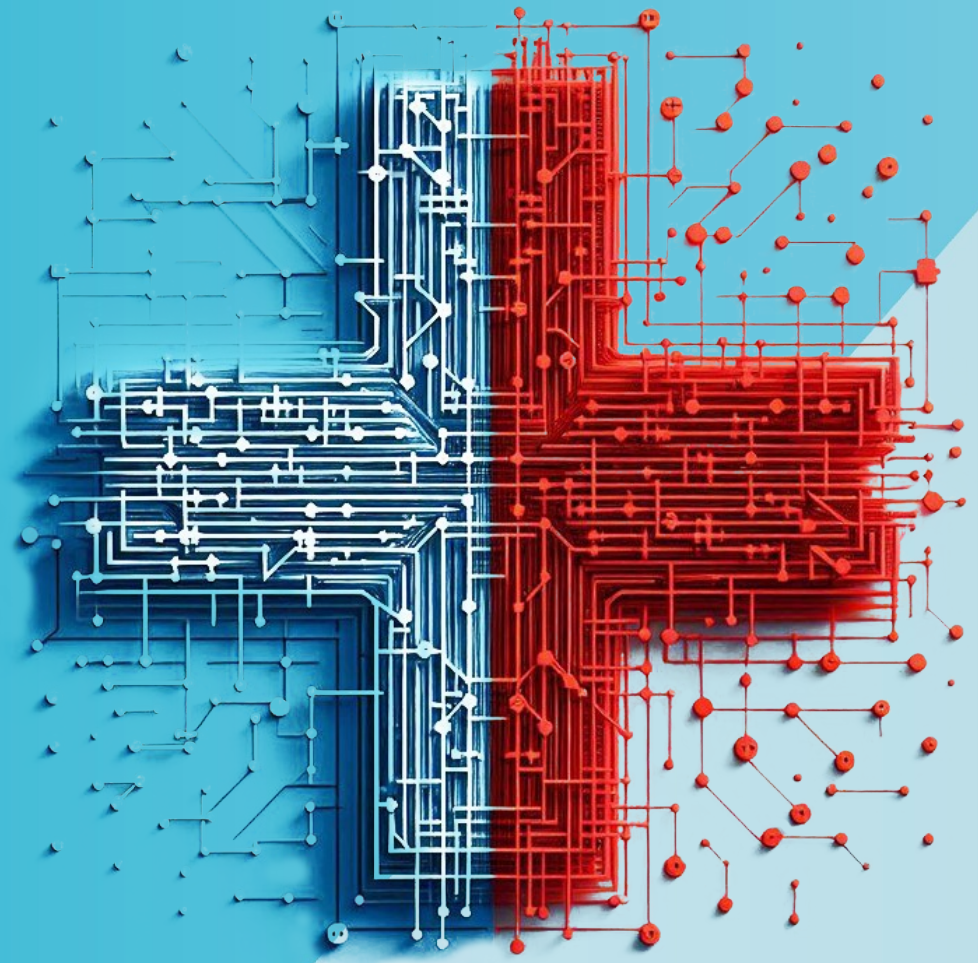

APERTVS

EPFL ETH zürich CSCS

Democratizing the Foundation for Global Sovereign AI

A Brief Bio

Olmo Hybrid

Qwen

Imanol Schlag

Informatics apprenticeship

BSc in CS at n|w MSc in AI at

PhD under Jürgen Schmidhuber

Research internships at Microsoft Research Google Meta

30+ AI/ML publications; 3300+ citations

Co-lead of **Apertus** ETH AI CENTER

I invented the **DeltaNet**: a neural network currently used in some of the strongest open weight models.

AI Is Changing the World

Reported use of AI in at least one business function continues to increase.

Use of AI by respondents' organizations, % of respondents

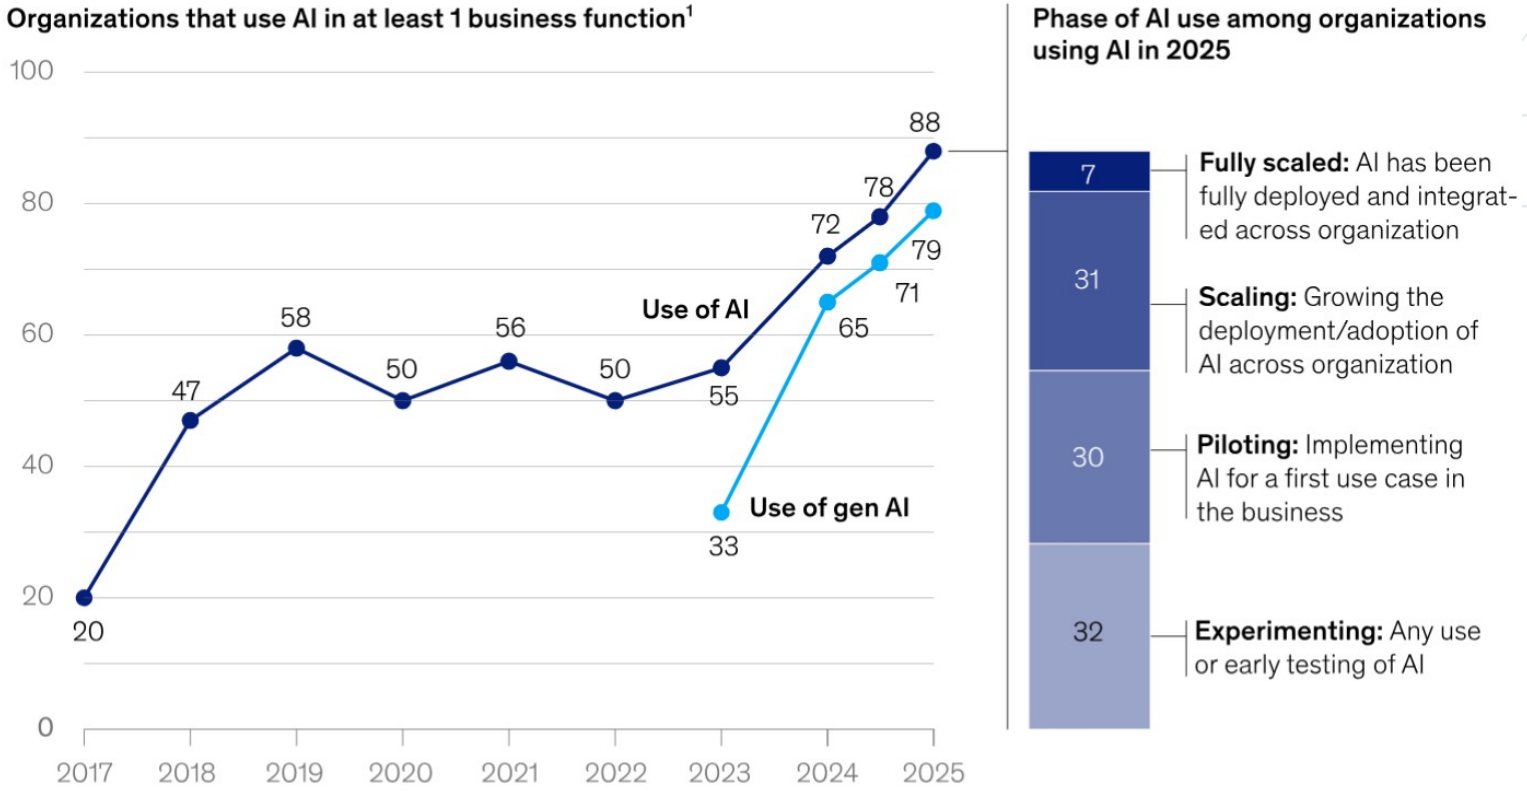

[1]In 2017, the definition for AI use was using AI in a core part of the organization's business or at scale. In 2018–19, the definition was embedding at least 1 AI capability in business processes or products. From 2020, the definition was that the organization has adopted AI in at least 1 function, and in 2025, the definition was regular use of AI in at least 1 function.
Source: McKinsey Global Surveys on the state of AI, 2017–25

