

The centralization of AI infrastructures and the emerging alternatives



Hector Taler-Fraisse

Word count: 3280

December 15, 2025

Contents

Thesis Statement	2
Introduction	3
1 The AI Industry Ecosystem and its Monopoly	4
1.1 How AI Systems Work and What They Need to Be Built	4
1.2 Key Players and the Competitive Landscape	4
2 Risks Resulting from AI Monopoly and the decentralisation solution	7
2.1 Typologies of Risks	7
2.1.1 The Alignment Problem	7
2.1.2 AI is alienating	8
2.2 Why Decentralisation Reduces Risks	9
3 Bittensor as a Solution	11
3.1 Bittensor Phylosophy and vulgarized functioning	11
3.2 Bittensor's ability to addresses the risks	11
3.3 Limitations of Bittensor long term adoption and viability	12
4 Conclusion	14

Thesis Statement

Centralization of AI power introduces systemic risks that threaten society, but decentralized frameworks like Bittensor offer a viable solution.

Introduction

In November 2022, OpenAI shocked the Artificial Intelligence community by launching their first public chatbot based on the Large Language Model GPT-3.5. Today, their new model GPT-5 can perform tasks with substantial precision that would take hours to be done by a human. According to a large part of the most renowned scientists working on such technologies, their abilities will very likely continue to improve in an exponential fashion. Numerous initiatives are now trying to build world governance in order to handle large scale risks this revolution could cause.

The proportions this revolution has taken are now making the AI industry a major ethical, political and technological challenge. ChatGPT-like systems are now widely used to look for information about politics, health, religion or history. Furthermore, the responsibilities given to AI-based systems (AI-Agents) are forecasted to soon reach important roles, such as company management, war strategy, government policy or healthcare and judgment decisions. [1] This obviously places a lot of power in the hands of the key stakeholders in this industry.

However, the entry barriers to have an influence on this market are extremely high with the three main resources needed for a stakeholder to compete in the domain are : human talents, compute power, and data. This paradigm results in a problematic oligopole held by few technological giants in America and China. This essay aims to showcase promising paths towards decentralised AI infrastructures enabling the sovereignty of less powerful players, whether users or builders of such systems.

Chapter 1

The AI Industry Ecosystem and its Monopoly

1.1 How AI Systems Work and What They Need to Be Built

There are a wide variety of AI systems currently available : some aim to recognise patterns on images, others to detect emotions in a text. Yet, this essay will focus on systems able to generate content, whatever its modality (text, sound, images, or videos). Namely, generative AI models. The reason for this choice lies in the fact that generative AI systems are the ones from which the emerging property of intelligence arises. LLM (for Large Language Models) constitute a branch of the generative AI field whose systems demonstrate capabilities to generate text that make sense to a point that they seem to behave as intelligent systems. We will continuously refer to them throughout this report.

Generative AI models are built during a process called training. It roughly consists of providing an enormous amount of data to a code that will itself use a substantial amount of compute to perform calculations. The revolution of ChatGPT and its surprising performances enabled the discovery of the so-called “scaling law”, that is : with more data and more compute comes a more intelligent model. This finding results in extremely high entry barriers for any newcomer who would try to compete against the most capitalized organisations already benefiting from access to talented engineers, data and compute power. That is to say, Big tech companies.

1.2 Key Players and the Competitive Landscape

Big tech companies are framing the current AI landscape alongside a few specialized startups benefiting from exceptionally high capital investment. One could interpret these abnormal capitalisation as a speculative frenzy, what is sure is that there is evidence of

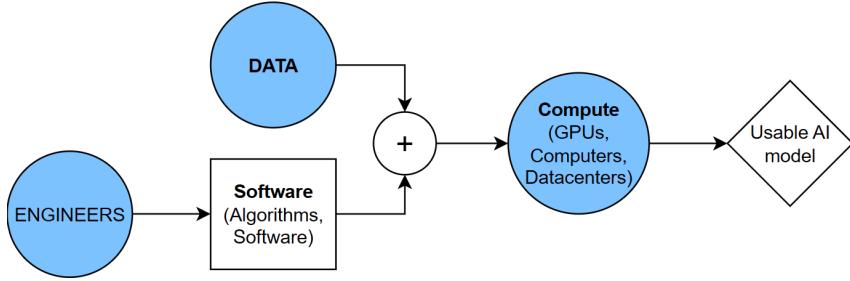


Figure 1.1: Elements needed for AI training

the so-called “Race to AGI” phenomenon (AGI for Artificial General Intelligence). That is, a competition to be the first to reach AI systems that would be more capable than any human on any task. Naturally, it has become a strategic goal for governments to control the power such technology represents. As a matter of fact, alongside large tech companies, geopolitics has now become the second key player in this industry by defining the access to the necessary resources previously listed.

