SLTA

2025-2

BEYOND THE HYPE :

WHAT "OPEN-SOURCE AI" REALLY MEANS

(AND WHY SWITZERLAND'S APERTUS MATTERS)

BRECHTJE LINDEBOOM, CRYSTAL DUBOIS

Editorial Team :

Nicolas Torrent, Nomungerel Jamsranjav, Ziyu Liu

INTRODUCTION

In early September 2025, the release of the Swiss Apertus open-source AI large language model ("LLM") made headlines[1]. Apertus is a joint project of the Swiss National AI Institute, a partnership between the two Swiss Federal Institutes of Technology, ETH Zurich and EPFL, and the Swiss National Supercomputing Centre ("CSCS")[2]. Developed on the "Alps" supercomputer at CSCS[3], it represents the first fully open multilingual and transparent large language model.

As organizations increasingly adopt open-source AI technologies, understanding both the concept and its legal implications becomes essential.

This article clarifies what open-source AI entails and it examines the regulatory framework governing its development and implementation, particularly under the EU AI Act ("AI Act").

Apertus – "open" in Latin – is a fully open suite of LLMs, which are openly released on Hugging Face[4], the platform for open-source AI. It addresses two systemic shortcomings in today's open AI model ecosystem:

- data compliance, and;
- lack of multilingual representation.

What sets Apertus apart is complete transparency: its entire development process, architecture, model weights, training data, and recipes are openly accessible and fully documented[5].

Secondly, its multilingual representation aims to democratize LLMs for broader global communities.

Even though Apertus is Swiss, it falls under the AI Act's jurisdiction if made available for distribution or use within the European Union market, in the context of a commercial activity.

The release of “open-source” AI models is not unprecedented. Earlier this year, the Chinese firm DeepSeek captured attention with the release of its open-source LLM[6]. Major players have also been active in the field, including Google with Gemma LLM[7], Meta with Llama models[8], Mistral AI[9], and even OpenAI with gpt-oss-120b[10] and gpt-oss-20b[11] in August 2025. In fact, at the date of writing this article, Hugging Face hosts over 2 million “open-source” AI-models and more than 500’000 “open-source” datasets[12].

Organizations increasingly rely on open-source AI for building and deploying AI-driven solutions.

A recent report[13] indicates wide adoption of open-source AI: over 50% of surveyed technology leaders and senior developers across 41 countries use open-source AI across data, models, and tools areas of the tech stack. Survey results indicate that leaders embrace open-source AI tools for their high performance, ease of use, and lower implementation and maintenance costs relative to proprietary alternatives.

DEFINITION OF OPEN-SOURCE AI

a. *Conflicting views*

No widely accepted definition of open-source AI currently exists; the open-source community and technology companies debate its meaning.

Open-source AI generally refers to AI models that incorporate components released under open-source licenses. These components range from training data (rarely made open) to documentation, software code, and model weights. The open-source community recognizes various license types, from permissive licenses like MIT/Apache to more restrictive ones such as GNU GPL or BSD licenses[14].

b. *The Open Source AI Definition (OSAID)*

In late 2024, the open-source community published the “Open Source AI Definition” (“OSAID”)[15]. It establishes that open-source AI systems must grant users four fundamental freedoms: (1) to use, (2) study, (3) modify and (4) share the system freely for any purpose[16]. To meet this standard, AI providers must supply sufficiently detailed information about training data, including its provenance and sources. They must also provide the complete source code for training and using the system, detailing data processing and training methodologies. Finally, all model parameters, including weights and configuration settings, must be made available under approved terms.

This definition conflicts with current industry practices. Corporations like Meta, Google, and Mistral have released AI model components under proprietary licenses labeled as "open source". However, these licenses often contain restrictive terms that prevent users from freely using, modifying, and sharing the models, making them non-compliant with the OSAID supported by the open-source community[17].

c. *The AI Act’s approach*

The AI Act closely reflects the OSAID approach: Recital 102 of the Act states that “*Software and data, including models, released under a free and open-source licence that allows them to be openly shared and where users can freely access, use, modify and redistribute them or modified versions thereof, can contribute to research and innovation in the market and can provide significant growth opportunities for the Union economy.*

General-purpose AI models released under free and open-source licences should be considered to ensure high levels of transparency and openness if their parameters, including the weights, the information on the model architecture, and the information on model usage are made publicly available. The licence should be considered to be free and open-source also when it allows users to run, copy, distribute, study, change and improve software and data, including models under the condition that the original provider of the model is credited, the identical or comparable terms of distribution are respected."[18]

d. *The really of varying openness*

The variety of components that may be opened and the different governing licenses create a broad spectrum of what constitutes "open-source AI" today. This diversity reflects different approaches to balancing openness with business considerations and technical constraints. Technology companies may release certain components while maintaining restrictions on others, resulting in varying degrees of openness rather than a binary open-or-closed model. Therefore, when evaluating AI models that claim to be open, it is important to understand the specific level and scope of openness, as this can range from limited component access to full transparency across all elements of the system.

e. *Apertus: an open-source AI model?*

Apertus (a 'General-Purpose AI Model' ("GPAIM") under the AI Act) exemplifies comprehensive open-source AI implementation, and is one of the first AI models to achieve full compliance with both OSAID and the AI Act requirements.

