March 2024

Meta's Response to the European Commission's Call for Contributions on Competition in Generative AI

I. Introductory Remarks

Meta appreciates the opportunity to respond to the European Commission's Call for Contributions on Competition in Generative AI. Artificial Intelligence (AI) has been at the heart of our work for over a decade, and is now front and center in the experiences available across our family of apps and devices. Recent advancements in generative AI have captured the public's imagination and promise to redefine how society can address major societal problems and create economic opportunities across a wide range of industries.

The excitement around generative AI has unleashed a flood of investment and innovation from both established and new players. At Meta, we believe that to fully realize its potential AI must be democratized across a wide range of stakeholders. Open sourcing AI models, if incentivised by policymakers, can lead to safer products and ensure state-of-art AI creates opportunities for everyone, especially for small developers with limited resources or even individuals.

With our submission, we would like to share our experience in building generative AI with the European Commission and other policymakers and highlight the benefits open models provide in democratizing the access and use of AI across everyone in society.

II. Our journey on AI

At Meta, we know that people use our technologies to get closer to what they care about – connecting with family, friends, customers, and new experiences. AI gives us the ability to deliver meaningful product advancements that drive greater engagement while unlocking new experiences for people and businesses alike. At the highest level, generative AI uses machine learning algorithms trained on real-world data to create images, sounds, videos, immersive 3D environments, and to write prose and to code. The

technology helps machines better understand the world and translate that understanding into creative expression. We see amazing possibilities ahead for users, creators, and businesses. In fact, Gartner estimates[1] that by 2025, generative AI will be producing 20 percent of all consumer-facing data.

Since the earliest days of News Feed in 2006, AI has been fundamental to our family of apps-whether it's feed ranking, community discovery, personalized ads, or content moderation. It would be hard to identify a single product across Meta – consumer or business facing – that doesn't depend on the work we're doing in AI. It powers everything from the systems that rank content in our apps and the tools advertisers use to reach customers, to how we keep hate speech off our services. AI helps deliver relevant content to people who use our services, helps small businesses better reach customers, and supports identifying and combating harmful content.

Our long-term investments in AI and the AI technologies that underpin our platforms today have been benefiting our users and business partners in a variety of ways for years now. AI powers our discovery engine for better content ranking, improves outcomes for consumers and advertisers in our ads systems, and will play a key role in how and what we build within the metaverse. For example, women's jewelry brand Jenny Bird when running a Meta Advantage+[2] shopping campaign together with its usual ad campaigns, the company saw a 14% lower cost[3] per purchase compared to its usual campaigns alone, and 17% higher conversions.

While generative AI is still at a nascent stage of development, it is already enabling new forms of connection and expression for people across our family of apps and devices. We're invested in generative AI to build new product experiences that will power expression and creativity for people, businesses, and creators. Many of these experiences are underpinned by technology from our Llama 2[4] and Emu[5] foundation models. We have

---

[1] Gartner (2021). 5 Impactful Technologies from Gartner Emerging Technologies and Trends Impact Radar for 2022: https://www.gartner.com/en/articles/5-impactful-technologies-from-the-gartner-emerging-technologies-and-trends-impact-radar-for-2022

[2] About Meta Advantage: https://www.facebook.com/business/help/733979527611858

[3] Jenny Bird. Boosting online sales using Meta Advantage+ shopping campaigns: https://www.facebook.com/business/success/jenny-bird

[4] Discover the power of Llama: https://ai.meta.com/llama/

[5] Emu: Enhancing Image Generation Models Using Photogenic Needles in a Haystack: https://ai.meta.com/research/publications/emu-enhancing-image-generation-models-using-photogenic-needles-in-a-haystack/

a strong track record in research breakthroughs across modalities with state-of-the-art models for language, image, and video generation. For example:

1. We are building a range of text, image, and video generation models that can generate rich 3D content, starting with objects and characters and all the way up to entire worlds in the metaverse.
2. We're investing in creating a new class of AI agents that bring a more human-centric approach to helping consumers, creators drive new forms of connection and expression. These generative AI-powered experiences are under development in varying phases. Our AI characters who have a diverse array of interests and personalities are available for you to chat with on WhatsApp, Messenger, and Instagram. Many of the AI characters will look familiar because we partnered with well known public figures to play them – but these familiar faces are representing the AI characters, not themselves.
3. We are launching AI stickers that let people create customized stickers in seconds from a simple text prompt and AI media editing where you can reimagine your photos with unique new styles and backgrounds.
4. Meta has also built an AI studio—a platform that supports the creation of our AIs and we plan to make it available for people outside Meta.

