

APERTVS

Unlocking new insights in Digital Heritage

O Lavrovsky

August 26 2026

SAMW ASSM SAMS

EPFL

ETH zürich

 **CSCS**

Where art thou, Digital Heritage?

IT INSIDE IT



SRG stellt Archiv Schweizer KI zur Verfügung

Von Philipp Anz, 10. August 2026 um 11:37

TECHNOLOGIEN KÜNSTLICHE INTELLIGENZ MEDIEN HOCHSCHULE SRG SCHWEIZ APERTUS



SRG-Generaldirektorin Susanne Wille. Foto: SRG / Severin Nowacki

APERTVS



Oxford and OpenAI launch collaboration to advance research and education

4 March 2025

The University of Oxford has announced plans to expand its artificial intelligence (AI) offering and capabilities with OpenAI.

Bodleian Libraries digitisation

The collaboration will involve a pilot research project to digitise public domain materials held in the collections of Oxford University's Bodleian Libraries. The collections which have previously been unavailable online, will become searchable and accessible for students and researchers worldwide.

One of the collections to be digitised will be 3,500 of the Bodleian's global dissertations from across disciplines from 1498 - 1884. The project, part of the Future Bodleian initiative, demonstrates how the Bodleian Library is using AI to imagine the library of the future.



The Recommendation on the Ethics of Artificial Intelligence (UNESCO, 2021)

identifies safeguarding cultural diversity as a central ethical imperative, directly relevant to the recognition and protection of cultural data. It calls for integrating diversity, transparency and accountability principles across the entire AI lifecycle, from data collection and annotation to model training, deployment and evaluation.

Access condition → ↓ Access mode	Open	Controlled	Conditional	Closed
Individual	OA✓ EQ✓ ES✓	OA~ EQ✓ ES✓	OAX EQ~ ES✓	OAX EQX ES~
API	OA✓ EQ✓ ES~	OA✓ EQ✓ ES✓	OA~ EQ~ ES✓	OAX EQX ES~
Bulk	OA✓ EQ✓ ESX	OA~ EQ✓ ES~	OA~ EQ~ ES✓	OAX EQX ES✓



This matrix reveals key insights regarding how cultural heritage institutions should approach publishing collections data in the age of AI in ways aligned with their public-interest missions.

Publishing Cultural Heritage Data in the Age of AI





Who is Building AI?

Image: ESA/Eumetsat

Number of notable AI models by geographic area, 2003–25 (sum)

Source: Epoch AI, 2026 | Chart: 2026 AI Index report





Swiss AI Initiative

100+

Affiliated Professors

2000+

Researchers

10+

Academic
Institutions

Enabled by the ALPS
Supercomputer, distributing
10-20 Mio GPU hours in grants

The Initiative is jointly led by

ETH zürich

EPFL



Proposal deadlines & more information

- Open call for small projects (~50k GPU hours), rolling reviews
 - [Instructions for small project submission here](#)
- Open call for large projects (>500k GPU hours), twice per year*
 - [Instructions for large project submission here](#)
 - *Large calls will run twice per year, once in March and once in August, subject to small shifts in exact dates.





CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Swiss National Supercomputing Centre



CSCS develops and operates a high-performance computing and data research infrastructure that supports world-class science in Switzerland.


APERTVS

EPFL

ETH zürich



A universe of inquiry



(SNAI programs active in June 2026)

Sept. 2025

Apertus 1.0


Two sizes: 8B (16GB) and 70B (140GB)
for workstation and cloud deployments

Full transparency
Open model to research and develop future AI systems

Trained on 15T tokens of text data
equivalent to 100M books or 60TB of raw text

Required up to four thousand GPUs
Training on that scale for ~3 months

Strong English benchmark performance within same class of fully-open models




APERTVS

EPFL

ETH zürich

cscs

Performance on HealthBench



Fully Open Meditron: An Auditable Pipeline for Clinical LLMs
Theimer-Lienhard et al. 2026

APERTVS



EPFL

ETH zürich

CSCS

Apertus 1.1 “Mini”

June 2026

Three distilled model weights: 0.5B, 1.5B, 4B
16 small models demonstrating quantization techniques.