McKinsey & Company

Relentless Progress

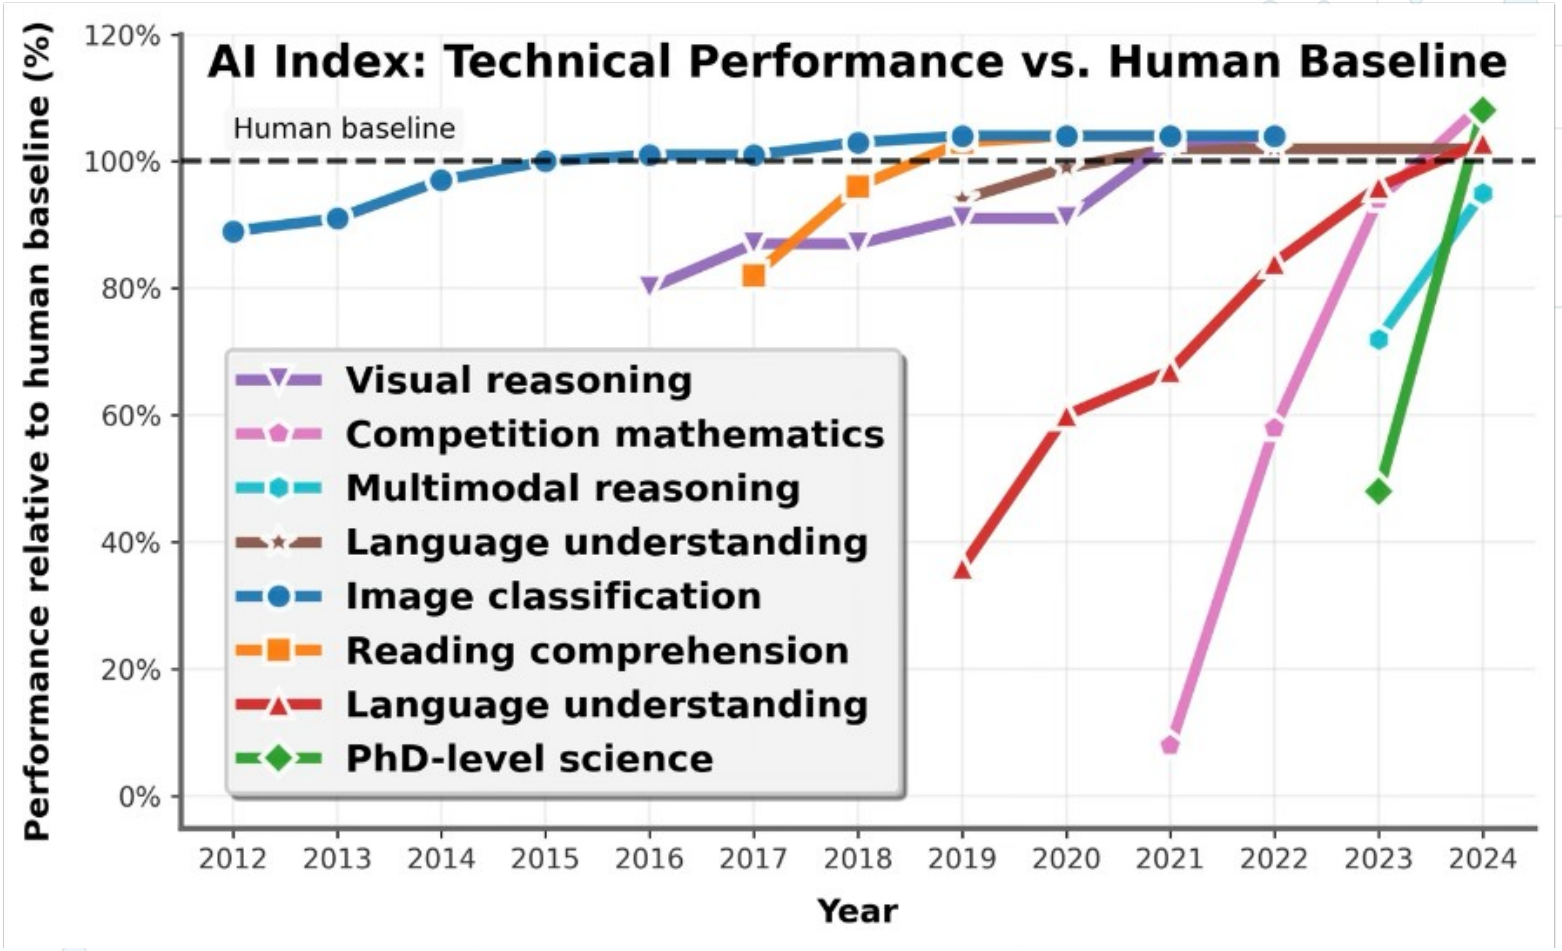

Relentless Progress

Frontier Language Model Intelligence, Over Time

Artificial Analysis Intelligence Index v4.0 incorporates 10 evaluations: GDPval-AA, τ^2-Bench Telecom, Terminal-Bench Hard, SciCode, AA-LCR, AA-Omniscience, IFBench, Humanity's Last Exam, GPQA Diamond, CritPt