Given that this is early days, and the rapid pace of developments, it's difficult to predict the direction for generative AI - the foundational technology itself, the downstream applications and use cases that will catch on, the business models, and even how these experiences will be delivered. While we expect to see continued development of larger models, we've recently seen greater focus on smaller models, improving performance and lowering compute requirements (and associated costs). Similarly, we've seen firms working to run AI models on smaller consumer devices, such as the Ray-Ban Meta Smart Glasses[6].

III. Developing AI Responsibly

As with all cutting edge technology, AI has promise and risk. We're incorporating lessons we've learned over the last decade into our new AI features and models, and we're building safeguards into our AI features and models before we launch them. We also stress test our products to improve safety performance and regularly work with

[6] Introducing the New Ray-Ban | Meta Smart Glasses (2023): https://about.fb.com/news/2023/09/new-ray-ban-meta-smart-glasses/

policymakers, experts in academia and civil society, and others in our industry to advance the beneficial and responsible use of this technology.

At Meta we're already building safeguards into our generative AI models and products from the beginning and working with others to collaborate on establishing guardrails, including:

1. Training and fine-tuning models to fit our safety and responsibility guidelines;
2. Red-teaming with external experts and internal teams to help ensure our models are safer and more inclusive;
3. Building in guardrails to limit inappropriate or harmful outputs;
4. And sharing system cards publicly so people better understand how these models work.

Our responsible AI approach is rooted in a set of principles that guide the development of our AI models:

1. **Transparency**. Developers should be transparent about how their systems work. At Meta, we have recently released 22 "system cards" for Facebook and Instagram, which give people insight into the AI behind how content is ranked and recommended in a way that does not require deep technical knowledge.
2. **Collaboration**. Openness requires collaboration across industry, government, academia and civil society. Meta is a founding member of Partnership on AI[7], alongside Amazon, Google, DeepMind, Microsoft, and IBM. We are participating in its Framework for Collective Action on Synthetic Media[8], an important step in ensuring guardrails are established around AI-generated content.
3. **Stress Testing**. AI systems should be stress tested. "Red teaming", common in cyber security circles, involves teams taking on the role of adversaries to hunt for flaws and unintended consequences.
4. **Openness**. Companies should share details of their work as it develops, be it through academic papers and public announcements, open discussion of the benefits and risks, or, if appropriate, making the technology itself available for research and product development.

[7] Partnership on AI: https://partnershiponai.org/partners/
[8] PAI's Responsible Practices for Synthetic Media: https://syntheticmedia.partnershiponai.org/#:~:text=A%20Framework%20for,generated%20or%20modified%20by%20AI.

Meta

For several years now, Meta has had interdisciplinary teams working on a variety of topics related to responsible AI[9], including privacy, governance, fairness, safety, and transparency, with the mission of ensuring that AI at Meta benefits people and society. Exploratory research, open science and cross-collaboration are foundational to Meta's AI efforts. Many of our AI applications are underpinned by breakthroughs made by our Fundamental AI Research (FAIR) team[10], which was established over a decade ago. Our research teams work closely with product teams across Meta to accelerate our research-to-production pipeline, enabling technology from scientific breakthroughs to go from the lab into product sometimes in a matter of months.

Most of our projects are done in collaboration with external researchers. We do it because we have a strong hypothesis that this is the fastest way to make progress in research:

1. We have more than 1,000 Meta open source projects.
2. FAIR has labs all over the world (US, Canada, France, UK, Israel).
3. Meta developed PyTorch, which is one of the fastest-growing open source deep learning frameworks.
4. More than 80,000 individuals and organizations have applied for access to LLaMA and we continue to receive new requests every day.
5. To date, we have already seen over 39,000 unique individuals who cloned (copied locally) the Segment Anything model.
6. We deliver AI breakthroughs for Meta, science, and society.
7. We've previously seen a number of our research projects positively impact our products such as NLLB (incorporating modeling techniques into our family of apps) and Few Shot Learner for Integrity.