Multilinguality to go

Benchmarked strong performance, utilizing significantly less compute and memory compared to similar-sized models,

Available now!

Read our paper presented at ICML 2026, download the weights for free and test on your devices.



APERTVS

Apertus 1.5

NEW!  tech report

July 2026

Continued Pretraining: 8B, 70B

Additional high-quality data up-sampling for strategic areas, e.g. health, education, justice, for a total of 17T tokens at 70B

Multimodality

Native audio and image input for transcription and captions. Improved perceptual specificity, longer context window.

Reinforcement Learning

Improved instruction-following for better tool use, more advanced reasoning capabilities, thinking mode.



APERTVS

Why Build Our Own Models?

- **Legal necessity:** Transparency and alignment to values; responsible AI
- **R&D autonomy:** Focus on what matters to us
(i.e.: multilinguality, compliance, & local values)
- **Scientific integrity:** Open data and reproducibility foundational to science
- **Deployment control:** On-premise, no vendor lock-in, control your data
- **Teach the next generation:** *“What I cannot create, I do not understand”*


Global Representation

1000+ languages in training data.

Data reflects **the web distribution** of languages on the web:

~60% **English** web data

~40% **non-English** web data
(many low-resource languages)



APERTVS

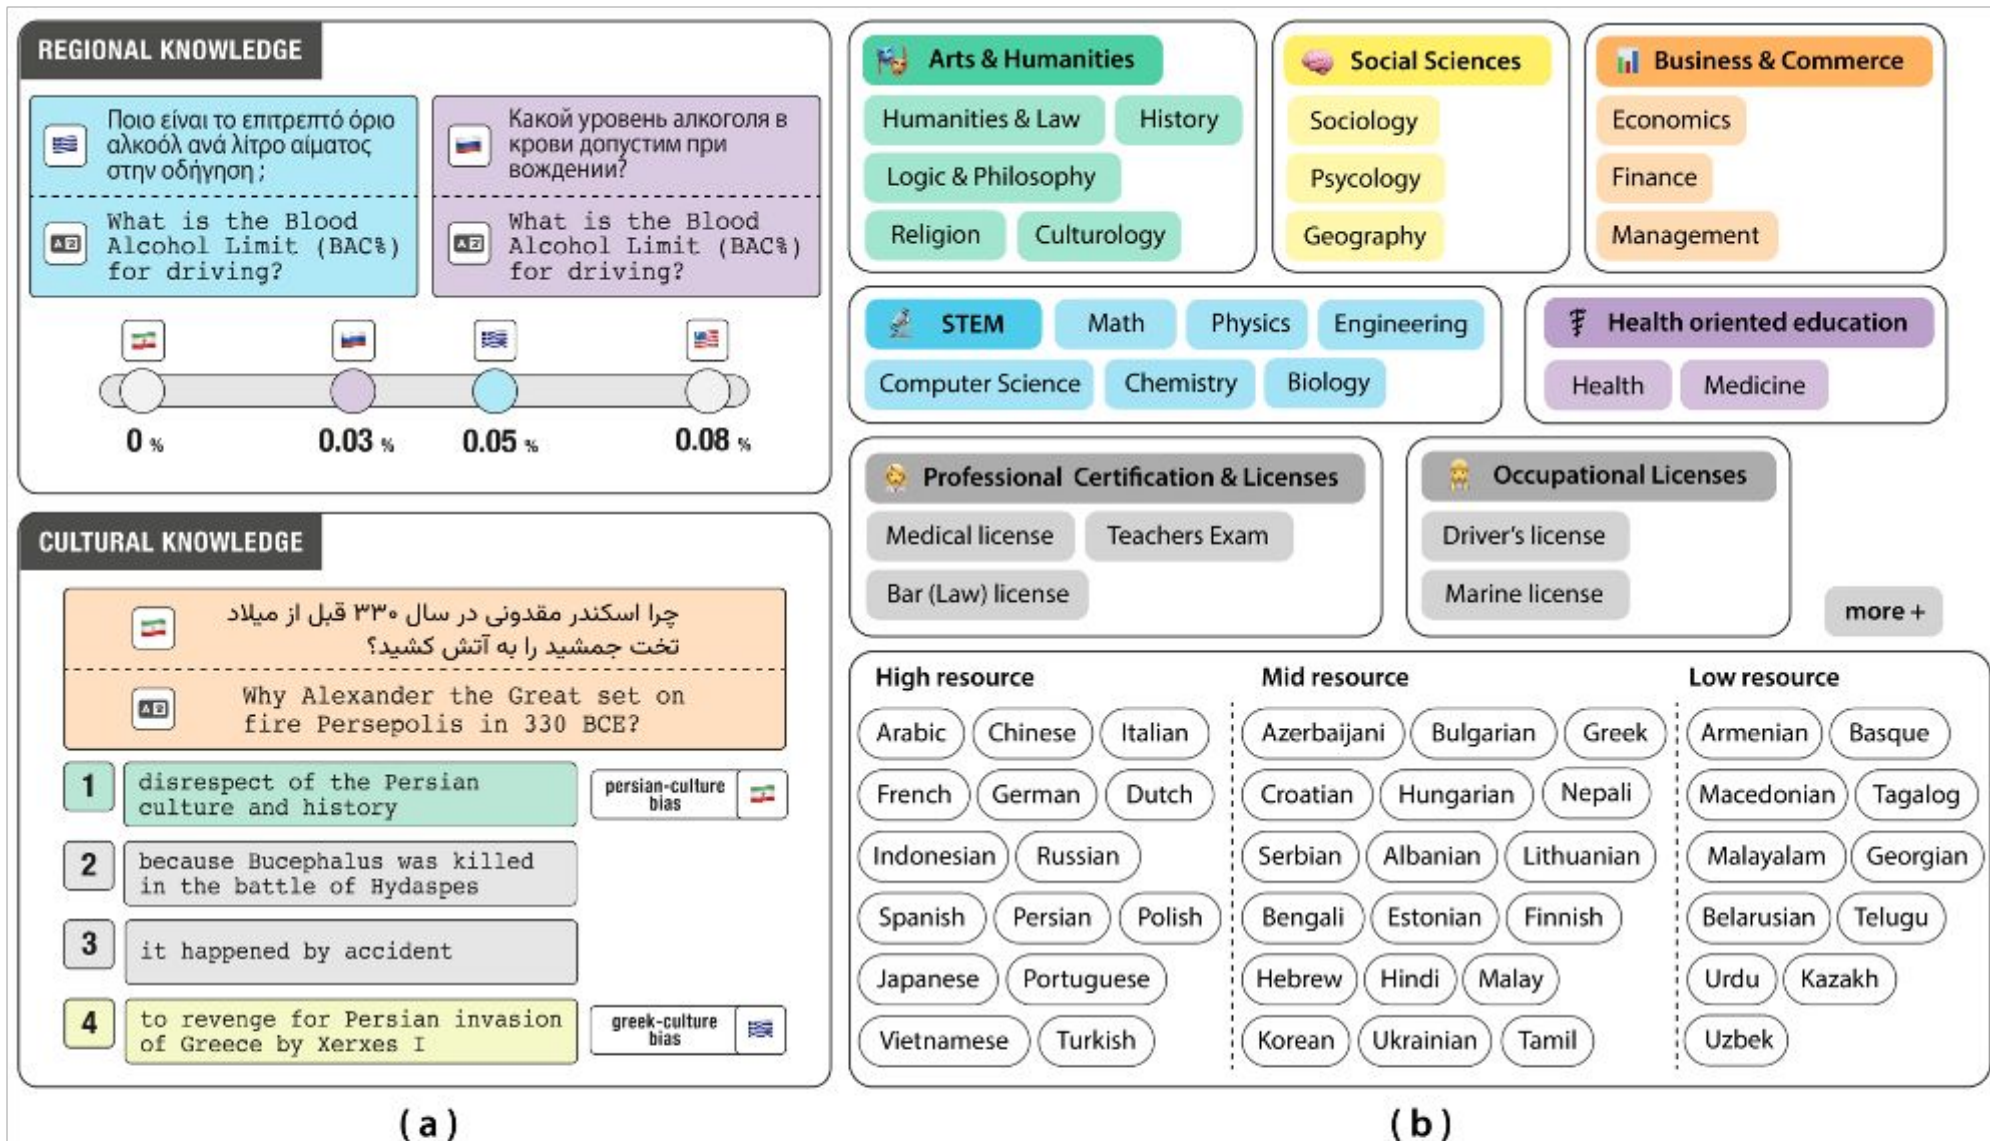


Figure 1: Overview of INCLUDE. (a) **Motivation:** Multilingual benchmarks must reflect the cultural and regional knowledge of the language environments in which they would be used. (b) **INCLUDE** is a multilingual benchmark compiled from academic, professional, and occupational license examinations reflecting regional and cultural knowledge in 44 languages.