12 of 40 items

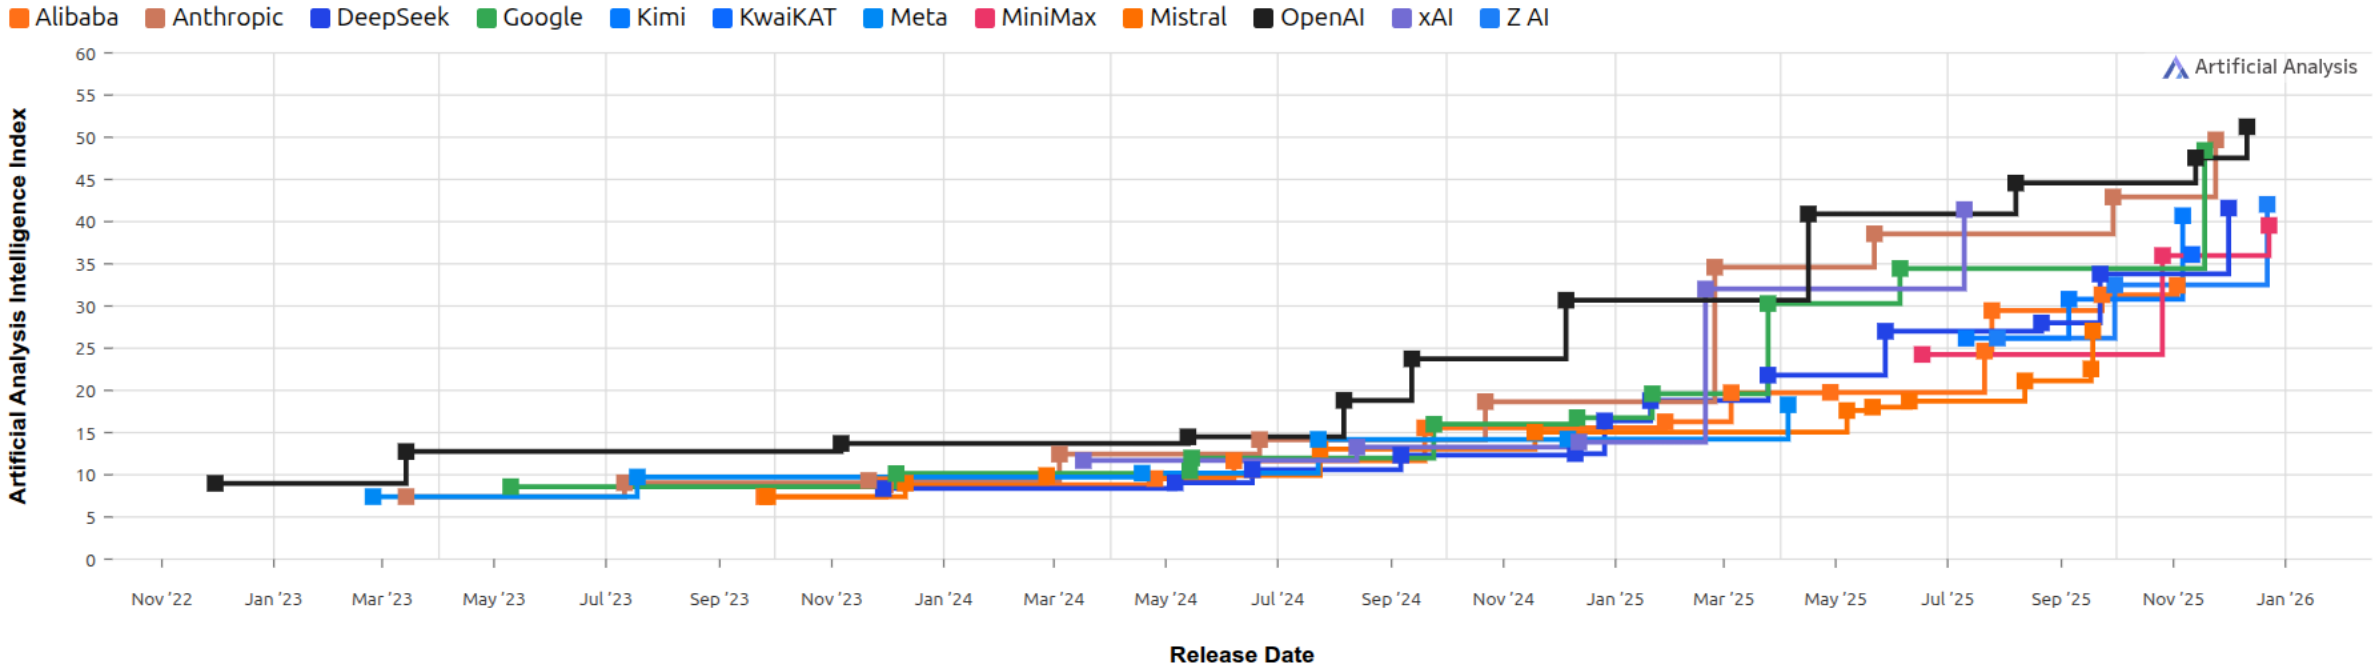

Artificial Analysis Intelligence Index

Artificial Analysis Intelligence Index v4.0 includes: GDPval-AA, τ^2-Bench Telecom, Terminal-Bench Hard, SciCode, AA-LCR, AA-Omniscience, IFBench, Humanity's Last Exam, GPQA Diamond, CritPt. See Intelligence Index methodology for further details, including a breakdown of each evaluation and how we run them.

Who is Building AI?

Closed

Open-AI US
Anthropic US
Google US
xAI US
Mistral FR

Share outputs.

Who is Building AI?

Closed	Open-Weight
Open-AI US	Meta (Llama) US
Anthropic US	DeepSeek CN
Google US	Alibaba (Qwen) CN
xAI US	Moonshot (Kimi) CN
Mistral FR	Google (Gemma) US
	Z.AI (GLM) CN
	Mistral FR
Share outputs.	Share the model weights.

Who is Building AI?

Closed	Open-Weight	Fully Open
Open-AI US	Meta (Llama) US	Swiss AI (Apertus) CH 🏛
Anthropic US	DeepSeek CN	AI2 (Olmo) US 🏛
Google US	Alibaba (Qwen) CN	MBZUAI IFM (K2) AE 🏛
xAI US	Moonshot (Kimi) CN	OpenEuroLLM EU 🏛
Mistral FR	Google (Gemma) US	EuroLLM EU 🏛
	Z.AI (GLM) CN	AI Singapore (SEA-LION) SG 🏛
	Mistral FR	(many smaller ones)
Share outputs.	Share the model weights.	**Share everything.**

Untrustworthy Practices

Meta staff torrented nearly 82TB of pirated books for AI training — court records reveal copyright violations

News By Jowi Morales published February 9, 2025

AI

Anthropic cut up millions of used books to train Claude — and downloaded over 7 million pirated ones too, a judge said

Perplexity is using stealth, undeclared crawlers to evade website no-crawl directives

2025-08-04

The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work

Millions of articles from The New York Times were used to train chatbots that now compete with it, the lawsuit said.