We've seen first-hand how making AI models available to researchers can reap enormous benefits. For example, AI has been used to speed up the discovery of antibiotics and to generate equally accurate and detailed MRIs using about a quarter of the raw data. 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

[9] Facebook's five pillars of Responsible AI (2021): https://ai.meta.com/blog/facebooks-five-pillars-of-responsible-ai/
[10] Meta Research. Innovating with the freedom to explore, discover and apply AI at scale: https://ai.meta.com/research/

amazing things they've achieved by building on top of it. We believe our open approach to Llama 2 and the thousands of businesses, researchers, academics and others building with it will help us iterate faster on these new generative AI experiences to enhance performance, safety, and creativity – and ultimately inform how we build experiences that are designed for many vs. a select few.

Furthermore, we continue to publish and open source our work not only to validate our results, but also to enable others to build off of our advances, accelerating the world's progress toward more capable AI.

1. We have an ambitious research program, aimed at achieving Autonomous Machine Intelligence[11]. We regularly share the results of our work through peer-reviewed papers and open science.
2. We have more than 1,000 Meta open source projects that are publically available through our github repository[12].
3. Our broader approach to open science and industry collaboration allows the world and our partners to hold us accountable to innovate responsibly.
4. And our commitment to open sourcing doesn't slow down our ability to implement the research into products. In fact, our research teams work closely with product teams across Meta.
5. This has helped us build a fast research to production pipeline, where scientific breakthroughs can go from the research lab to products reaching billions of people, sometimes in a matter of months.
6. Meta is on the bleeding edge, with world-class research and technical leadership matched with deep investment and global scale. This means our teams are uniquely positioned to work on some of the most challenging problems out there—and to apply that work to building real products that solve real problems.

We are working together with various stakeholders, including governments, industries, AI experts in academia and civil society, parents, mental health experts, privacy experts and advocates, to establish responsible guardrails for AI. We have also joined the Synthetic Media Framework of the Partnership on AI (PAI) and their draft Model Deployment Guidance, which is now open for public comment. Additionally, we have joined the White House Voluntary AI Commitments for AI and submitted Llama 2 to DEFCON for nearly

[11] LeCun, Y. (2022). A Path Towards Autonomous Machine Intelligence: https://openreview.net/pdf?id=BZ5a1r-kVsf
[12] Meta Research github repository: https://github.com/facebookresearch

2,500 hackers to stress test the model. Our bug bounty program incentivizes researchers to stress test and report vulnerabilities.

We are also involved in several initiatives to promote safe and responsible AI. For example, we established the AI Alliance, a community of over 50 technology creators, developers, and adopters collaborating to advance safe and responsible AI rooted in open innovation. We are also a founding member and funding support of the Partnership on AI (PAI), a non-profit community of academic, civil society, industry, and media organizations creating solutions so that AI advances positive outcomes for people and society. Furthermore, we are a founding member of ML Commons, an AI engineering consortium built on a philosophy of open collaboration to improve AI systems.

IV. Open Innovation

Meta and many other AI leaders believe that AI should be built to benefit everyone across society. Making AI models available openly can deliver benefits across a wider range of stakeholders and unlock opportunities for everyone. As Meta's President for Global Affairs Sir Nick Clegg has said, *"Openness isn't altruism — Meta believes it's in its interest. It leads to better products, faster innovation, and a flourishing market, which benefits us as it does many others."*[13]

An open approach increases transparency and trust in the technology, facilitates research, and ensures participation from a wide range of creators and developers, which leads to the development of better and safer products and increases competition. Developers will be able to draw upon the collective wisdom, diversity, and ingenuity of AI practitioners across academia and businesses alike to harness the potential of foundation models to help ensure the technology built is inclusive and safe.