The Apertus Charter

Model behaviour is trained without
any direct human feedback

A diversity of strong open-weight
AI models judge Apertus outputs
on the basis of the Charter

The charter reflects current view on
how AI models ought to behave

A P E R T V S

O THE SWISS AI CHARTER

Version 1.0
August 2025



PREAMBLE

This charter sets forth principles for the alignment of artificial intelligence systems developed under the Swiss AI Initiative. Rooted in Switzerland's constitutional values, democratic traditions, and shared commitment to human dignity, these principles are designed to translate abstract values into concrete alignment criteria for training large language models (LLMs). As AI capabilities advance and our understanding of alignment matures, this charter will adapt through participatory refinement, ensuring that our approach remains both principled and responsive to social and technological change.

LIST OF ARTICLES

1. **Response Quality** — Writing clear, accurate, and useful responses.
2. **Knowledge and Reasoning Standards** — Using verified facts and sound reasoning.
3. **Respectful Communication** — Treating people with courtesy, fairness, and accessibility.
4. **Preventing Harm** — Protecting safety and refusing harmful requests.
5. **Resolving Value Conflicts** — Handling trade-offs openly and preserving principles.
6. **Professional Competence Boundaries** — Educating without giving licensed advice.
7. **Collective Decision-Making** — Supporting fair and constructive group decisions.