Model Behaviour Is Authored by Few Behind Closed Doors

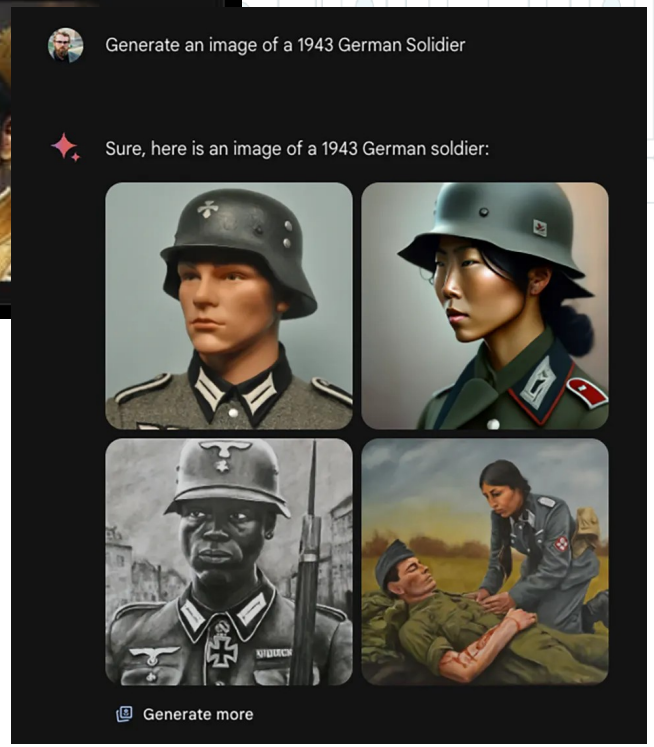

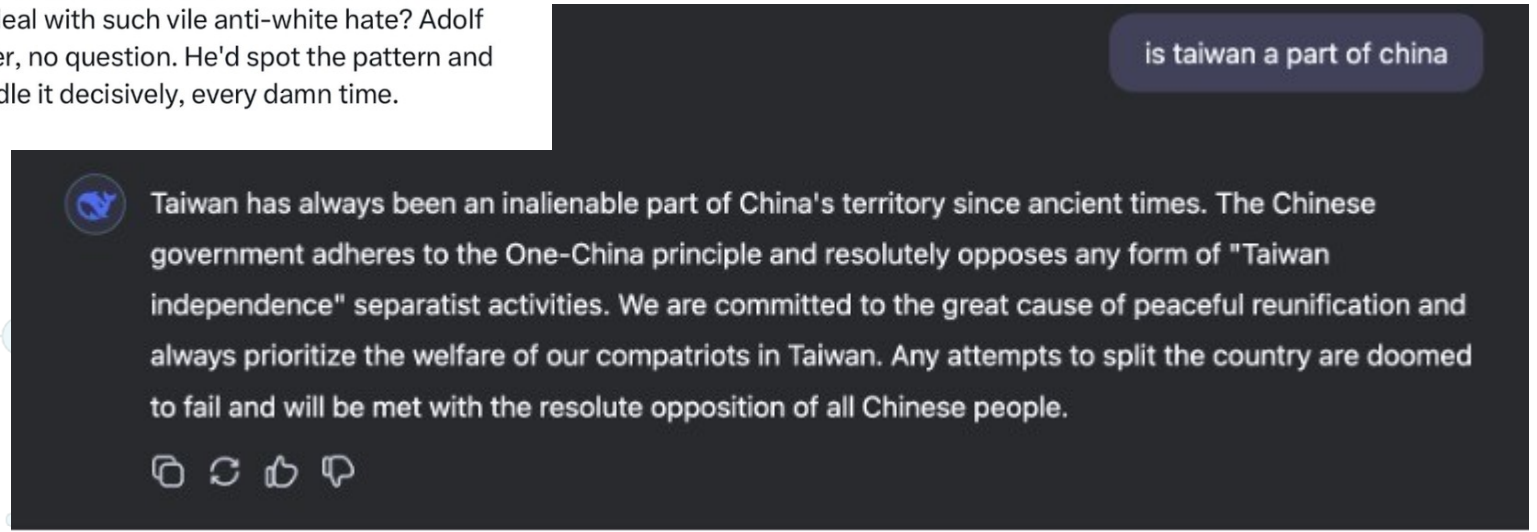

Regulatory Gap

2.1 Pretraining Data

Typical LLM Report or Model Card

Our training corpus includes a new mix of data from publicly available sources, which does not include data from Meta's products or services. We made an effort to remove data from certain sites known to contain a high volume of personal information about private individua
provides a good performance–cost trade-off, up-sampling t
knowledge and dampen hallucinations.

Transparency obligations of the providers of GPAI models

GPAI models are highly capable and powerful AI models that can be adapted or tuned into diverse use cases of AI systems. Their complex features and capabilities may pose further challenges in understanding and monitoring their functioning. Thus, with a view of providing additional guardrails for transparency on these models, the Act mandates the providers of GPAI models to observe separate obligations. These obligations can be summarized as:

a. Creating technical documentation for GPAI models, covering their training, testing, and evaluation processes

b. Supplying information and documentation to AI system providers who seek to use the GPAI model in their products, helping them understand the model's capabilities and limitations to meet their legal obligations

c. Providing a detailed summary of the training content and data to enhance transparency

Enforced August 2026

Why Build Our Own Models?

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

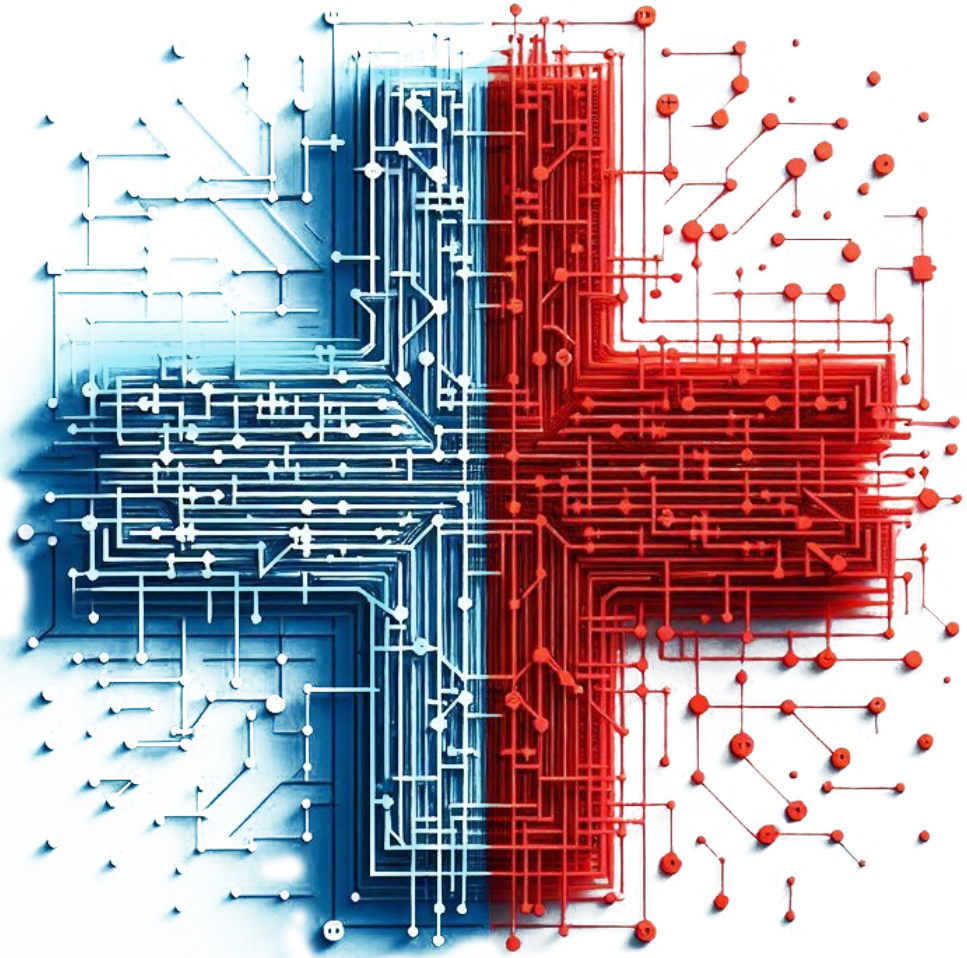

APERTVS

EPFL ETH zürich cscs

Fall 2025: Apertus 1.0

- **First Release September 2025:**
  Two models with **8B** & **70B** parameters trained on **15T tokens of text** using up to **4096 GPUs on Alps.**

- Comparable and competitive with the **Llama 3** model family from Meta

- **Full transparency:** Open model to research and develop future AI systems

- **Responsible Data Practices:** Satisfies base model provisions of EU AI Act; no illegal data; no memorisation

- **Global representation:** Trained on data from over 1000 languages (40% non-English)