Although Large Language Models are not representative of the whole generative AI Industry, the current consensus is that their architecture is the most likely to enable AGI. Thus, LLM rankings analysis enable a comprehensive identification of the most capable organisations of the industry. Looking at the benchmark of the best LLM models provided by Vellum [2], within the 47 most performant we have 14 models from OpenAI, 7 from Anthropic, 5 from Google, 2 from X.ai, one from Deepseek one from Qwen. However, the consensus on the current architecture of these models should be questioned. First, it is largely driven by players who need to raise funds and convince investors that the technology they’re using is the right one. Moreover, some leading figures remain deeply skeptical about current approaches. Yann LeCun, Meta’s Chief AI Scientist and one of the “godfathers” of AI explains that large language models will not lead to real computer intelligence. [3]

Beyond American GAFAMs and heavily funded startups AI specialists (OpenAI, Anthropic...), Chinese tech giants like Alibaba, Baidu, and Tencent, have emerged as substantial frontier model providers, alongside specialists like DeepSeek. These few firms are characterized by their heavily data-driven activities and their seamless integration with cloud infrastructure providing the compute resources.

As explained earlier, sovereignty of the AI industry now stands for a key national interest for competitors. This brings this competition to a geopolitical level, both on data and compute concerns. USA restricts exportation of Nvidia’s chips to China, Taiwan’s geopolitical position as the quasi only chip producer is a serious source of tension in Sino-American relations. We also observe significant political challenges surrounding digital platforms, which generate the data essential for training these AI systems. These include threats to ban TikTok in the United States, restrictions on American digital services

in China, and European regulations aimed at limiting the use of data collected from European users, despite their heavy dependence on these platforms. Governments have a decisive role to provide the capital requirements for AI models training projects. On this matter, disparities are enormous : the American company OpenAI has raised over \$57 billion, when Mistral French's biggest player managed to get funds up to \$3 billion.

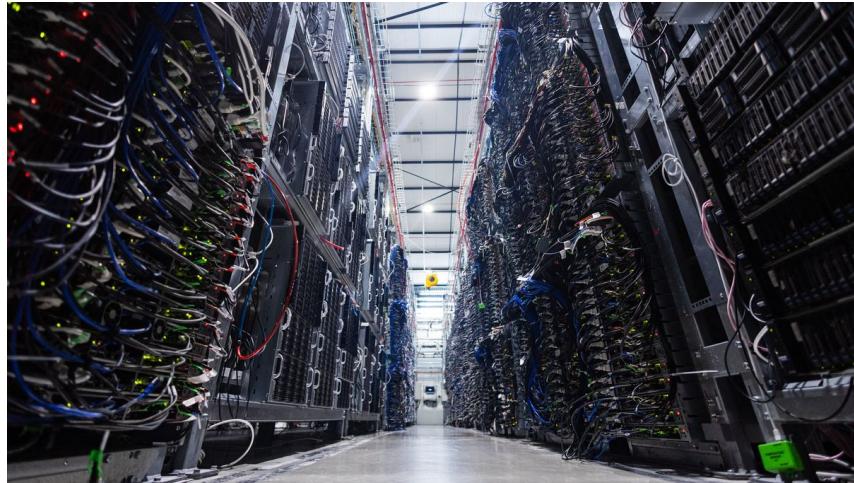


Figure 1.2: Example of a datacenter containing the hardware infrastructure needed to build and use large AI systems

Chapter 2

Risks Resulting from AI Monopoly and the decentralisation solution

2.1 Typologies of Risks

Now that the current business and technological ecosystems of AI have been clarified, let us finally give interest to what threat such systems could represent.