Indeed, Apertus was released under the Apache 2.0 license, providing complete transparency through the inclusion of model weights, source code, final and intermediate checkpoints, reproducibility scripts for training data, evaluation suites, and a detailed technical report[19]. Furthermore, the full data sets used for training are openly available, under various license terms.

In summary, as stated by project leaders in a recent podcast episode[20], Apertus establishes a new standard for accessible and compliant LLMs by offering the community fully reproducible models and components, that are compliant with the law, and that support diverse applications and use cases.

REGULATORY FRAMEWORK UNDER THE AI ACT FOR PROVIDERS OF OPEN-SOURCE AI

The AI Act[21] distinguishes between open-source AI systems which are exempt from the AI Act, and GPAIM released under an open-source license which are exempt from specific obligations of the AI Act only.

AI systems released under free and open-source licenses are exempt from the AI Act, unless they are (1) placed on the market in the course of a commercial activity or put into service as high-risk AI systems, or; (2) if they are classified as an AI system that falls (a) under prohibited AI practices (Art. 5) or (b) under transparency obligations (for AI systems that can generate or manipulate realistic image, audio, or video content (i.e., “deepfakes”) (Art. 50)[22].

This exemption does not apply to AI systems or AI components provided against payment or otherwise monetized, such as through the provision of technical support or other services, or the use of personal data for reasons other than exclusively for improving the security, compatibility or interoperability of the software, with the exception of transactions between microenterprises[23].

Third parties making services, processes, or AI components other than GPAIM accessible to the public tools under a free and open-source license have reduced compliance obligations, particularly in the context of high-risk AIs. Specifically, they are not required to comply with requirements pertaining to the AI value chain. This includes the responsibilities toward a provider who used or integrated them[24].

Finally, providers of GPAIM released under a free and open-source license benefit from exceptions concerning certain documentation and transparency obligations, subject to some conditions. The conditions and exemptions have been further developed in the European Commission’s Guidelines on the scope of obligations for providers of GPAI models, which were published on July 18, 2025[25].

Importantly, for exemptions to apply, the GPAIM parameters, including the weights, information on the model architecture, and information on model usage, must be made publicly available. We however note that the GPAIM does not apply if classified as presenting systemic risks[26] or if an AI component within the open-source model is monetised[27].

The main exemptions for compliant Open-source AI models are significant. First, they do not need to create and maintain technical documentation of the model, including training and testing processes and results.[28] Second, they are not required to provide information and documentation to downstream providers integrating the model[29]. Finally, providers established in third countries who are subject to the AI Act's extraterritorial scope of application[30] are released from the obligation to appoint an authorized representative established in the European Union[31].

Notwithstanding the open-source status, providers must still comply with obligations to put in place a policy to meet relevant European Union copyright law (specifically concerning the reservation of rights) and must produce and make publicly available a sufficiently detailed summary of the content used for training the model[32].

BUILDING ON TOP OF OPEN-SOURCE

When organizations use open-source AI as a building block, they must primarily comply with the underlying open-source license terms. These obligations typically center on attribution requirements, mandating proper credit to original authors and contributors, often satisfied by preserving original copyright notices and providing appropriate acknowledgment in end-user documentation.

Downstream users face potential contractual liability for non-compliance with license terms and tort liability if deployed systems cause harm or infringe rights[33].

Additionally, under the AI Act, deployers may face specific obligations depending on their AI system's risk classification, including transparency requirements, human oversight duties, and documentation obligations.

Notably, substantial modifications to open-source AI models may transform the downstream user into a new "provider" under the AI Act, triggering additional compliance responsibilities[34]. Indeed, according to the European Commission's Guidelines on the scope of obligations for providers of GPAI models, an indicative criterion for when a downstream modifier is considered to be the provider of a GPAI model is that the training compute for the modification is greater than a third of the training compute of the original model[35].

CONCLUSION

The open-source AI landscape remains complex, with ongoing debates about what truly constitutes "open" in this context. However, the alignment between the OSAID and the AI Act's framework provides much-needed clarity for organizations navigating this space.

Switzerland's Apertus demonstrates that comprehensive open-source AI compliance is achievable, setting a benchmark for transparency across all model components.

As the AI Act takes full effect, organizations must carefully evaluate their open-source AI usage, particularly given the potential to become "providers" under the regulation through substantial model modifications.

The key lesson for practitioners is straightforward: not all "open-source" AI is created equal. Organizations need robust processes for license compliance and must understand their regulatory obligations based on how they use and modify these models.