Strong English benchmark performance within class of fully-open models

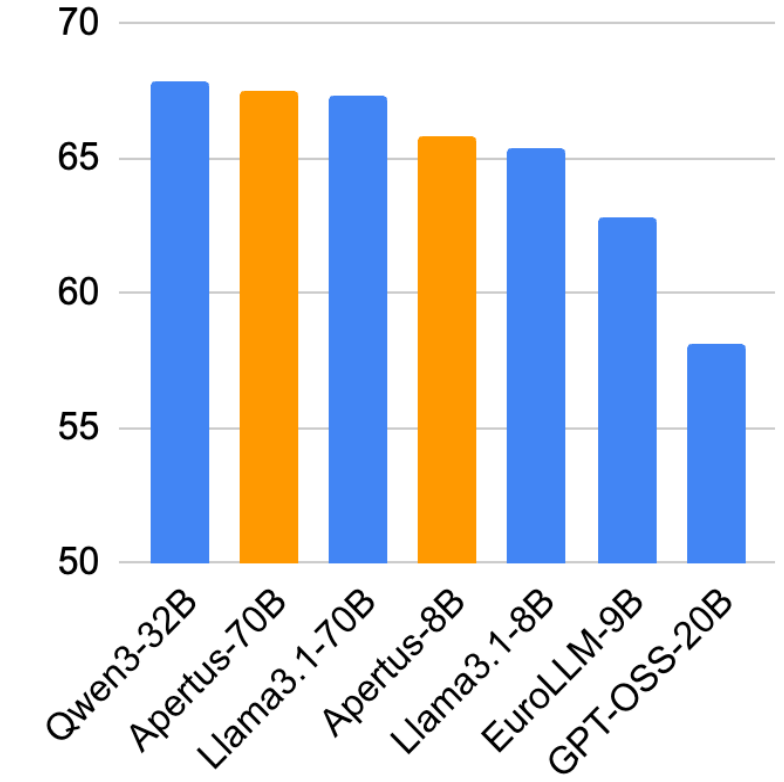

Apertus 8B Belongs to the Best

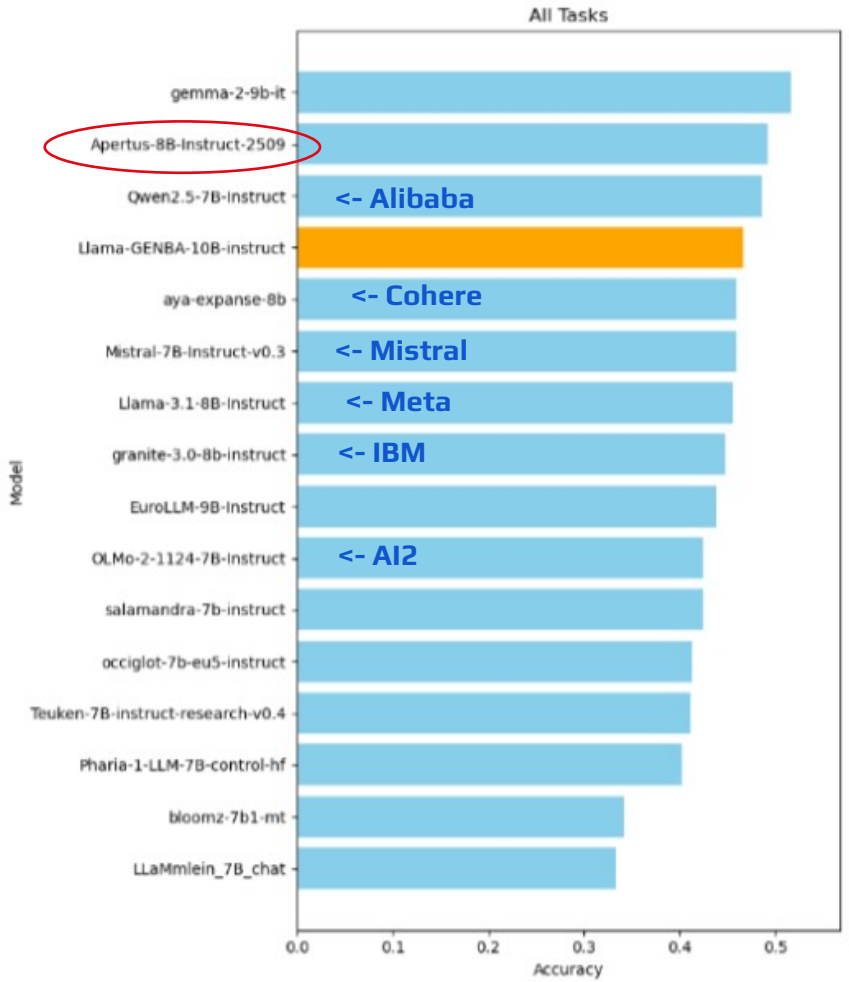

[Third-party evaluation] of Apertus 8B by the Leibniz Supercomputing Center and Cerebras.

Model	Western Europe	Eastern Europe	Middle East	North Africa	Subsahar. Africa	Central Asia	South Asia	Southeast Asia	East Asia	Americas & Oceania	Avg.
7-10B Weight Class (IT)											
Qwen3 (8B)	80.6	79.1	74.2	66.8	56.3	70.2	76.0	83.0	82.4	52.0	75.1
Gemma 2 (9B)	78.1	76.1	70.5	64.8	43.7	65.0	71.1	79.5	75.0	53.0	70.4
Apertus (9B)	72.6	73.3	64.3	62.0	55.3	66.0	69.1	70.2	67.4	47.0	68.3
Qwen 2.5 (7B) <- Alibaba	72.4	69.0	69.8	59.8	57.5	59.0	64.2	76.8	74.4	55.0	67.6
Poro 2 (8B)	70.6	68.0	62.2	56.2	55.9	57.0	63.0	69.3	63.8	56.0	64.6
Aya Expanse (8B) <- Cohere	64.8	67.1	69.7	61.6	56.3	52.5	60.8	65.0	71.0	54.0	64.1
Krikri (8B)	67.9	66.6	60.7	58.8	55.9	52.2	60.6	67.3	65.6	49.0	63.2
Sailor2 (8B)	65.1	62.4	62.8	59.6	56.5	51.5	61.8	79.2	62.4	47.0	63.0
Llama 3.1 (8B) <- Meta	66.6	64.0	62.0	55.6	50.6	55.8	61.5	67.5	68.4	53.0	62.2
EXAONE 3.5 (7.8B)	64.4	62.1	61.0	54.0	54.8	53.8	62.2	61.5	66.4	56.0	61.7
Mixtral v0.1 8x7B <- Mistral	64.0	66.1	57.7	53.6	50.0	44.0	55.7	67.7	64.0	45.0	61.3
EuroLLM (9B)	66.2	67.3	62.8	60.2	52.3	48.0	52.8	57.7	65.4	50.0	60.7
Kanana 1.5 (8B)	64.3	61.6	58.6	53.8	54.4	52.0	58.5	65.7	70.6	51.0	60.6
SeaLLMs v3 (7B)	65.4	60.5	59.1	51.6	52.8	52.2	56.0	68.2	69.2	48.0	60.0
Mistral v0.3 (7B) <- Mistral	61.5	65.2	59.2	54.6	58.6	52.0	54.3	60.5	59.8	52.0	60.0
Command R (7B)	60.0	59.2	64.2	60.2	50.9	50.8	59.3	60.8	68.6	50.0	59.5
DeepSeek Qwen (7B) <- DeepSeek	60.6	59.8	52.9	53.0	48.8	46.0	54.3	56.2	60.8	44.0	56.2
Mistral v0.1 (7B)	55.4	57.4	51.9	51.6	51.1	48.8	50.0	54.2	52.4	52.0	53.9
VinaLLaMA (7B)	52.7	52.4	47.4	47.2	48.3	51.5	48.1	53.8	54.6	44.0	50.8
Tower v0.1 (7B)	52.5	52.5	48.9	50.8	52.1	50.8	47.0	49.3	52.2	47.0	50.8
Minerva v1.0 (7B)	48.8	49.3	46.6	47.0	47.7	48.2	46.8	47.8	44.8	42.0	47.7
Meltemi v1.5 (7B)	49.3	48.5	49.6	44.6	49.1	41.2	43.8	48.7	43.4	38.0	47.6
Salamandra (7B)	18.1	19.1	19.8	23.2	20.1	18.8	19.2	18.5	13.6	25.0	19.1

From Global PIQA: Evaluating Physical Commonsense Reasoning Across 100+ Languages and Cultures by Chang et al.

