

Open development, open models: transparent generative AI for Europe

Andrey Kutuzov
Language Technology Group

University of Oslo

11 February 2026



1 LLM-as-a-service

2 Open software and open LLMs

3 Aspects of openness

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

1) large language model (LLM) itself, doing all the work

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

1) large language model (LLM) itself, doing all the work and

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

- 1) large language model (LLM) itself, doing all the work and
- 2) point of access to the model(s): interface

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

- 1) large language model (LLM) itself, doing all the work and
 - 2) point of access to the model(s): interface
- What we are dealing with nowadays is usually a combination of the actual model and a user interface.

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

- 1) large language model (LLM) itself, doing all the work and
- 2) point of access to the model(s): interface
 - ▶ What we are dealing with nowadays is usually a combination of the actual model and a user interface.
 - ▶ This allows to 'serve' an LLM or LLMs as a service
 - ▶ good for business!

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

- 1) large language model (LLM) itself, doing all the work and
- 2) point of access to the model(s): interface
 - ▶ What we are dealing with nowadays is usually a combination of the actual model and a user interface.
 - ▶ This allows to 'serve' an LLM or LLMs as a service
 - ▶ good for business!
 - ▶ Not so good for public understanding of differences between models.
 - ▶ The model is becoming something very far from real life, almost mythical: unlike the inference interface.

For most end users, 'generative AI' means 'chatbot with a web interface or an app'.



This blurs an important distinction

- 1) large language model (LLM) itself, doing all the work and
- 2) point of access to the model(s): interface
 - ▶ What we are dealing with nowadays is usually a combination of the actual model and a user interface.
 - ▶ This allows to 'serve' an LLM or LLMs as a service
 - ▶ good for business!
 - ▶ Not so good for public understanding of differences between models.
 - ▶ The model is becoming something very far from real life, almost mythical: unlike the inference interface.

In fact, **models** are powering 'AI apps'. It is extremely important who owns them.

1 LLM-as-a-service

2 Open software and open LLMs

3 Aspects of openness

Open software and open LLMs



Let's recall **open source** / **free software** (not 'freeware'!).

Open software and open LLMs




Let's recall **open source** / **free software** (not 'freeware'!).

Four essential freedoms, according to Free Software Foundation

1. The freedom **to run the program** as you wish, for any purpose.
2. The freedom **to study how the program works, and change it to make it do what you wish.**
3. The freedom **to redistribute copies** so you can help others.
4. The freedom **to distribute copies of your modified versions** to others.

Most of modern computing infrastructure (including Internet) is built on **open source software**.

Openly licensed content powers foundational projects like Wikipedia.



Open software and open LLMs



Let's recall **open source** / **free software** (not 'freeware'!).


Four essential freedoms, according to Free Software Foundation

1. The freedom **to run the program** as you wish, for any purpose.
2. The freedom **to study how the program works, and change it to make it do what you wish.**
3. The freedom **to redistribute copies** so you can help others.
4. The freedom **to distribute copies of your modified versions** to others.

Most of modern computing infrastructure (including Internet) is built on **open source software**.

Openly licensed content powers foundational projects like Wikipedia.

What does 'open' and 'closed' mean for **large language models**?



- 1 LLM-as-a-service
- 2 Open software and open LLMs
- 3 Aspects of openness



Obviously closed models

- ▶ ChatGPT, Gemini, Claude, Grok ≥ 3 etc are entirely closed.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training



Obviously closed models

- ▶ ChatGPT, Gemini, Claude, Grok ≥ 3 etc are entirely closed.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training
 - ▶ Meaningful comparison and studying impossible.
 - ▶ E.g., is the GPT-5 model larger than GPT-4? One can only guess.



Obviously closed models

- ▶ ChatGPT, Gemini, Claude, Grok ≥ 3 etc are entirely closed.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training
 - ▶ Meaningful comparison and studying impossible.
 - ▶ E.g., is the GPT-5 model larger than GPT-4? One can only guess.
- ▶ Roughly equivalent to closed proprietary software like MS Windows...