An open approach also increases competition by providing an alternative to proprietary or closed models, and democratizing access to AI technology. Broad availability of pre-trained and fine-tuned generative AI models will allow even smaller firms to offer AI-powered products and services in a range of industries, competing head-to-head with leading tech companies and other large firms, by innovating on top of open models. In turn, this innovation will spur new market competition, increasing output, lowering prices, creating more choices, and overall benefiting consumers in a whole range of ways.

[13] Sir Nick Clegg: Openness on AI is the way forward for tech (2023). Appeared in the FT: https://www.ft.com/content/ac3b585a-ce50-43d1-b71d-14dfe6dce999

Meta

One of the biggest challenges of working with AI is the level of resources required (e.g., research, data, computing power). Meta and other developers of open models have released pre-trained and fine-tuned models in a range of sizes and capabilities, which allows firms to pick the option that works for them without needing the expertise, data, or computing infrastructure to build and train a model from scratch. In this way, open innovation and open models provide a wide range of stakeholders who may not have the capital or resources to make these investments independently (developers, academics, civil society, etc.) with access to AI models that will allow them to compete.

Open Innovation will create opportunities for a broad range of firms to build whatever they can imagine, including smaller enterprises that otherwise would not have access to these tools. Open models also allow firms to customize them as they see fit, including by using their own data to fine-tune models for various use cases (e.g., a firm using customer data to fine-tune a customer service chatbot, or a pharma firm using research data to fine-tune a model for drug discovery). Open Innovation will contribute to the development of a larger number of AI use cases representing diverse interests, diverse backgrounds, and diverse goals ultimately democratizing the AI industry. Applications built on Llama 2 include:

- Lyrise[14]: built a chatbot on Llama 2 that helps match job candidates from Africa with AI experience with tech firms looking for talent.
- Niantic Labs[15]: maker of Pokemon Go created "Peridot," an augmented reality game using Llama 2 to generate character reactions.
- Cornell University: released a paper on Radiology-Llama 2[16], a model trained on radiology reports to support clinical decision-making.
- Yale University: created Meditron[17], an open source LLM built using Llama 2 and trained on medical data sources, to help guide clinical decision-making.

Open Innovation will help ensure the competitive process remains dynamic by creating opportunities for a wide range of market and societal stakeholders. As such, open innovation will help close the inequality gap by more widely distributing economic

[14] Lyrise website: https://lyrise.ai/
[15] Niantic Labs website: https://nianticlabs.com/?hl=en
[16] Radiology-Llama2: Best-in-Class Large Language Model for Radiology: https://arxiv.org/abs/2309.06419
[17] MEDITRON-70B: Scaling Medical Pretraining for Large Language Models: https://arxiv.org/abs/2311.16079

opportunity; in turn, economic growth will be more sustainable with so many market participants pursuing many different ideas and approaches.

We've released more than 1,000 open source AI projects so far, which levels the playing field and fosters innovation for people and businesses while increasing safety. Examples include:

1. OPT-175B: With the release of OPT-175B in May 2022, we opened up direct access to the large-scale models to the research community for the first time, so that scientific discourse on LLMs can be conducted on reproducible results. We called on the broader research community to help develop rigorous tooling to detect and mitigate potential safety issues in these models so that proper extraction techniques can be developed (and mitigation efforts can be upstreamed).
2. LLaMA: We released LLaMA (Large Language Model Meta AI)[18], a state-of-the-art foundational large language model, to the AI research community in February 2023. The goal of the release was to accelerate responsible research in this important, fast-changing field. Smaller models such as LLaMA enable others in the research community who don't have access to large amounts of infrastructure to study these models. LLaMA was released under a noncommercial license focused on research use cases. Access to the model has been granted on a case-by-case basis to those who have applied. To date, 80,000+ individuals and organizations have applied for access to LLaMA and we continue to receive new requests every day.
3. LLaMA 2: In response to commercial demand for open LLMs, we released several iterations of our successor to LLaMA, known as LLaMA 2. This set of models was released under a permissive license allowing commercial use, and has already been downloaded more than 100 million times and led to the creation of over 10,000 models derived from LLaMA 2.
4. Segment Anything: In April 2023 we released Segment Anything. The Segment Anything Model (SAM) can segment, or mask, almost any object from a photo or video with just a simple prompt — without needing additional training or examples to learn from. As part of this announcement, we released the model and the largest-ever segmentation dataset (SA-1B) containing 1 billion masks. To date, we have already seen over 39,000 unique individuals who cloned (copied locally) the model. Feedback has been extremely positive and we have already seen the

