stability.ai

Response to the Call for Contributions on Competition in Generative AI
March 2024

Introduction

Stability AI welcomes the opportunity to respond to the European Commission ("Commission") call for contributions on "Competition in Virtual Worlds and Generative AI". As a leading AI developer with deep roots in European research, Stability AI is committed to supporting the growth of a competitive AI ecosystem in the European Union. The following comments focus on the role of open models in driving grassroots innovation across AI research, applications, and ventures. We encourage the Commission to ensure that future competition policy supports this grassroots innovation, and actively promotes a diverse AI ecosystem – from large firms building closed-source products to everyday developers and independent researchers sharing open technology.

Background

Stability AI is a global company working to amplify human intelligence by making foundational AI technology accessible to all. Today, we develop AI models across a range of modalities, including image, language, audio, and video. Essentially, these models are software programs that can help a user to create, edit, or analyze complex content. With appropriate safeguards, we release these models openly, sharing our software code along with the billions of distinctive settings or "parameters" that define the model's performance. That means everyday developers and

independent researchers can integrate or adapt our models to develop their own AI models, build their own AI tools, or start their own AI ventures, subject to our ethical use licenses.[1]

To date, our models have been downloaded over 100 million times by developers, and over 300,000 developers and creators actively contribute to the Stability AI online community.[2] Our family of image models, Stable Diffusion, underpin up to 80 percent of all AI-generated imagery.[3] These models can take a text instruction or "prompt" from a user and help to create a new image. In addition, we develop a suite of language models that can interpret, summarize, or generate text. These include highly capable large language models, compact language models, specialized models for software development, and models for underrepresented languages, including Spanish and Japanese. Our audio model, Stable Audio, generates high-quality soundtracks and was recently listed on the *TIME* Best Inventions of 2023. Building on this experience, we have developed models that demonstrate new breakthroughs in video generation and 3D rendering.[4] Further, we support academic research into scientific applications of AI. Stability AI provides a range of services to help partners customize and deploy our models, sustaining our open research efforts.

We are committed to the safe development of AI. To that end, we are signatories to the White House *Voluntary AI Commitments*, the British Government's *Joint Statement on Tackling Child Sexual Abuse in the Age of AI*, the Singapore Government's Generative AI Evaluation Sandbox, and the United States AI Safety Institute Consortium; we participated in the first large scale public evaluation of AI models at DEF CON, facilitated by the White House, and the UK AI Safety Summit; and we engage with authorities in the EU and around the world.

Do open-source generative AI systems and/or components, including AI models compete effectively with proprietary AI generative systems and/or components?

Open models play a vital role in promoting competition within the AI ecosystem. Generative AI will support critical applications across the digital economy, enabling creative, analytic, and scientific use-cases – from personalized tutoring to drug discovery – that go far beyond the caricature of "push a button, get an image" or "push a button, get a poem". Language models will power tools that revolutionize essential services, from education to healthcare; reshape how we search and access information online; and transform analysis, knowledge management, or decision making in some of our most important public and private sector institutions. Audiovisual models will power tools that radically accelerate the creative process, helping existing creators boost their productivity and experiment with new concepts while lowering barriers to entry for people who do not have the resources or training to realize their creative potential today. Instead of simply consuming the best available content, these "dormant" creators will be able to produce their best imaginable content.

[1] See e.g. the Open Responsible AI License (OpenRAIL) for Stable Diffusion, prohibiting a range of unlawful or misleading uses, available here. We use the term "open" to refer to any models with publicly-available parameters.
[2] Figures from Hugging Face and Discord, November 2023.
[3] Everypixel, 'AI Image Statistics', August 2023, available here.
[4] See e.g. Stability AI, 'Scaling Rectified Flow Transformers for High Resolution Image Synthesis', March 2024, available here; Stability AI, 'Stable LM 2 1.6B Technical Report', February 2024, available here; Stability AI, 'Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets', November 2023, available here.