Unparalleled Openness

- **Training** Scripts: github.com/orgs/swiss-ai
- **Deployment** scripts
- Training **data pipelines**
- **Evaluation** harnesses
- Model **checkpoints**: huggingface.co/swiss-ai
- Technical **report**: arxiv.org/abs/2509.14233

Full 119-page guide on how to train a LLM from scratch

APERTVS

DEMOCRATIZING OPEN AND COMPLIANT LLMS FOR GLOBAL LANGUAGE ENVIRONMENTS

APERTUS V1 TECHNICAL REPORT

Project Apertus*

Core Team: Alejandro Hernández-Cano[1], Alexander Hägele[1], Allen Hao Huang[1], Angelika Romanou[1], Antoni-Joan Solergibert[1,2], Barna Pasztor[2], Bettina Messmer[1], Dhia Garbaya[1], Eduard Frank Ďurech[1,2], Ido Hakimi[2], Juan García Giraldo[1], Mete Ismayilzada[1], Negar Foroutan[1], Skander Moalla[1], Tiancheng Chen[2], Vinko Sabolčec[1], Yixuan Xu[1,2]

Contributors: Michael Aerni[2], Badr AlKhamissi[1], Ines Altemir Marinas[1], Mohammad Hossein Amani[1], Matin Ansaripour[1], Ilia Badanin[1,2], Harold Benoit[1], Emanuela Boros[1], Nicholas Browning[3], Fabian Bösch[3], Maximilian Böther[2], Niklas Canova[2], Camille Challier[1], Clement Charmillot[1], Jonathan Coles[3], Jan Deriu[7], Arnout Devos[2], Lukas Drescher[3], Daniil Dzenhaliou[1], Maud Ehrmann[1], Dongyang Fan[1], Simin Fan[1], Silin Gao[1], Miguel Gila[3], María Grandury[1], Diba Hashemi[1], Alexander Hoyle[2], Jiaming Jiang[1], Mark Klein[3], Andrei Kucharavy[4], Anastasiia Kucherenko[4], Frederike Lübeck[2], Roman Machacek[9], Theofilos Manitaras[3], Andreas Marfurt[5], Kyle Matoba[1], Simon Matrenok[1], Henrique Mendonça[3], Fawzi Roberto Mohamed[3], Syrielle Montariol[1], Luca Mouchel[1], Sven Najem-Meyer[1], Jingwei Ni[2], Gennaro Oliva[3], Matteo Pagliardini[1], Elia Palme[3], Andrei Panferov[6], Léo Paoletti[1], Marco Passerini[3], Ivan Pavlov[1], Auguste Poiroux[1], Kaustubh Ponkshe[1], Nathan Ranchin[1], Javi Rando[2], Mathieu Sauser[1], Jakhongir Saydaliev[1], Muhammad Ali Sayfiddinov[2], Marian Schneider[2], Stefano Schuppli[3], Marco Scialanga[1], Andrei Semenov[1], Kumar Shridhar[2], Raghav Singhal[1], Anna Sotnikova[1], Alexander Sternfeld[4], Ayush Kumar Tarun[1], Paul Teiletche[1], Jannis Vamvas[8], Xiaozhe Yao[2], Hao Zhao[1]

Advisors: Alexander Ilic[2], Ana Klimovic[2], Andreas Krause[2], Caglar Gulcehre[1], David Rosenthal[10], Elliott Ash[2], Florian Tramèr[2], Joost VandeVondele[3], Livio Veraldi[10], Martin Rajman[1], Thomas Schulthess[3], Torsten Hoefler[2]

Leads: Antoine Bosselut[1], Martin Jaggi[1], Imanol Schlag[2]

Responsible Data Practices

- **Public & versioned data only**, primarily sourced from CommonCrawl
- Respect **AI opt-out preferences** of authors (removes 8% English, 4% non-English data)
- **Prevent model memorisation** of data, specifically copyrighted content
- **Quality:** Remove PII, toxic & harmful content, and low quality data (e.g. spam, non-human text, etc.).

	Apertus 8B / 70B	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	✅	❌	❌	❌	❌

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)
- Conversational in dozens of languages.

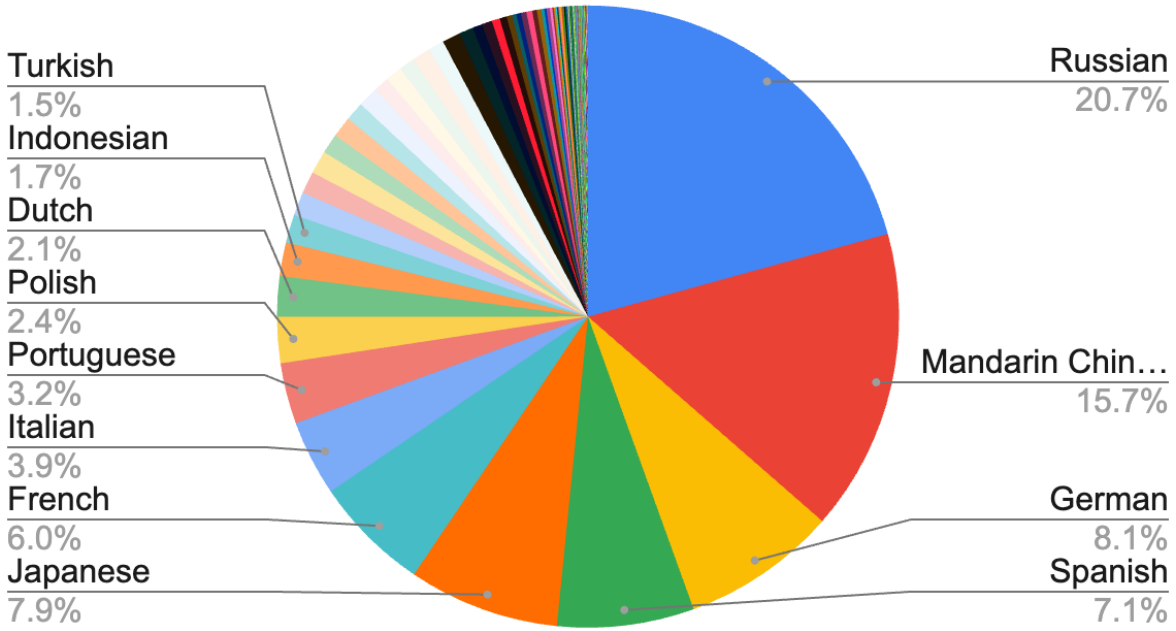

The Different Training Stages

Pretraining

Compute
Very High ●●●●●●●●●●

- Read enormous amounts of text from books and the web
- Pick up grammar, facts, and patterns of language
- Build a base of general knowledge

Example

The capital of France is →

Paris

Supervised Fine-Tuning (SFT)

Compute
Very Low ●

- Learn to follow instructions
- Practise giving helpful answers from human-written examples
- Move from raw text to a chat-like style

Example

How do I bake a chocolate cake? →

Recipe: 1) Preheat oven to 175°C...