Obviously closed models

- ▶ ChatGPT, Gemini, Claude, Grok ≥ 3 etc are entirely closed.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training
 - ▶ Meaningful comparison and studying impossible.
 - ▶ E.g., is the GPT-5 model larger than GPT-4? One can only guess.
- ▶ Roughly equivalent to closed proprietary software like MS Windows...
 - ▶ ...or, rather, Office 365.



Obviously closed models

- ▶ **ChatGPT, Gemini, Claude, Grok ≥ 3** etc are **entirely closed**.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training
 - ▶ **Meaningful comparison and studying impossible.**
 - ▶ E.g., is the GPT-5 model larger than GPT-4? One can only guess.
- ▶ Roughly equivalent to **closed proprietary software** like MS Windows...
 - ▶ ...or, rather, Office 365.

'Freeware' LLMs: you can have the model but not the data

Some actors advertise their models as 'open source':

- ▶ **Llama** by Meta [Touvron et al., 2023],
- ▶ **Deepseek** by HighFlyer [DeepSeek-AI, 2025]...



Obviously closed models

- ▶ **ChatGPT, Gemini, Claude, Grok ≥ 3** etc are **entirely closed**.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training
 - ▶ **Meaningful comparison and studying impossible.**
 - ▶ E.g., is the GPT-5 model larger than GPT-4? One can only guess.
- ▶ Roughly equivalent to **closed proprietary software** like MS Windows...
 - ▶ ...or, rather, Office 365.

'Freeware' LLMs: you can have the model but not the data

Some actors advertise their models as 'open source':

- ▶ **Llama** by Meta [Touvron et al., 2023],
- ▶ **Deepseek** by HighFlyer [DeepSeek-AI, 2025]...
- ▶ You can indeed download these models and self-host them. But...



Obviously closed models

- ▶ **ChatGPT, Gemini, Claude, Grok ≥ 3** etc are **entirely closed**.
 - ▶ No direct access to the models: only via API or web interfaces.
 - ▶ Training datasets unknown, no docs or peer-reviewed papers on training
 - ▶ **Meaningful comparison and studying impossible.**
 - ▶ E.g., is the GPT-5 model larger than GPT-4? One can only guess.
- ▶ Roughly equivalent to **closed proprietary software** like MS Windows...
 - ▶ ...or, rather, Office 365.

'Freeware' LLMs: you can have the model but not the data

Some actors advertise their models as 'open source':

- ▶ **Llama** by Meta [Touvron et al., 2023],
- ▶ **Deepseek** by HighFlyer [DeepSeek-AI, 2025]...
- ▶ You can indeed download these models and self-host them. But...

More correct to use the term '**open-weights models**'. For LLMs, **openness** is more complex.

‘Openness’ of an LLM is a multi-faceted property

- ▶ are **model weights** available? **intermediate checkpoints**?
- ▶ are **training data** available or reproducible?
- ▶ is the architecture and **training methodology** transparent?
- ▶ are **inference endpoints** easily available?
- ▶ what are **licences / terms of use**? [Liesenfeld et al., 2023]

‘Openness’ of an LLM is a multi-faceted property

- ▶ are **model weights** available? **intermediate checkpoints**?
- ▶ are **training data** available or reproducible?
- ▶ is the architecture and **training methodology** transparent?
- ▶ are **inference endpoints** easily available?
- ▶ what are **licences / terms of use**? [Liesenfeld et al., 2023]




<https://osai-index.eu/database>

Artificial Analysis Openness Index: Components

Openness Index underlying score contribution by components, up to a maximum of 18 (higher is more open)

■ Model Availability ■ Transparency - Methodology ■ Transparency - Post-training Data ■ Transparency - Pre-training Data



<https://artificialanalysis.ai/models/open-source>

Why we should prefer **more open LLMs**?

- ▶ They stimulate **progress and common good**:
 - ▶ reproducible and transparent research for scientists
 - ▶ 'everyone can download, use, study and improve' paradigm leads to a better world for everyone

Why we should prefer **more open LLMs**?