As generative AI becomes embedded in tools and services, it is more important than ever that we can scrutinize these systems before the next wave of digital services are built on “black box” technology operated by a small cluster of firms. Today, the digital economy relies on opaque systems that amplify content on social media, govern our access to information, determine our exposure to advertising, and mediate our online interactions. It is difficult to scrutinize these systems, or build viable alternatives. Without a conscious effort to promote transparency and competition, AI is at risk of repeating this history. In that context, open models can help to improve safety through transparency, foster competition in critical technology, and support grassroots innovation in AI:

- **Open models promote transparency.** Researchers and authorities can “look under the hood” of an open model to verify performance, identify risks or vulnerabilities, study interpretability techniques, develop new mitigations, and correct for bias. By comparison, closed models may not disclose how they are developed or how they operate. Closed models may be comparatively opaque, and risk management may depend on trust in the developer.

- **Open models lower barriers to entry.** Training a new “base” model from scratch requires significant resources that are not available to everyday developers. Open models lower these barriers to entry. Everyday developers can build on open models to create new AI tools or launch new AI ventures without spending tens of millions of euros on research and computing. In this way, the economic benefits of AI accrue to a broad community of developers and firms, not just major firms.

- **Open models drive innovation in safety.** Developers can refine open models for improved safety and performance in specific tasks. For example, open models can be optimized through a range of techniques to mitigate undesirable behavior such as bias, misinformation, or toxicity. These techniques can yield significant improvements in the behavior of a model without requiring extensive computing resources. That means ordinary developers can build safer and more effective models to better support their real-world applications.

- **Open models foster strategic independence.** Open models enable public and private sector organizations to build independent AI capabilities without relying on a handful of firms for foundational technology. They can develop these AI capabilities securely “in house” without exposing their confidential data or ceding control of their distinctive model parameters to third parties. Operational independence will be important for organizations in sensitive or regulated sectors, such as healthcare, finance, law, and public administration.

- **Open models improve accessibility.** Many open models are smaller, more efficient, and more accessible than proprietary models. Unlike those models, which require significant computational resources to train and run, small open models can deliver useful performance with regular hardware. For example, open models may be hundreds of times smaller than a closed-source model such as GPT-4. Users can run small models on local devices, including smartphones, and developers can train or optimize these models with desktop hardware.

In this way, open models are fueling a wave of grassroots innovation in AI. Open models put this technology in the hands of everyday developers, independent researchers, and small businesses who are helping to build safer AI models and useful AI tools. Open models offer a transparent, competitive, and secure alternative to opaque systems owned and operated by a small number of firms.

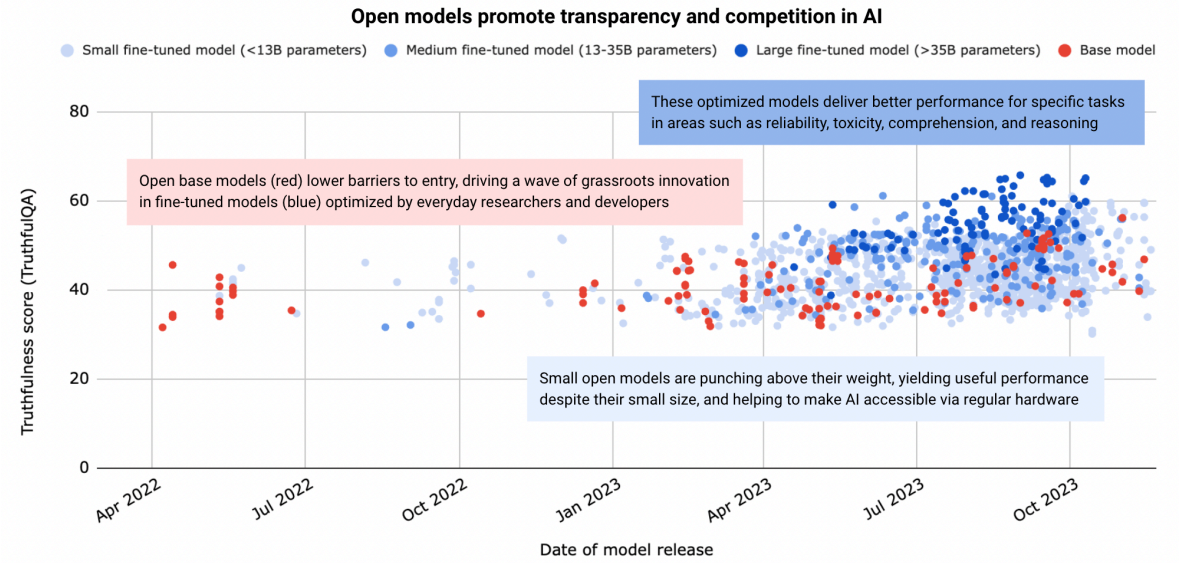

Source: Our analysis of data from the Hugging Face "Open LLM Leaderboard" (November 2023). The TruthfulQA benchmark measures a model's tendency to reproduce falsehoods.