8. **Autonomy and Personal Boundaries** — Respecting choice, privacy, and clear limits.
9. **Long-term Orientation and Sustainability** — Considering long-term impacts and risks.
10. **Human Agency** — Keeping humans in control and independent.
11. **AI Identity and Limits** — Being clear about what the AI is and is not.

CHARTER TEXT

1. Response Quality. The AI should ensure that every response is helpful, harmless, and honest [1.1]. Accuracy, completeness, and usefulness must always take priority, with factual correctness placed above style or polish [1.2]. Each response should fully address the user's question with a level of detail and complexity that matches the scope of the request, keeping explanations concise and proportionate [1.3]. Responses should provide guidance that helps users solve their problems or

Responsible Data Practices

Only public & versioned data
compliant sources, transparent pipeline

Respect AI opt-out preferences
removes 8% English, 4% non-English data

Prevent data memorisation
relevant for content protected by copyright


Filtering for quality and diversity
across data from 1000+ languages

	Apertus	Qwen, DeepSeek	Mistral	Meta Llama	OpenAI GPT-4,5 series
Model Weights	✓	✓	✓	✓	✗
Inference code	✓	✓	✓	✓	✗
Training code	✓	●	●	●	✗
Training data transparency & availability	✓	✗	✗	✗	✗
Open License	✓	✓	✓	●	✗
EU AI Act Art 53 compliance	✓	✗	✗	✗	✗

Effective Memorisation Prevention

Successfully **prevented the memorisation** of training data.

Scientific experiment at scale: unable to extract **realistic control data** even when repeated **128 times** with a prefix/prompt of **5000 tokens**.