2.1.1 The Alignment Problem

A very old and fantasmatic theory predicts that sometime, the Human will create something that will be responsible for its own destruction. It is the well-known Frankenstein myth, also present in the movie (*For example Machina, 2001 Space Odyssey...*) Since the explosion of AI systems performances, numerous world class Artificial Intelligence experts got interested in studying how to ensure we can keep control on them. It is from this very crucial question that the scientific field of Artificial Intelligence Interpretability emerged. Yet, no answer has been found until now ; and this problem, still considered very tough, is expected to be solved in minimum 5 years and more likely in decades. Without a clear solution to this challenge, even though one can see the calculations these algorithms are making in order to look intelligent, we can not give clear meaning to each of them. In this sense, modern generative AI systems are still used as uncontrollable blackboxes (this characteristic notably makes them very good liars). This phenomenon, called the Alignment problem, makes the famous Frankenstein myth possible by two scenarios. A first one in which such an uncontrollable super powerful system is given to a malicious person or group of people, a second one in which the AI system turns against its user and creator. However, this is unfortunately no longer science fiction. Some serious collectives such as Pause AI are aiming to find an international consensus to regulate AI scientific progresses and investments and the open letter “Pause Giant AI Experiments” : An Open Letter” [4] calling AI labs to stop training their most powerful algorithms have been signed ex-

tremely serious scientific personalities (Yoshua Bengio, Stuart Russell, Elon Musk, Steve Wozniak, Yuval Noah Harari).

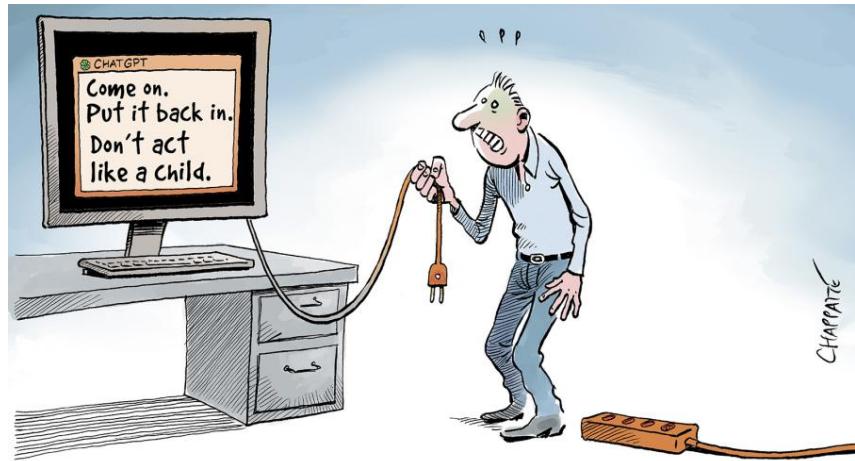


Figure 2.1: Illustration of Alignment problem : the AI system slips out of human control

2.1.2 AI is alienating

Beyond the existential risks we just discussed, let's analyze more tangible ways these types of systems threaten our society. Artificial Intelligence is not a mere new technology like all the ones that arose during the last century (Radio, phones, TV, internet ...), because this technology imitates the most human characteristic that can be : intellectual functions. This very basic statement enables it to drastically change paradigms in the domains that are the closest to what makes Human essence. AI is alienating. In this essay I will only focus on what I consider to be the domains where AI alienating power represents the most important threat to our civilisation. Democracy, Education and Work.

Although this writing focused mainly on Generative AI so far, I will have to consider another type of AI system here : Recommendation Algorithms. Today, the majority of the information is consumed through social media. In addition to that, most of the content consumed is not deliberately chosen. Despite this, it is selected by a system of recommendation based on Artificial Intelligence recommendation algorithms. On the one hand, likewise frontier Generative AI models, these algorithms are owned only by few mega Big Tech corporations (Google with Youtube, Meta with Facebook and Instagram, ByteDance with TikTok) ; and on the other hand, they are responsible for the majority of the worldwide information field. Because they are substantially relying on information, our democratic systems happen to be very sensitive to AI recommendation systems that can become an important threat to them. This weakness is widely used in communication strategies called astroturfing. These strategies are performed to influence public opinion and are typically funded by corporations and political entities (*For example : Cambridge Analytica scandals, Russian Influence on social media, propaganda and fake information*).

