

Meta releases Llama 2 weights free for research and commercial use

Source copy kept by wheresthe.ai for the story "Meta releases Llama 2 weights free for research and commercial use" (<https://wheresthe.ai/d/meta-llama-2-open-release-2023-07/>).

Source: <https://about.fb.com/news/2023/07/llama-2/>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: 36c9c05937cac711704a42e3aa58a7b0ad59b1a83a91ea33db1b117e7e92eba1

[Back to Newsroom](#)

Meta

July 18, 2023 July 18, 2023

Listen to Article 10 Pause 10 0.25x 0.5x 1x 1.25x 1.5x 2x

Takeaways

Today, we're introducing the availability of Llama 2, the next generation of our open source large language model.

Llama 2 is free for research and commercial use.

Microsoft and Meta are expanding their longstanding partnership, with Microsoft as the preferred partner for Llama 2.

We're opening access to Llama 2 with the support of a broad set of companies and people across tech, academia, and policy who also believe in an open innovation approach to today's AI technologies.

We're committed to building responsibly and are providing resources to help those who use Llama 2 do so too.

Recent breakthroughs in AI, and generative AI in particular, have captured the public's imagination and demonstrated what those developing these technologies have long known — they have the potential to help people do incredible things, create a new era of economic and social opportunities, and give individuals, creators, and businesses new ways to express themselves and connect with people.

We believe an open approach is the right one for the development of today's AI models, especially those in the generative space where the technology is rapidly advancing. By making AI models available openly, they can benefit everyone. Giving businesses, startups, entrepreneurs, and researchers access to tools developed at a scale that would be challenging to build themselves, backed by computing power they might not otherwise access, will open up a world of opportunities for them to experiment, innovate in exciting ways, and ultimately benefit from economically and socially.

And we believe it's safer. Opening access to today's AI models means a generation of developers and researchers can stress test them, identifying and solving problems fast, as a community. By seeing how these tools are used by others, our own teams can learn from them, improve those tools, and fix vulnerabilities.

Meta has put exploratory research, open source, and collaboration with academic and industry partners at the heart of our AI efforts for over a decade. We've seen first-hand how innovation in the open can lead to technologies that benefit more people. Dozens of large language models have already been released and are driving progress by developers and researchers. They're being used by businesses as core ingredients for new generative AI-powered experiences. We've been blown away by the huge demand for Llama 1 from researchers — with more than 100,000 requests for access to the large language model — and the amazing things they've achieved by building on top of it .

We're now ready to open source the next version of Llama 2 and are making it available free of charge for research and commercial use . We're including model weights and starting code for the pretrained model and conversational fine-tuned versions too. As Satya Nadella announced on stage at Microsoft Inspire, we're taking our partnership to the next level with Microsoft as our preferred partner for Llama 2 and expanding our efforts in generative AI. Starting today, Llama 2 is available in the Azure AI model catalog, enabling developers using

Microsoft Azure to build with it and leverage their cloud-native tools for content filtering and safety features. It is also optimized to run locally on Windows, giving developers a seamless workflow as they bring generative AI experiences to customers across different platforms. Llama 2 is available through Amazon Web Services (AWS), Hugging Face, and other providers too.

People and businesses have benefited from the longstanding partnership between Microsoft and Meta. Together we've introduced an open ecosystem for interchangeable AI frameworks, and we've co-authored research papers to advance the state of the art in AI. We've collaborated to scale the adoption of PyTorch — today's leading AI framework created by Meta and the AI community — on Azure, and we're among the founding members of the PyTorch Foundation. Microsoft and Meta recently joined a cohort of supporters that endorse the Partnership on AI's framework for collective action in the creation and sharing of synthetic media. Our partnership extends outside of AI and into the metaverse too to deliver immersive experiences for the future of work and play.

Now, with this expanded partnership, Microsoft and Meta are supporting an open approach to provide increased access to foundational AI technologies to the benefits of businesses globally. It's not just Meta and Microsoft that believe in democratizing access to today's AI models. We have a broad range of diverse supporters around the world who believe in this approach too — including companies that have given us early feedback and are excited to build new products with Llama 2, cloud providers that will include Llama 2 in their offerings for customers, research institutions who are collaborating with us on the safe and responsible deployment of large generative models, and people across tech, academia, and policy who see the benefits as we do.

A Focus on Responsibility

Our open source approach promotes transparency and access. We know that while AI has brought huge advances to society, it also comes with risk. We are committed to building responsibly and are providing a number of resources to help those who use Llama 2 do so too.

Red-Teaming Exercises: Our fine-tuned models have been red-teamed — tested for safety — through internal and external efforts. The team worked to generate adversarial prompts to facilitate model fine-tuning. In addition, we commissioned third parties to conduct external adversarial testing across our fine-tuned models to similarly identify gaps in performance. These safety fine-tuning processes are iterative; we will continue to invest in safety through fine-tuning and benchmarking and plan to release updated fine-tuned models based on these efforts.

Transparency Schematic: We explain our fine-tuning and evaluation methods for the model and identify its shortcomings. Our transparency schematic, which is located within the research paper, discloses known challenges and issues we've experienced and provides insight into mitigations taken and future ones we intend to explore.

Responsible Use Guide: We created this guide as a resource to support developers with best practices for responsible development and safety evaluations. It outlines best practices reflective of current, state-of-the-art research on responsible generative AI discussed across the industry and the AI research community.

Acceptable Use Policy: We put a policy in place that prohibits certain use cases to help ensure that these models are being used fairly and responsibly.

Meta has also created new initiatives to harness the insight and creativity of individuals, researchers, and developers around the world to get feedback on how the models are performing and how they might be improved.

Open Innovation AI Research Community: Today, we also launched a new partnership program for academic researchers that aims to deepen our understanding of the responsible development and sharing of large language models. Researchers may apply to join a community of practitioners to share learnings on this important topic, and the community will form a research agenda to pursue going forward.

Llama Impact Challenge: We want to activate the community of innovators who aspire to use Llama to solve hard problems. We are launching a challenge to encourage a diverse set of public, non-profit, and for-profit entities to use Llama 2 to address environmental, education and other important challenges. The challenge rules will be available prior to the start of it.

Conclusion

Throughout our company's history, we've experienced the benefits of an open source approach when innovating in other areas of the business. Our engineers developed and shared frameworks that are now industry standards

— like React, a leading framework for making web and mobile applications, and PyTorch, which is now the leading framework for AI. These became commonly used infrastructure for the entire technology industry. We believe that openly sharing today's large language models will support the development of helpful and safer generative AI too.

We look forward to seeing what the world builds with Llama 2.

Categories :

Meta

Technology and Innovation

Tags:

AI

DOWNLOAD ALL IMAGES

Share this article

Related News