- ▶ They stimulate **progress and common good**:
 - ▶ reproducible and transparent research for scientists
 - ▶ 'everyone can download, use, study and improve' paradigm leads to a better world for everyone
- ▶ Imagine *a world without Wikipedia*
 - ▶ pay per-token for access to encyclopaedic knowledge

Why we should prefer **more open LLMs**?

- ▶ They stimulate **progress and common good**:
 - ▶ reproducible and transparent research for scientists
 - ▶ 'everyone can download, use, study and improve' paradigm leads to a better world for everyone
- ▶ Imagine *a world without Wikipedia*
 - ▶ pay per-token for access to encyclopaedic knowledge
- ▶ To compare, **closed models** owned by BigTech monopolies naturally increase **centralization** of the LLM ecosystem
 - ▶ ...thriving on open community contributions, but not giving back to the society
 - ▶ ...monetizing via surveillance-based advertising and pricing schemes (same as search engines and social networks 20 years ago)
 - ▶ ...capitalizing on vendor lock-ins



Joint Statement: Apple and Google have entered into a multi-year collaboration under which the next generation of Apple Foundation Models will be based on Google's Gemini models and cloud technology. These models will help power future Apple Intelligence features, including a more personalized Siri coming this year.

After careful evaluation, Apple determined that Google's AI technology provides the most capable foundation for Apple Foundation Models and is excited about the innovative new experiences it will unlock for Apple users. Apple Intelligence will continue to run on Apple devices and Private Cloud Compute, while maintaining Apple's industry-leading privacy standards.

Last edited 6:08 PM · Jan 12, 2026 · **10.9M** Views



1/ **Reasoning across sources**: Ever need to buy parts for your car but don't have the info handy?

Ask Gemini to "Recommend tires for my car."

By referencing connected apps like @Gmail and Photos, it can understand your car's make and model, and even the types of trips you take, to give recommendations of tires and info like your license plate number to make your visit to the auto shop go more smoothly.

openai.com

Testing ads in ChatGPT

5-6 minutes

Today, we're beginning to test ads in ChatGPT in the U.S. The test will be for logged-in adult users on the Free and Go subscription tiers. Plus, Pro, Business, Enterprise, and Education tiers will not have ads. Ads do not influence the answers ChatGPT gives you, and we keep your conversations with ChatGPT private from advertisers. Our goal is for ads to support broader access to more powerful ChatGPT features while maintaining the trust people place in ChatGPT for important and personal tasks. We're starting with a test to learn, listen, and make sure we get the experience right.


The community strikes back



- ▶ In the last 1-2 years, **open models** are starting to seriously compete with closed LLMs
 - ▶ Notable: **OLMo** family of models by Ai2 [Olmo et al., 2025]
- ▶ Even BigTech players open up some of their artifacts
 - ▶ OpenAI releases <https://hf.co/openai/gpt-oss-120b>

The community strikes back

- ▶ In the last 1-2 years, **open models** are starting to seriously compete with closed LLMs
 - ▶ Notable: **OLMo** family of models by Ai2 [Olmo et al., 2025]
- ▶ Even BigTech players open up some of their artifacts
 - ▶ OpenAI releases <https://hf.co/openai/gpt-oss-120b>
- ▶ **Chinese open models** are currently taking over the users from the **US labs**
- ▶ **Europe** is (slowly) catching up:
 - ▶ **HPLT** <https://hplt-project.org>
 - ▶ **TildeOpen** <https://hf.co/TildeAI/>
 - ▶ **Apertus/SwissAI** <https://publicai.co/>
 - ▶ **OpenEuroLLM** <https://hf.co/openeurolm>
 - ▶ **NorMistral** and **NorOLMo**
<https://chat.llm.sigma2.no/>
- ▶ ...still severely underspending compared to China and US.



