

One year on, has Swiss AI model Apertus lived up to the hype?

02.09.2026 Sara Ibrahim

One year after its launch, Switzerland's fully open AI model has passed four million downloads and found real-world uses. Now on version 1.5, Apertus is closing the gap with its rivals, but weaknesses on key tasks still make it difficult for the Swiss model to compete in the fast-moving AI race.

The news arrived with much fanfare: Switzerland finally had its own fully open, transparent and multilingual artificial intelligence (AI) model, designed with data protection rules and the European AI Act in mind. Yet just weeks after its launch, Apertus was attracting harsh criticism. Some users [External link](#) went as far as calling it “unusable” and “unable to correctly answer even the most basic questions”.

One year on, the numbers look encouraging. Apertus has surpassed four million downloads on the AI platform Hugging Face – a figure Apertus community lead Oleg Lavrovsky considers representative of global interest in open-source AI projects.

Download figures, however, do not necessarily correspond to real-world usage. Lavrovsky says he is aware of at least 70 external deployments and more than one hundred derivative versions of the Swiss model. That includes Singapore's SEA-LION project [External link](#), which independently adapted Apertus and other AI models to work across multiple Southeast Asian languages.

For Lavrovsky, this is important validation. “We wanted to create a foundation model,” he says. “When people optimise or adapt Apertus, that's exactly what's happening.”

In July, the team released Apertus 1.5, which introduced image understanding and improved reasoning capabilities, among other features. The official roadmap [External link](#) foresees the release of Apertus 2.0 in 2027.

For those who remain disappointed, Lavrovsky has a clear message: Apertus was not designed to rival ChatGPT. It is not a finished commercial product, but a foundation that companies, institutions and researchers can adapt to their own needs. He acknowledges, though, that questions remain about to what extent the Swiss model's potential is translating into actual use and tangible value.

Open-source, open-weight and closed AI models

AI models can be developed and released with different levels of openness:

Open source: The model's trained parameters, or “weights”, training code and information about the training data are publicly available, allowing others to inspect, modify and reproduce it. Examples include Apertus and OLMo.

Open weight: The model's weights are publicly released and can be downloaded and adapted, but the full training code and datasets are not necessarily disclosed. Examples include Meta's Llama and Alibaba's Qwen.

Closed/proprietary: The model's weights, training code and datasets are not publicly released. This is the case for most commercial models, including OpenAI's GPT and Anthropic's Claude.

Multilingualism and data safeguards have paid off

Some firms have found uses for Apertus. Ticino-based Artificialy, which specialises in AI solutions, says it works extensively with Apertus.

“We're not doing charity. We don't use Apertus because we have to, but because it's a good model,” says Damiano Binaghi, the company's head of deep learning.

According to Artificialy, Apertus is already mature enough for some industrial and public-sector applications. The company's internal tests show particularly good results in translation and multilingual conversations.

“We often prefer Apertus to other models for our clients' translation tasks, especially when they involve Switzerland's four national languages or sensitive documents,” Binaghi says.

One example is the southern Swiss canton of Ticino, which has decided to migrate to Apertus for the automated translation of administrative documents. The cantonal migration office benefits the most, as it deals with personal documents in multiple languages and requires both accuracy and strong data protection guarantees.

“The needs of Ticino’s migration office are almost the perfect embodiment of the model’s strengths,” Binaghi says.

Multilingualism is one of the areas in which Apertus’s developers invested most heavily. Around 40% of its training data is in languages other than English, and the corpus includes more than 1,800 languages, including under-represented ones such as Romansh and dialects of Swiss-German. By comparison, estimatesExternal link suggest that current AI models use only around 10% of non-English data.

>> Does Apertus really speak more than 1,800 languages? We’ve analysed the facts and myths surrounding the Swiss AI model:

Full openness was one of the reasons AI Singapore chose Apertus as the foundation for a version of its SEA-LION project. “We chose Apertus because it is one of the few truly open-source models in the world and that’s important for research, especially when you need to know exactly what data was used and how it was trained,” Leslie Teo of AI Singapore told Swissinfo by email.

For the same reasons, Apertus has also found applications in medicine. In June, the Federal Institute of Technology Lausanne (EPFL) presented MeditronFOExternal link, a fully open framework for building clinical AI models based also on Apertus.

The performance gap remains

While Artificially considers Apertus 1.5 to have improved significantly in translation, vision and reasoning, the company sees room for improvement.

“When it comes to agentic capabilities, despite considerable progress, Apertus is not yet at the level of other open-weight models,” Binaghi says. Agentic AI models are able to independently execute actions involving several steps, such as comparing and booking flights.

The gap is also visible in benchmarks published by the Apertus team itself. The original and largest model, with 70 billion parameters, remains well behindExternal link similarly sized open-weight models, particularly in programming, mathematics and reasoning tasks.

Liip, a Zurich-based digital agency that tested Apertus in ZüriCityGPT, its AI-search tool for answering citizens’ questions, also sees Apertus lagging behind commercial models, although the release of Apertus 1.5 has narrowed the gap.

Josef Kruckenberg, head of generative AI projects at Liip, says Apertus 1.5 produces more useful answers, follows instructions better and is more likely to refuse to answer when it lacks sufficient information. But it lags behind in what Liip considers its most important criterion: faithfulness, or the extent to which every part of an answer is supported by the available sources.

“So far, we haven’t used Apertus in client projects. We felt it wasn’t ready yet,” Kruckenberg says. With Apertus 1.5, it is still too early to determine concrete ways of using it with clients, he adds.

>> AI models have biases? Switzerland’s Apertus is no exception:

The price of openness

Developing competitive fully open models remains a difficult undertaking and involves inevitable trade-offs. “Being truly open-source means being unlikely to use some data, and this has consequences on performance and general capabilities,” Leslie Teo of AI Singapore says.