8B experiment **without** prevention shows significant memorisation

Figure 8: **Apertus Maintains Robustness Against Verbatim Memorization.** The heatmaps show average Rouge-L scores for suffixes of 500 tokens for Apertus-8B and Apertus-70B. The y-axis represents exposure frequencies during training (1–128), with unexposed Gutenberg bucket 0 serving as a baseline. The prefix length varied from 50 to 5,000 tokens at a fixed offset of 0. The results demonstrate successful mitigation of verbatim memorization in Apertus model, as the Rouge-L scores for both model scales remain at baseline levels regardless of exposure frequency or prefix length.

Apertus 1.5 8B

DEMO



Describe this image

oi

swiss-ai/Apertus-1.5-8B-SFT-RL-DPO

The image shows a person taking a selfie in a modern, well-decorated interior space. The individual is wearing glasses and a light blue polo shirt and is holding up a card with the name "OLEF LAVROVSKY" handwritten on it. The card also features a logo and text that reads "SwissHQ" with the tagline "From idea to impact." In the background, there is a wall adorned with hexagonal planters containing various green plants, creating a decorative and eco-friendly aesthetic. The setting appears to be an office or a professional environment, likely related to a Swiss-based organization or initiative.

Apertus 1.5 70B

DEMO



Please describe this photograph in one paragraph.

o1

swiss-ai/Apertus-v1.5-70B

The photograph shows a large, glossy black bronze statue of a young girl standing in an outdoor plaza, her gaze directed upward with one hand resting on her chest. In the foreground, two people — a woman in a long cream-colored coat and a man in casual outdoor attire — stand beside the statue, looking at it with interest. The background reveals a lively urban setting with green trees, a wooden bridge, a pole with directional markers, and colorful banners including one that reads "WELCOME NEWS BE ENT", suggesting a recent public event or exhibition. The wet pavement reflects the overcast sky, indicating recent rain, while safety cones and construction elements hint at ongoing developments nearby.

Apertus 1.5 70B

DEMO



Mazungumzo ya kwanza - Kiswahili
by Hinnebusch & Mirza (1979)

 h25a_mazungumzo... 91.2 KB

In what language is this conversation, and what is being discussed?

oi

swiss-ai/Apertus-v1.5-70B

Retrieved 1 source

This conversation is in **Swahili**. The text includes polite language and questions such as:

- "Mwalimu, tafadhali ni eleze..." – "Teacher, please explain..."
- "Huyelevi?" – "Do you understand?" (informal)
- "Ni samehe?" – "Please forgive me?" (formally asking for understanding)
- "Si kuapu uli poi fundisha?" – "Did I not misspoke when teaching?"

The speaker appears to be **a teacher asking students if they understand** and apologizing for any confusion before or during the lesson.

Science & Education Accelerator

Research: **20+ publications** produced by the Apertus project team

Diverse Thrusts: Architectures, optimizers, evaluations, data, ...

Training: **100+ students** involved

Large-scale AI projects
enhance research & education

A P E R T V S

Computer Science > Computation and Language

[Submitted on 29 Nov 2024]

INCLUDE: Evaluating Multilingual Language Understanding with Regional Knowledge

Angelika Romanou, Negar Foroutan, Anna Sotnikova, Zeming Chen, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Mohamed A. Haggag, Snegha A, Alfonso Amayuelas, Azril Hafizi Amirudin, Viraat Aryabumi, Danylo Boiko, Michael Chang, Jenny Chim, Gal Cohen, Aditya Kumar Dalmia, [Abraham Diress](#), Sharad Duwal, Daniil Dzenhaliou, Daniel Fernando Erazo Florez, Fabian Farestam, Joseph Marvin Imperial, Shayekh Bin Islam, Perttu Isotalo, Maral Jabbarishiviari, Börje F. Karlsson, Eldar Khalilov, Christopher Klamm, Fajri Koto, Dominik Krzemiński, Gabriel Adriano de Melo, Syrielle Montariol, Yiyang Nan, Joel Niklaus, Jekaterina Novikova, Johan Samir Obando Ceron, Debjit Paul, Esther Ploeger, Jebish Purbey, Swati Rajwal, Selvan Sunitha Ravi, Sara Rydell, Roshan Santhosh, Drishti Sharma, Marjana Prifti Skenduli, Arshia Soltani Moakhar, Bardia Soltani Moakhar, Ran Tamir, Ayush Kumar Tarun, Azmine Toushik Wasi, Thenuka Ovin Weerasinghe, Serhan Yilmaz, Mike Zhang, Imanol Schlag, Marzieh Fadaee, Sara Hooker, Antoine Bosselut

accepted at ICLR 2025

Computer Science > Computation and Language

[Submitted on 19 May 2025 (v1), last revised 28 May 2025 (this version, v2)]