Crossroads

- ▶ Generative AI (well, LLMs) is changing the society

Crossroads

- ▶ Generative AI (well, LLMs) is changing the society
- ▶ We are now at the crossroads:
 - ▶ diverse, decentralized and transparent ecosystem of **open LLMs for public good**, or
 - ▶ centralized and monopolized ecosystem of BigTech **LLM services for profit**: a dystopian future

Crossroads

- ▶ Generative AI (well, LLMs) is changing the society
- ▶ We are now at the crossroads:
 - ▶ diverse, decentralized and transparent ecosystem of **open LLMs for public good**, or
 - ▶ centralized and monopolized ecosystem of BigTech **LLM services for profit**: a dystopian future
- ▶ The real threat is not '*AI destroying humanity to optimize paperclips production*'

Crossroads

- ▶ Generative AI (well, LLMs) is changing the society
- ▶ We are now at the crossroads:
 - ▶ diverse, decentralized and transparent ecosystem of **open LLMs for public good**, or
 - ▶ centralized and monopolized ecosystem of BigTech **LLM services for profit**: a dystopian future
- ▶ The real threat is not '*AI destroying humanity to optimize paperclips production*'
- ▶ The real threat is the **concentration of technological power and knowledge in the hands of a few corporations, totally unaccountable to the people.**

Crossroads

- ▶ Generative AI (well, LLMs) is changing the society
- ▶ We are now at the crossroads:
 - ▶ diverse, decentralized and transparent ecosystem of **open LLMs for public good**, or
 - ▶ centralized and monopolized ecosystem of BigTech **LLM services for profit**: a dystopian future
- ▶ The real threat is not '*AI destroying humanity to optimize paperclips production*'
- ▶ The real threat is the **concentration of technological power and knowledge in the hands of a few corporations, totally unaccountable to the people.**
- ▶ Europe is now exploring LLM regulations (**EU AI Act**); this can be tricky

Crossroads

- ▶ Generative AI (well, LLMs) is changing the society
- ▶ We are now at the crossroads:
 - ▶ diverse, decentralized and transparent ecosystem of **open LLMs for public good**, or
 - ▶ centralized and monopolized ecosystem of BigTech **LLM services for profit**: a dystopian future
- ▶ The real threat is not '*AI destroying humanity to optimize paperclips production*'
- ▶ The real threat is the **concentration of technological power and knowledge in the hands of a few corporations, totally unaccountable to the people.**
- ▶ Europe is now exploring LLM regulations (**EU AI Act**); this can be tricky
 - ▶ e.g., it is easy to completely kill the emerging European LLMs by adopting **too restrictive copyright laws**
 - ▶ but if done right, the society at large can actually benefit from really **open and transparent AI**.

References I



DeepSeek-AI (2025).

Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning.



Liesenfeld, A., Lopez, A., and Dingemanse, M. (2023).

Opening up ChatGPT: Tracking openness, transparency, and accountability in instruction-tuned text generators.

In Proceedings of the 5th International Conference on Conversational User Interfaces, pages 1–6.



Olmo, T., Ettinger, A., Bertsch, A., Kuehl, B., Graham, D., Heineman, D., Groeneveld, D., Brahman, F., Timbers, F., Ivison, H., Morrison, J., Poznanski, J., Lo, K., Soldaini, L., Jordan, M., Chen, M., Noukhovitch, M., Lambert, N., Walsh, P., Dasigi, P., Berry, R., Malik, S., Shah, S., Geng, S., Arora, S., Gupta, S., Anderson, T., Xiao, T., Murray, T., Romero, T., Graf, V., Asai, A., Bhagia, A., Wettig, A., Liu, A., Rangapur, A., Anastasiades, C., Huang, C., Schwenk, D., Trivedi, H., Magnusson, I., Lochner, J., Liu, J., Miranda, L. J. V., Sap, M., Morgan, M., Schmitz, M., Guerquin, M., Wilson, M., Huff, R.,

Bras, R. L., Xin, R., Shao, R., Skjongsberg, S., Shen, S. Z., Li, S. S., Wilde, T., Pyatkin, V., Merrill, W., Chang, Y., Gu, Y., Zeng, Z., Sabharwal, A., Zettlemoyer, L., Koh, P. W., Farhadi, A., Smith, N. A., and Hajishirzi, H. (2025).

Olmo 3.



Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. (2023).

Llama: Open and efficient foundation language models.