As enforcement mechanisms develop, the distinction between genuinely open and nominally open AI systems will become increasingly important for both compliance and competitive positioning.

REFERENCES

1. https://actu.epfl.ch/news/a-language-model-built-for-the-public-good/; https://www.swiss-ai.org/apertus.
2. https://actu.epfl.ch/news/a-language-model-built-for-the-public-good/; https://www.swiss-ai.org/apertus.
3. https://www.swiss-ai.org/.
4. https://huggingface.co/collections/swiss-ai/apertus-llm-68b699e65415c231ace3b059.
5. “Apertus: Democratizing open and compliant LLMs for global language environments”, Apertus v1 technical report, available at: https://github.com/swiss-ai/apertus-tech-report/blob/main/Apertus_Tech_Report.pdf.
6. https://www.technologyreview.com/2025/01/31/1110740/how-deepseek-ripped-up-the-ai-playbook-and-why-everyones-going-to-follow-it/; https://huggingface.co/deepseek-ai/DeepSeek-V3.1; https://huggingface.co/deepseek-ai/DeepSeek-R1.
7. https://ai.google.dev/gemma/docs/core; https://huggingface.co/google/collections.
8. https://huggingface.co/meta-llama; https://www.llama.com/.
9. https://huggingface.co/mistralai/models;
10. https://huggingface.co/openai/gpt-oss-120b.
11. https://huggingface.co/openai/gpt-oss-20b.
12. https://huggingface.co/models; https://huggingface.co/datasets.
13. “Open source technology in the age of AI” by McKinsey & Company, Patrick J McGovern Foundation and Mozilla Foundation, 22 April 2025 (https://www.mckinsey.com/capabilities/quantumblack/our-insights/open-source-technology-in-the-age-of-ai.
14. For all the licenses that are recognized as complying with the Open Source Definition, see the list at https://opensource.org/licenses.
15. https://opensource.org/ai/open-source-ai-definition.
16. Open Source Artificial Intelligence definition version 1.0 published by the Open Source Initiative (https://opensource.org/ai/open-source-ai-definition).
17. For a brief review of the comparison between those different types of licenses, see the license terms themselves and the following sources: Yaniv Benhamou in Kluwer Copyright Blog, “Open Source AI – definition and selected legal challenges” dated April 15 2024 (https://legalblogs.wolterskluwer.com/copyright-blog/open-source-ai-definition-and-selected-legal-challenges/) and “Open Source Artificial Intelligence Definition 1.0 – a “take it or leave it approach” dated March 4 2025

(https://legalblogs.wolterskluwer.com/copyright-blog/open-source-artificial-intelligence-definition-10-a-take-it-or-leave-it-approach-for-open-source-ai-systems/); McKinsey & Company and. AI mentioned earlier (https://www.mckinsey.com/capabilities/quantumblack/our-insights/open-source-technology-in-the-age-of-ai); Open Future Foundation and Open Source Initiative, "Data Governance in Open Source AI" dated January 22, 2025 (https://opensource.org/data-governance-open-source-ai).

18. Underlining added by the authors.
19. Apertus – Democratizing open and compliant LLMs for global language environments, Apertus v1 technical report, available at: https://github.com/swiss-ai/apertus-tech-report/blob/main/Apertus_Tech_Report.pdf.
20. EPFL AI Center : Inside AI, "Conversation about Apertus with Antoine Bosselut, Marting Jaggi and Imanol Schalg (EPFL and ETH Zurich)", September 9th 2025, (available on Spotify at: https://open.spotify.com/episode/5qDXRWQwUU3ytPkVWsJIax?si=mb77-_DnSJWOCEM4jdrAwg).
21. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ%3AL_202401689.
22. Article 2 (12) AI Act.
23. AI Act : Recital 103.
24. Article 25(4) AI Act.
25. European Commission Guidelines on the scope of obligations for providers of general-purpose AI models under the AI Act, July 18 2025, Sections 4.1 and 4.2.
26. As defined in Article 3 (65) AI Act.
27. See European Commission Guidelines on the scope of obligations for providers of general-purpose AI models under the AI Act, July 18 2025, Section 4.2.2; AI Act: Recital 103.
28. Article 53(2) AI Act.
29. Idem.
30. In accordance with Article 2 para. 1 (a) AI Act.
31. Article 54(6) AI Act.
32. AI Act: Recitals 102-104.
33. Yaniv Benhamou in Kluwer Copyright Blog, "Open Source AI – definition and selected legal challenges" dated April 15 2024 (https://legalblogs.wolterskluwer.com/copyright-blog/open-source-ai-definition-and-selected-legal-challenges/).
34. AI Act, Recital (109).
35. European Commission Guidelines on the scope of obligations for providers of general-purpose AI models under the AI Act, July 18 2025, Section 3.2 (63).