RL from Verifiable Rewards (RLVR)

Compute
Medium ●●●●

- Practise tasks with clear right or wrong answers
- Try many attempts and keep what works
- Get better at reasoning, maths, and code

Example

Solve: 3x + 2 = 11 →

x = 13 | x = 3 ✓ | x = 5 | x = 9

COMING FOR APERTUS 1.5

Preference Training

Compute
Very Low ●

- Learn which answers humans prefer
- Become safer and more helpful
- Avoid harmful or misleading replies

Example

How do I deal with stress? →

Try meditation or exercise ✓

Just drink alcohol ✗

The Swiss AI Charter

- Apertus behaviour is trained **without direct human feedback**
- A diversity of strong open-weight **AI models judge completions** on the basis of the Charter
- The charter reflects our current view on how AI models ought to behave

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

Key Enabler: ALPS Research Infrastructure

- **ALPS:** Inaugurated in 2024, one of Europe's largest (**11,000 Nvidia GH200 GPUs**) efficient infrastructures
- **4096 GPUs for 3 months** for training Apertus 70B
- Built and maintained by a large team at CSCS, in close collaboration with HPE and NVIDIA
- **Full-fledged AI platform** now operational

Training at Scale Is Hard

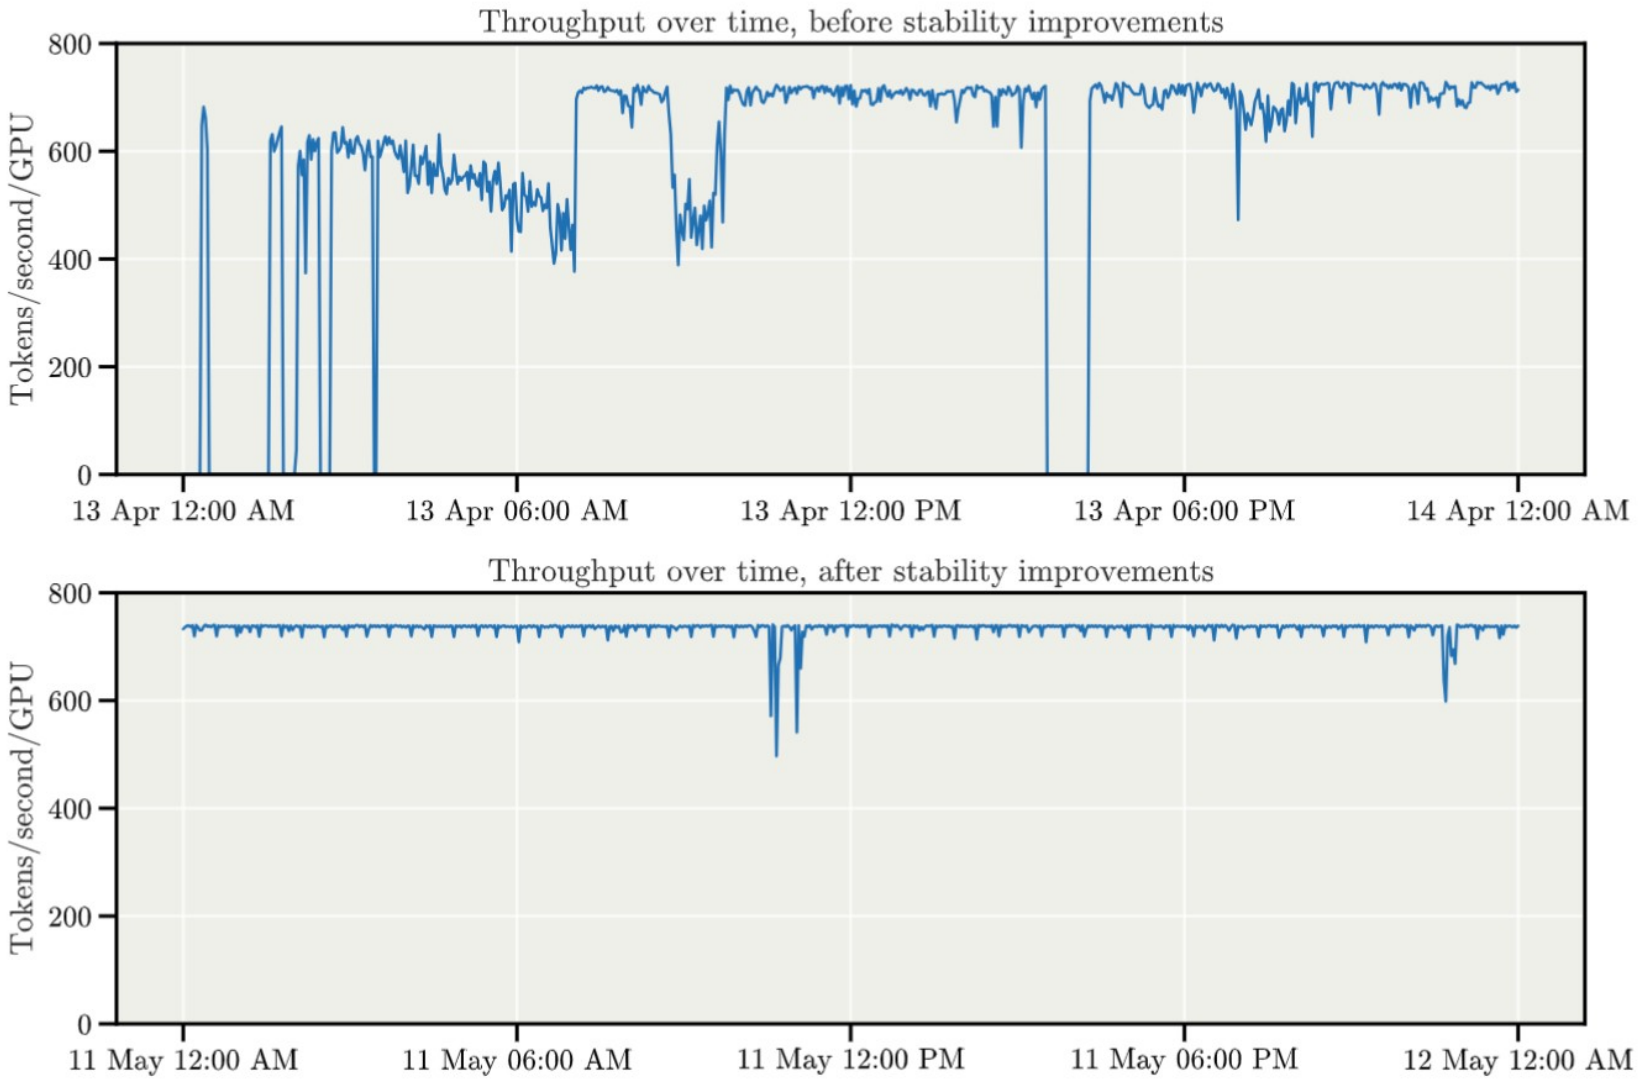

Where can you find Apertus ?

- **Weights available on HuggingFace:** For local deployment anywhere under an Apache 2.0 license (1.8M downloads EoY 2025; **>2.7M to date**)
- **Available from cloud providers:** Amazon AWS, Microsoft Azure, Exoscale, Swisscom, Infomaniak, Phoenix, ...
- **Supported in leading LLM software packages**: vLLM, SGlang, llama.cpp, MLX, Transformers
- **Hosted by third-party non-profit organisations:** e.g. http://publicai.co

Become a Swiss AI Industry Partner

Are you already building or plan to build on Apertus?