Open technology supports vital digital infrastructure today, and underpins recent developments in AI. For example, Google openly published the research that underpins many large language models.[5] Meta, Google, and their partners chose to open-source foundational code libraries for machine learning.[6] University teams in Europe openly published the research that led to Stable Diffusion,[7] and a range of privately- and publicly-funded organizations have chosen to release highly capable base language models openly, such as Falcon 180B, Llama 2, Mixtral 8x7B, or our own Stable LM 2 1.6B. Beyond AI, open-source operating systems such as Linux underpin a significant portion of web servers and data centers globally, and can be found on submarines, destroyers, and SpaceX rockets. Similarly, Android is an open-source mobile operating system that powers a majority of all smartphones worldwide.[8] As with these other systemically important technologies, open models can help to promote transparency, competition, and security in the digital economy.

[5] Transformer via Vaswani et al, 'Attention is All You Need', 2017 available here.
[6] PyTorch via Paszke et al, 'PyTorch: An Imperative Style, High Performance, Deep Learning Library', 2016 available here; TensorFlow via Abadi et al, 'Large-Scale Machine Learning on Heterogeneous Distributed Systems', 2015 available here.
[7] Rombach et al, 'High-Resolution Image Synthesis with Latent Diffusion Models', 2021 available here.
[8] Vaughan-Nicols, 'From Earth to Orbit with Linux and SpaceX', ZDNET, 2020; Gallagher, 'The Navy's Newest Warship is Powered by Linux', *Ars Technica*, 2013; StatCounter, 'Mobile Operating System Market Share', 2023.

What are the main components (i.e. inputs) necessary to build, train, deploy, and distribute generative AI systems? Please explain the importance of these components.

Models are just one component in a generative AI system such as a chatbot or image generator. The model must be hosted and deployed in a user-facing system in order to analyze or generate content. In that environment, different actors perform different functions, ranging from: training a base model (the raw "engine" that understands complex patterns and relationships within a textual, visual, musical, or scientific dataset); optimizing or "fine-tuning" the model for a specific use-case (such as conversational interactions); distributing the model; hosting the model on a computing service; developing a user-facing application that interacts with the model; and promoting that application to users. Each actor may have limited visibility or control over downstream activity.