[18] Introducing LLaMA: A foundational, 65-billion-parameter large language model (2023): https://ai.facebook.com/blog/large-language-model-llama-meta-ai/

community experimenting with SAM across a wide range of use cases including medical imaging and geosciences.

5. PyTorch: In 2016, Meta AI (then Facebook) researchers set out in collaboration with the AI research community to build a better framework for AI research. This led to the creation and open sourcing of the PyTorch ML framework. Today, Pytorch has become one of the leading platforms for AI research as well as commercial production use with over 18k organizations using PyTorch. Here at Meta, PyTorch powers 50 on-device AI models across different mobile applications.

An open approach to innovation enables research in nearly limitless areas of AI. Allowing researchers access to state of the art models vastly increases the potential sources of innovation.

Meta's open approach to foundation models has already led to impressive technical innovations by the community:

1. Researchers from the University of Edinburgh, Google Research, and Macquarie University determined two factors[19] they propose are major sources of "hallucination" (providing incorrect or inappropriate information in a factual manner) among Llama and other LLMs, so companies can use these factors to further fine-tune their models.
2. Researchers from the University of Rome Tor Vergata, Idiap Research Institute, Campus Bio-Medico University, and Sapienza University developed a de-biasing method[20] that mitigates category bias in LLMs. Crucially, the method maintains performance on downstream tasks commonly used for benchmarking.
3. Researchers from the University of Illinois at Urbana-Champaign, Stanford University, the University of California, Berkeley, the Center for AI Safety, and Microsoft have performed a comprehensive study of the amount of trustworthiness of LLMs[21], by creating a new benchmark of adversarial texts using three open-source models fine-tuned from LLaMa.

[19] McKenna, N. et al. (2023). Sources of Hallucination by Large Language Models on Inference Tasks: https://arxiv.org/pdf/2305.14552.pdf
[20] Ranaldi, L. et al. (2023). A Trip Towards Fairness: Bias and De-Biasing in Large Language Models: https://arxiv.org/pdf/2305.13862.pdf
[21] Wang, B. et al. (2024). Decoding Trust: A Comprehensive Assessment of Trustworthiness in GPT Models: https://arxiv.org/pdf/2306.11698.pdf

4. Researchers from Harvard University developed a method to improve LLM truthfulness[22] during text generation, and tested it on a LLaMa model fine-tuned for instructions.

Open innovation benefits those that develop and use generative AI as it can also lead to safer products. A mistaken assumption is that releasing source code or model weights makes systems more vulnerable. On the contrary, an open community of external developers and researchers can more quickly identify and address problems than individual teams holed up inside company silos. Potential toxicity, bias, bugs, and vulnerabilities can be continuously identified and mitigated in a transparent way by an open community. We can learn from that work to build safer, more robust, and more effective products. The transparency of open models helps build trust in the technology.

V. Final Remarks

An open approach to AI allows a wide range of stakeholders – developers, academics, civil society, nonprofits, and more – to both realize the benefits of AI technologies and improve our understanding of how to manage and mitigate the potential risks. By allowing more people to use, research, inspect, and improve these models, an open approach helps ensure that this technology benefits the world at large.

Crucially, AI built on open innovation leads to more innovation, better and safer products, and more robust competition. Customers and developers have a wide range of models to choose from, open and proprietary, with various sizes, attributes, and capabilities. As models begin to be deployed in more applications, including on consumer devices, the Commission may want to monitor how that shift affects the choices available.

[22] Li, K. et al. (2023). Inference-Time Intervention: Eliciting Truthful Answers from a Language Model: https://arxiv.org/pdf/2306.03341.pdf