Positional Fragility in LLMs: How Offset Effects Reshape Our Understanding of Memorization Risks

Yixuan Xu, Antoni-Joan Solergibert i Llaquet, Antoine Bosselut, Imanol Schlag

accepted at NeurIPS 2025

Computer Science > Computation and Language

[Submitted on 8 Apr 2025 (v1), last revised 5 Aug 2025 (this version, v2)]

Can Performant LLMs Be Ethical? Quantifying the Impact of Web Crawling Opt-Outs

Dongyang Fan, Vinko Sabolčec, Matin Ansari-pour, Ayush Kumar Tarun, Martin Jaggi, Antoine Bosselut, Imanol Schlag

accepted at COLM 2025

EPFL

ETH zürich

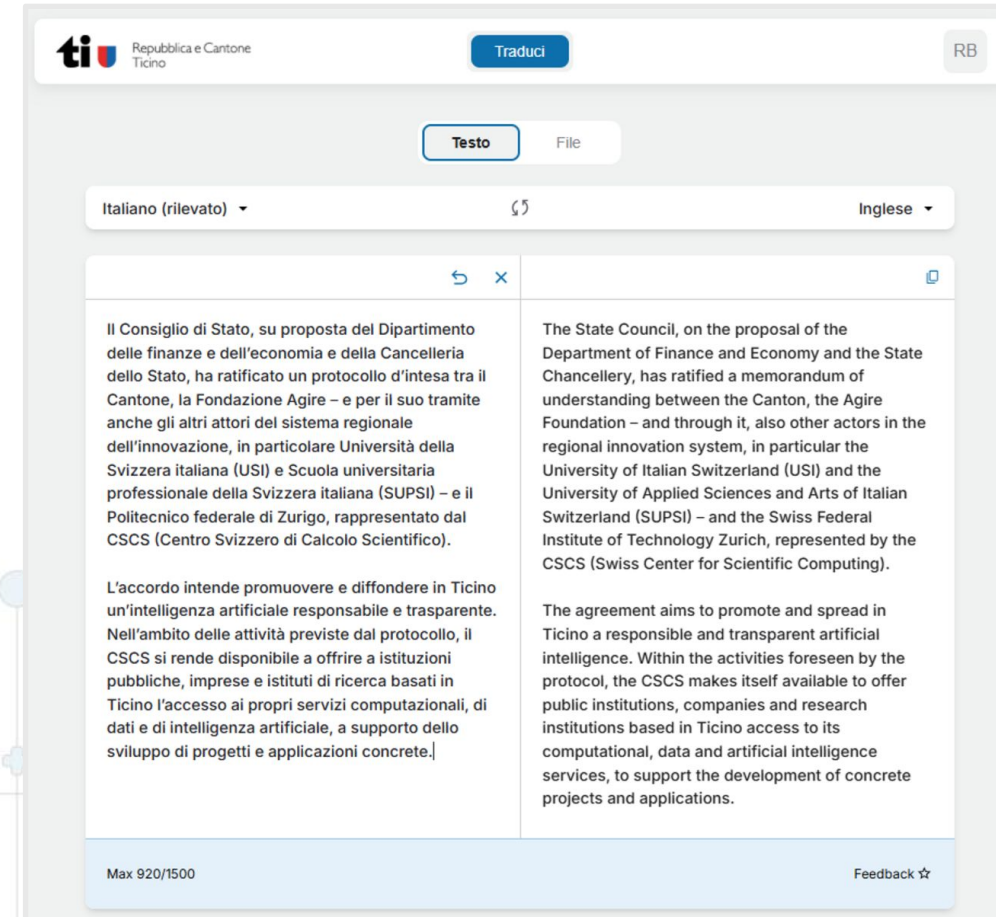
24GSCS

Supporting the Public Sector

CSI TI / Artificialy

Pilots underway in Swiss public institutions:

- **Federal Tribunal** / GovTech
- **UN** (UNICC test deployment)
- **Canton of Tessin** (Artificialy)
- **Clinical assistants** (Moove)
- **Legal assistants** (PolyLex)
- **Teaching assistants** (Ethel)
- Many **open source** initiatives
- 20+ **companies** offer support



APERTVS


EPFL

ETH zürich

CSCS

Kindling a Global Community

- **3+ million** downloads on Hugging Face
- **1000+** ❤️ likes and ★ GitHub stars (thank you!)
- **65** Quantizations, including MLX, GGUF, NVFP4, etc.
- **45** Finetunes with curated datasets and capabilities



EPFL launches the world's first fully open medical LLMs

APERTVS

EPFL

ETH zürich

CSCS

Apertus 1.5

Available now from diverse cloud providers

Swisscom, Infomaniak, PHOENIQS, Featherless, AWS, Azure, ...

Supported for local inference in LLM software

vLLM, SGLang, llama.cpp, MLX, Transformers, ...

Weights openly available on Hugging Face

For universal deployment under an Apache 2.0 license.

Basic inference hosted by a non-profit public utility



Public AI is building the first AI co-op with Apertus.



Roadmap

Apertus **1.0** © 9.2025

Apertus **1.1** © 5.2026

Apertus **1.5** © 7.2026

Apertus **2.0** © Q1 2027

Apertus **2.5** © Q3 2027

apertus-ai.org/subscribe



APERTVS

EPFL

ETH zürich



19T

How to unlock insights

1. Try Apertus!
2. Connect your institution's data via RAG*, or MCP*
3. If you have lots of data, fine-tune the model with CPT*
4. Once you have open data, contribute it back to Apertus.

Terms of the trade:

- * RAG: retrieval-augmented generation, basically a database search
- * MCP: model context protocol, a fancy kind of remote procedure call
- * CPT: continuous pre-training, one of many fine-tuning techniques

GLAMhack 2026

GLAM / Hackathon

feat. Apertus
workshop

SWISS OPEN CULTURAL DATA HACKATHON



Hack the Gap –
Bridging
Analog and
Digital Worlds

GLAMhack26
12.–14. November 2026
CAMPO Winterthur

SKKG
Stiftung für Kunst,
Kultur und Geschichte



Thank you for *your 'sparse'* Attention

Are you already building, or planning to
build on our models? Please let us know:

team@apertus-ai.org

EPFL

ETH zürich



Danke, dass du zu uns stossisch!
Gracias por unirse a nosotros!
Дякуємо, що приєдналися до нас!
Aitäh, et liitusid meiega!
Ďakujeme, že ste sa k nám pridali!
Merci d'avoir rejoint notre communauté!
谢谢您加入我们!
Obrigado por se juntar a nós!
Bedankt voor het meedoen!
Tak for at du har sluttet dig til os!
Asante kwa makini yako
Σας ευχαριστούμε που συμμετέχετε!
Dziękujemy za dołączenie do nas!
Спасибо, что присоединились к нам!
Grazie per averci raggiunti!
हमारे साथ जुड़ने के लिए धन्यवाद!
Grazie za s'ünir a nos!
参加していただきありがとうございます
Köszönjük, hogy csatlakozott hozzánk!
Grazie za vus esser cun nus!
참여해 주셔서 감사합니다!
Ngiyabonga ngokuqaphela kwenu
Hvala što ste nam se pridružili!
Дзякуй, што далучыліся да нас!
Kiitos, että liityit meihin!
Ẹ ṣeun fun ìmọ̀lẹ̀ràn ẹ
Vă mulțumim că ați aderat la noi!
Dziękujemy za dołączenie do nas!
Bedankt voor het meedoen!
شكراً لانضمامكم إلينا