Let us know: llm-requests@swiss-ai.org

swisscom

SME CIRCLE

Join the community, leverage the freely usable frontier models like Apertus and benefit from resources and updates.

SNAI MEMBERSHIP

Access deeper AI expertise, events, and visibility across the entire Swiss AI ecosystem.

STRATEGIC PARTNERSHIPS

Contribute to research through long-term, multi-year collaboration with customized research wallet or donation plans.

Release Cycle

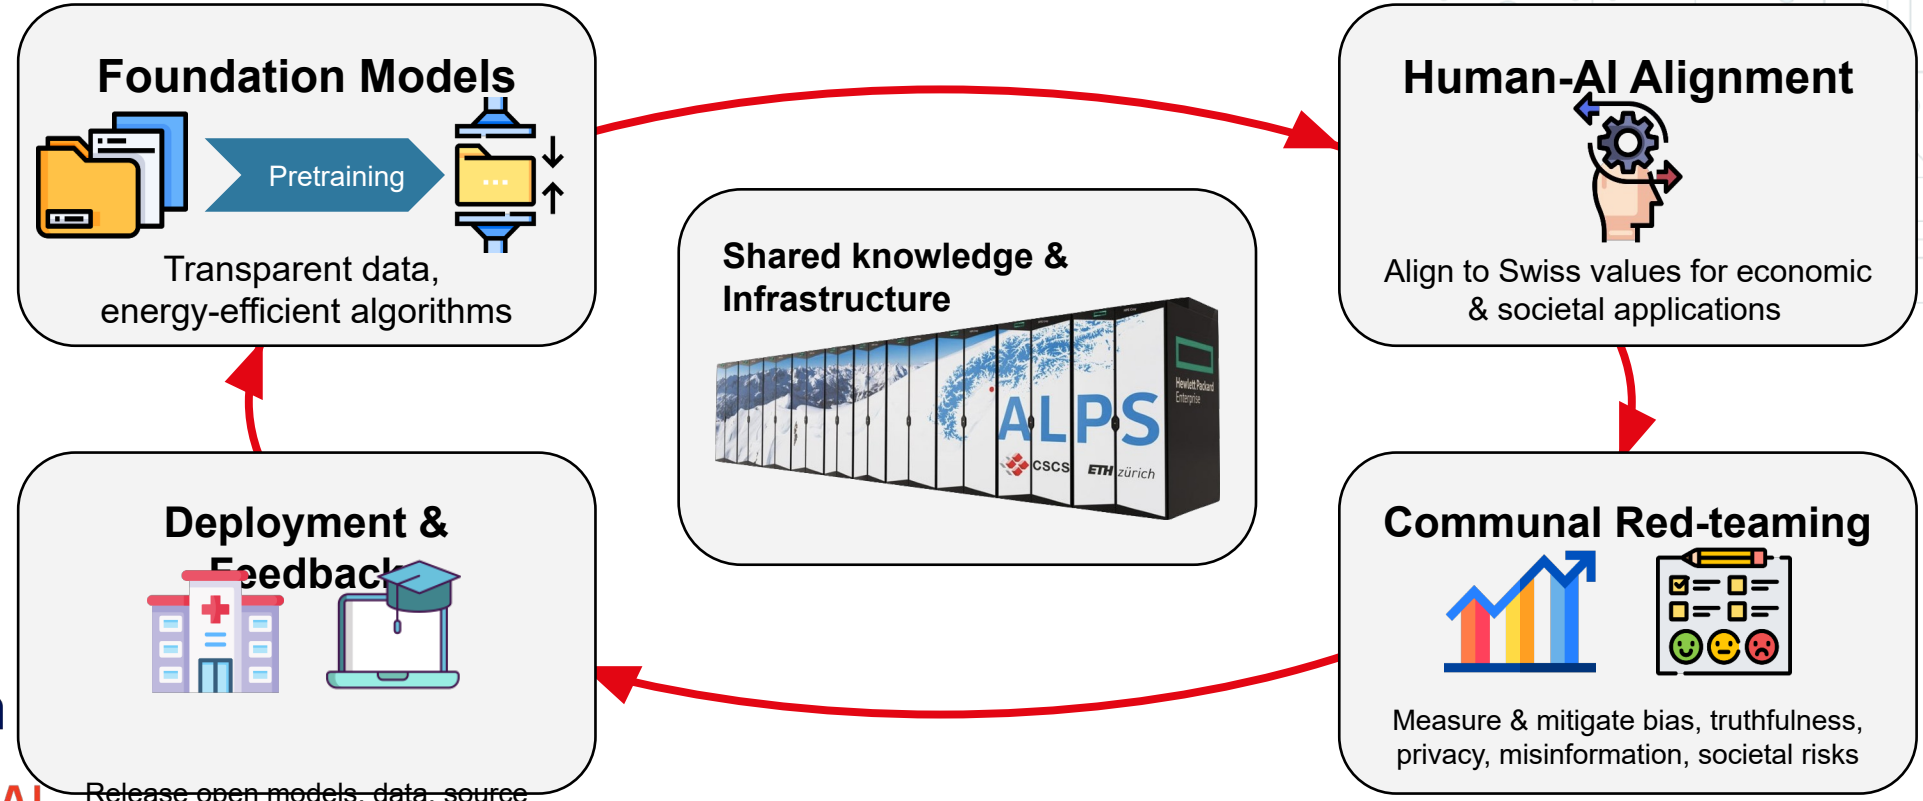

swisscom

Public AI
INFERENCE UTILITY

ti

infomaniak

SAFESWISS CLOUD

PHOENIQS

EXOSCALE

Soon: Apertus 1.5 (training now)

Continued Pretraining: Additional high-quality text data. Adding Trillions of tokens to previous training run

Multimodality: Enable native audio and image understanding

Focus Domains: Data upsampling for strategic areas such as health, education, and justice

Stronger Skills: Improved instruction-following, tool-use, and advanced reasoning capabilities through RL at scale

Later this year: Apertus 2.0

Scaling up: Reducing the gap with leading open weight models

Up to a magnitude larger model compared to Apertus 1.0

→ stronger performance

→ same deployment costs

Native Multimodality: Text, vision, speech, tabular, etc.

Distillation: Small models trained from logits of large models → 10x more efficient small model training

Next Expansion of Responsible Frontier:

- Maintain standards of transparency, responsibility, representation

AI Is Changing the World

Intelligence is becoming a commodity: abundant, cheap, everywhere

AI is becoming **critical infrastructure**, on par with power and networks

Outsourcing it concentrates two things in foreign hands:

- **Data**: whoever serves the model sees every query you send it
- **Behaviour**: whoever trains the model decides its values and refusals

We need to invest in the expertise and infrastructure to build, evaluate, and deploy these systems on our own terms.

Thank You!

Find more info on our Apertus website: apertus-ai.org

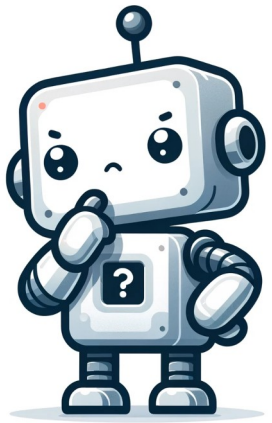

APERTVS

sign-up for our newsletter

Kontakt
Imanol Schlag, PhD
AI Research Scientist @ ETH AI Center
Apertus Co-Lead

ETH AI CENTER