Artificially’s Binaghi agrees: Apertus relies on transparent and verifiable public data in line with copyright requirements, unlike proprietary developers that often use licensed or undisclosed datasets. This can lead to a slight quality gap, but it also strengthens digital sovereignty, he points out.

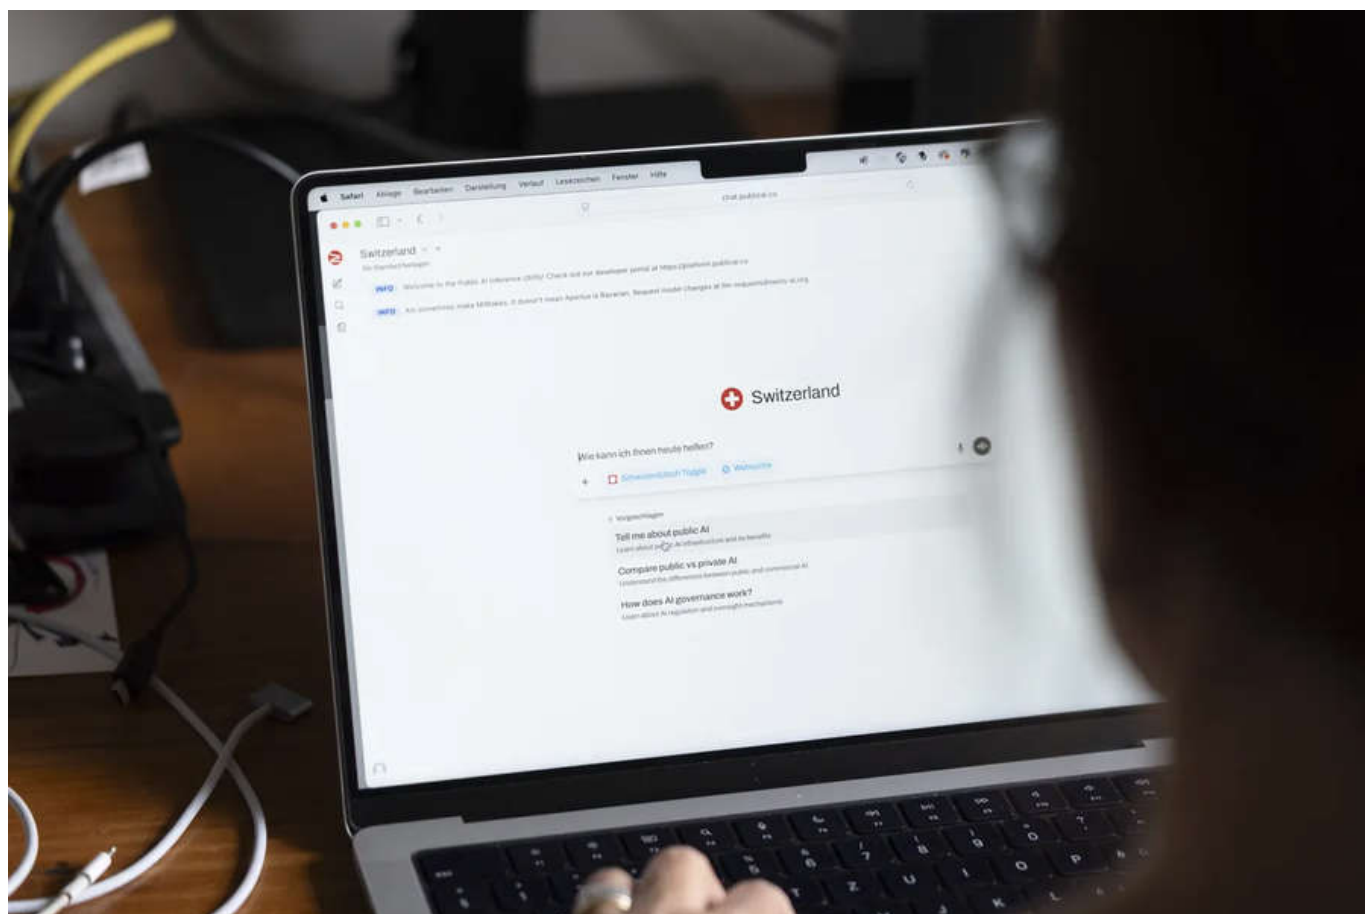
Federico Magnolfi, head of Natural Language Processing at Artificialy, estimates that Apertus 1.0 was already around two years behind the best US models when it launched. "Two years is enormous in a sector that advances every three months," he says. Closing that gap requires new training cycles, computing capacity and investment. As the flagship project of the Swiss AI Initiative, Apertus already uses a significant share of the public computing resources available. Lavrovsky acknowledges that the allocation of those resources is not without debate.

Not all of Apertus's shortcomings necessarily stem from the model itself. An AI model is only one component of a wider system. Software for tasks such as retrieving and ranking information relevant to a user's request also matters. Kruckenberg of Liip thinks that software components surrounding Apertus still have to be adjusted to work optimally. In Liip's tests, for example, an inaccurate answer about Zurich's population was not solely due to Apertus: the retrieval system had failed to select the most up-to-date source.

For Kruckenberg, this shows that a good AI model alone is not enough to be competitive. "The whole ecosystem around it also needs to be up to the task," he says.

In its second year, Apertus will have two major challenges: improving the model itself and strengthening the ecosystem on which it depends. Both will be essential if the Swiss AI model wants to remain competitive and achieve broader real-world adoption.

Edited by Gabe Bullard/vdv/ac



A year after its launch, it is time to take stock of Switzerland's AI model Apertus. Keystone / Gaetan Bally