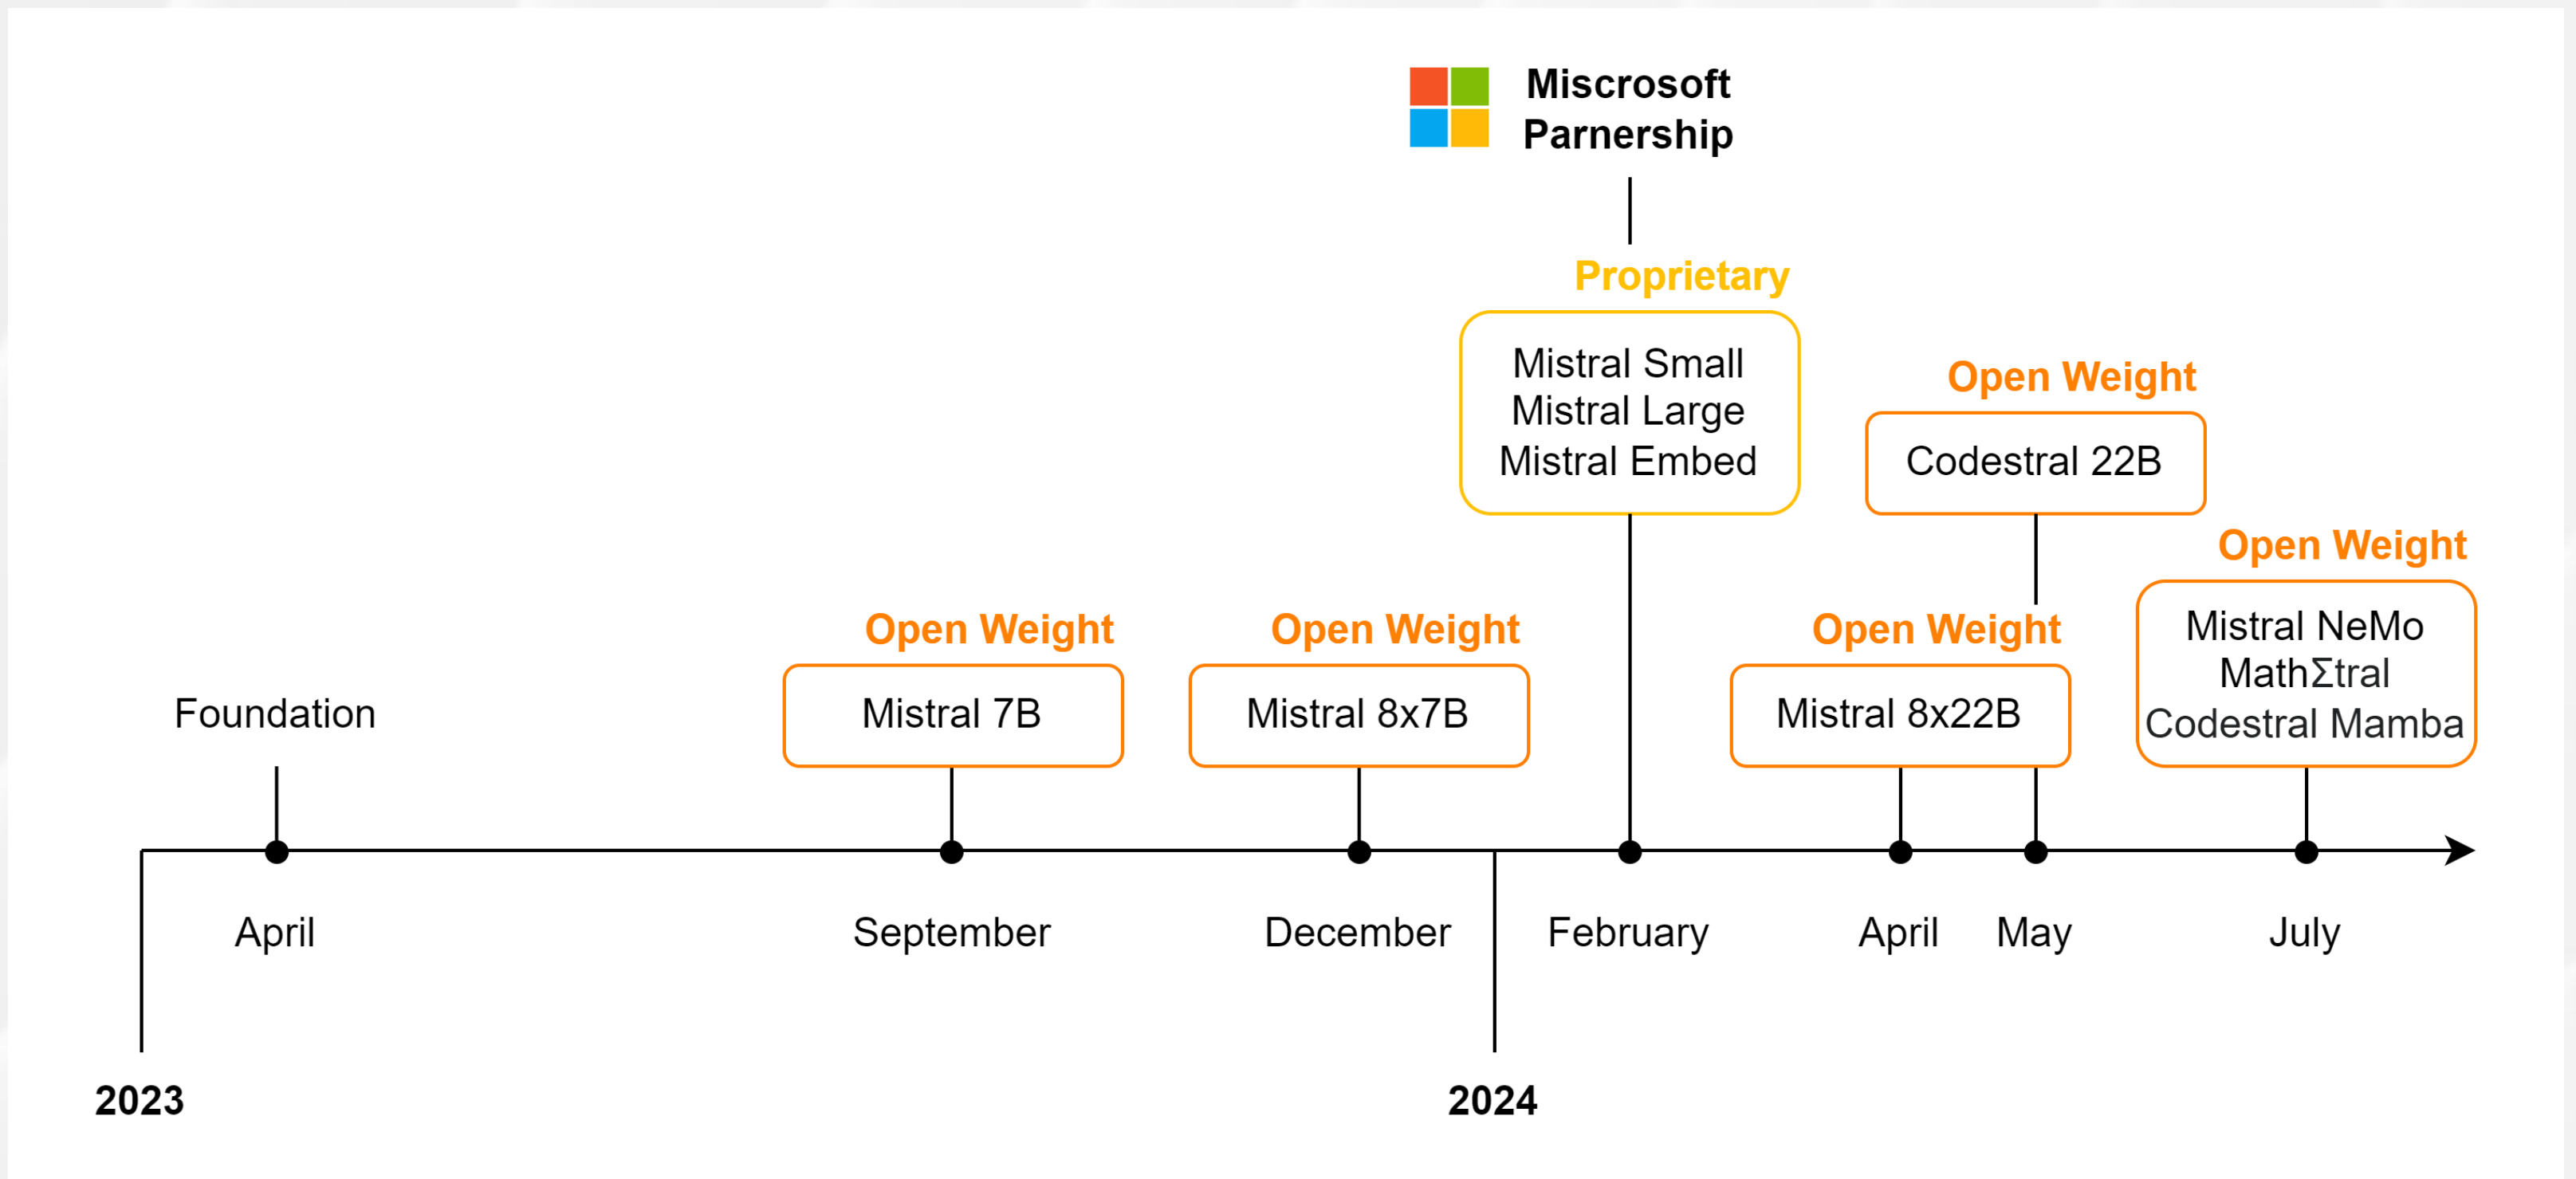
Mistral Embed, Release date: 26/02/2024

API



Let's take a closer look at the **Mistral Large** model...

TIMELINE



What does **Open Weight** and **Proprietary** mean?

An **open weight model** refers to a machine learning model where the weights (the parameters learned during training) are publicly accessible and can be freely used, modified, and shared. This contrasts with **proprietary models** where the weights are kept secret or restricted by licenses.

*Data as of July 20th, 2024

MODEL ARCHITECTURE



MAMBA

This architecture enhances how the model accesses its memory, optimizes how it processes information, and speeds up its response to queries. It makes the language model faster, more responsive, and capable of handling large context without slowing down.



MoE (MIXTURE OF EXPERTS)

This architecture is composed of a collection of smaller, specialized programs known as “**expert**”. Each expert excels in specific types of tasks or questions. When a question is posed to the model, it identifies and utilizes the best experts suited for that particular query.

The primary benefit of this system is its efficiency. By engaging only a subset of experts at any given time, it uses fewer computational resources. This makes the model faster and more cost-effective.

*Data as of July 20th, 2024

KEY FEATURES



CAPABLE MODEL

Mistral Large is the most powerful model of the Mistral AI family. With top-tier reasoning for high-complexity tasks, this model is natively fluent in English, French, Spanish, German, and Italian. It is natively capable of function calling.



FAST MODEL

Getting a quality response quickly is what users expect from Gen AI-based applications. Tests on a variety of questions have shown that Mistral Large can generate text very quickly compare to other large-size models.



AVAILABLE ON AZURE

Mistral AI has partnered with Microsoft to offer Mistral Large, and other models, on Azure. This collaboration ensures a seamless integration for companies and organizations already within the Microsoft ecosystem.

*Data as of July 20th, 2024



Contact us!



Fabio Solazzo

f.solazzo@reply.it



Elena Baiocco

e.baiocco@reply.it