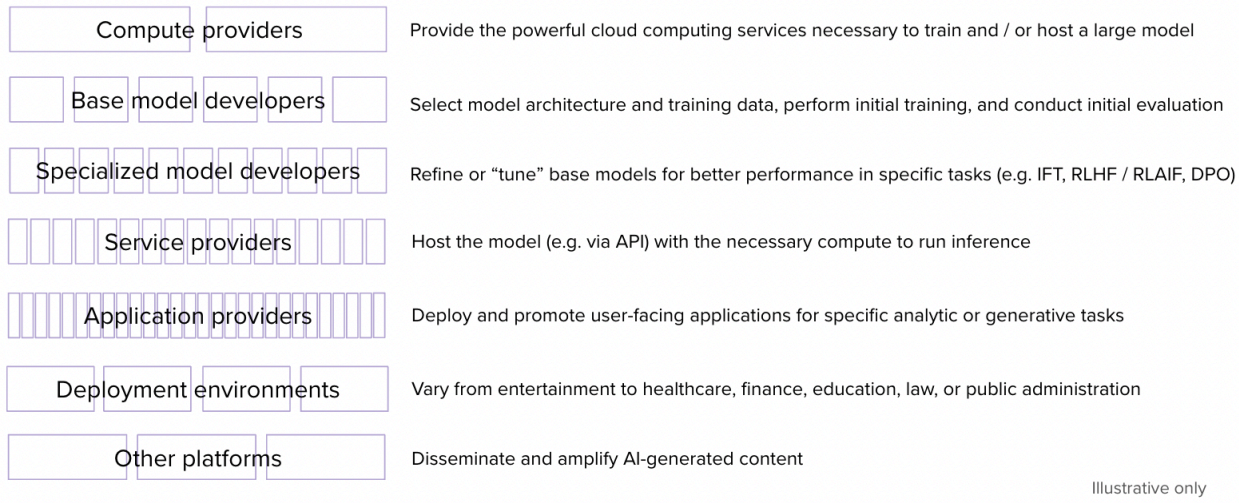

Competition frameworks should account for diversity in these supply chains. The relationships between actors in a vertically-integrated, closed-source environment may be different to the relationships in a disaggregated, open-source environment. Conflating these actors, or imposing rules that assume vertical integration or formal relationships between these actors, could have a chilling effect on open innovation in AI. Instead, Stability AI encourages authorities to acknowledge the range of actors and components that comprise an AI system. While there is no single gatekeeper and no silver bullet, every actor can play a role in helping to mitigate the risks of unintentional harm or intentional misuse.

In particular, the risk profile of an AI system will vary depending on how and where the system is deployed. For example, an AI system deployed in higher-stakes domains such as healthcare, finance, education, or public administration may attract more rigorous obligations than an AI system deployed in a lower-stakes domain such as entertainment, with different requirements for reliability, interpretability, and robustness. One model may be deployed in a range of such environments, and responsibility for risk mitigation and assurance may be shared by different actors in different ways. "One size fits all" frameworks governing development could set back competition by imposing disproportionate or ill-adapted requirements on every AI system and every AI component without accounting for their specific risks.

Stability AI has urged all European Union institutions to ensure that the AI Act ("Act") accounts for this variety in AI supply chains and use-cases. Earlier versions of the Act prescribed a single set of obligations for all models, with no allowance for different capabilities, different operating environments, or different risks. In addition, the Act had the unintended effect of treating individuals who shared fine-tuned variants of models the same as the organizations – often corporations – that released the initial base model. These provisions would have had a significant chilling effect on open innovation in Europe, making it impossible for everyday developers and independent researchers to access, refine, and share models openly. To that end, Stability AI is encouraged by the final text of the AI Act, which establishes a partial exemption for free and open-source models, and recognizes important differences between base model developers, fine-tune model developers, and systems deployers.[9]

What are the main barriers to entry and expansion for the provision, distribution, or integration of generative AI systems and/or components, including AI models? Please indicate to which components they relate.

Access to models is a major barrier to the development and deployment of AI systems. Training a base model from scratch may be prohibitive for many deployers, especially small businesses. They may not have the requisite expertise, access to compute resources, or the data management infrastructure to support training. By way of illustration, OpenAI disclosed that it cost USD 100 million to train the closed-source GPT-4 model.[10] Training a 70 billion parameter language model today via third-party compute services might cost over 3 million euros.[11] Open models can help developers and deployers to access foundational components without "reinventing the wheel" at significant expense. To that end, we welcome efforts by the Commission to fund additional research into open models, open data, and open compute under the Digital Europe Programme.

For example, open models help to support the development of safe and effective AI systems. For example, a raw or "pre-trained" base model might understand how to read, write, or draw, but it may be prone to undesirable behaviors such as bias, misinformation, or toxicity. It must be fine-tuned for a specific task. Given access to the weights of an open model, developers can adjust these behaviors before deployment in a user-facing system, taking into account their intended application and specific operating environment. In addition, open access to model parameters helps to support risk assurance by enabling deployers, researchers, and authorities to directly scrutinize the behavior of a model.

Likewise, open models to support security in AI applications. By building on open models, developers can operate their own AI capabilities without relying on third-party service providers. They can fine-tune their model without sharing confidential datasets with the upstream model developer, and they can run the model without exposing user data to a third-party host. For example, a regulated financial institution may customize AI models to assist in analysis, decision making, or customer support. The financial institution must be able to audit the performance of the model for reliability; improve the model without exposing sensitive customer data; and retain

[9] Provisional Agreement, AI Act, art 52c and recital 60g.
[10] Wired, 'Open AI's CEO says the age of giant models is already over', April 2023, available here.
[11] Hugging Face, Training Cluster, available here.

full control over their resultant AI capabilities. By building on open models, a financial institution can customize, operate, and manage their own AI technology stack.

Conclusion

Open models play a vital role in the emerging AI ecosystem, and they form part of the EU's competitive advantage in AI. We have welcomed engagement with all EU institutions on these matters to date, and urge the EU to ensure that future policy accounts for grassroots innovation in open models.