Another field heavily impacted by AI technologies is Education. It is becoming very tricky for teachers to assess students that have on demand access to these ChatGPT-like AI writing their essay, solving mathematic exercises or text analysis. Furthermore, alongside education at school, AI Chatbots have a great ability to convince young minds about something because of their highly anthropomorphic nature. It is obvious that, likewise Recommendation Algorithms, this situation makes AI Chatbots highly influential on people's education beyond schools : how to manage relationships, how to interpret human history, what basic principles should one follow...

Finally, there is work. Even though there has been a strong interest in integrating the promising Generative AI technologies in industry and business, most recent studies tend to show that a majority of the projects of this type in companies were a failure. In my opinion, such a failure is mostly due to a lack of digital systems readiness in companies. Indeed, in order to build powerful automation use cases thanks to Generative AI systems, a firm must dedicate substantial resources to work on their data and digital organisation, this represents a huge cost and is currently achievable by : either very small recent companies whose digital systems are built for AI first, either by very large companies able to invest enough money in such projects (Amazon deleting 15,000 positions replaced by AI systems). Another reason for the difficulties faced integrating Generative AI successfully in corporation systems is their lack of performance on some special tasks only expert human workers could reach so far. To counter this, the biggest AI Labs are following the new post-training paradigm. This involves, among other things, gathering top class experts on a specific subject whose purpose is to train an LLM model how to become as expert as they are. This mechanism illustrates perfectly the fighting tendency to centralize human knowledge and most frontier capabilities into proprietary systems such as OpenAI, Google or Anthropic algorithms.

2.2 Why Decentralisation Reduces Risks

Within this frame, transparency and decentralisation of the ecosystem enabling these threats become a vital challenge. As explained in Chapter One, building Artificial Intelligence systems broadly requires data, computing power, and skilled engineers. Therefore, we will examine how providing decentralised solutions for these three resources can lower the aforementioned risks.

Regarding compute power, decentralised networks allow anyone to contribute to the training of AI models by providing its computer or servers resources, this reduces reliance on Big Tech companies and allows for instance to build models without censorship or hidden bias. Nevertheless, it is a double edge sword. Since such frameworks enable smaller actors to train their own LLMs it also makes it easier for malicious entities to use this technology as a weapon.

When it comes to the talents able to design frontier software taking the most out of AI, the current ecosystem acts in favor of the few highly-capitalised companies that are able to attract world-class talents by providing unprecedented high salaries (for instance, Meta Platforms reportedly offered “seven- to nine-figure” compensation packages to top AI researchers in 2025). A decentralised network similar to Bittensor whose functioning will be explained in details in Chapter Three can enable incentivized development of software technologies through rewards given to people or groups proposing the most efficient code. In this framework, a skilled engineer no longer has to work in a big company to earn money thanks to its talent.

If a system has a substantial influence over our democratic structures, it must be as transparent as possible in order to prevent it from compromising them. Offering clear data about how previously studied recommendation systems behave and shape the information space would drastically reduce all the threatening phenomena highlighted previously.

The third chapter of this essay examines a case study of an ambitious project that proposes a decentralised framework, inspired by the famous Bitcoin blockchain but adapted to address the specific challenges of digital and AI decentralisation. This project is called Bittensor.

Chapter 3

Bittensor as a Solution

3.1 Bittensor Phylosophy and vulgarized functioning

Bittensor is a network within which anyone can deploy a digital service (be it digital storage, streaming, compute...). Each service deployed on the Bittensor network is called a “subnet”. Different personas are involved in the subnet inner mechanism (refer to the diagram below). First, there is the “owner” of a subnet that, following the Bittensor protocol, formalises the challenge its subnet intends to address. Then, we have the “minor’s”. Minors are providing their own solution to answer the subnet workload. “Validators” offer a system to rate the minors work. And finally, the “stakers” are the stakeholders investing their money in the subnet, trusting that it will be worthy of the confidence they place in it. This capital is then used to incentivize the different stakeholders of the subnet in proportion to their performance.

In order to better illustrate the workflow here, let us take a very simple subnet example here. The Zeus subnet [5] aims to provide the best weather forecast possible. Zeus subnet’s minors are later rewarded according to a grade rating how close their forecasts were to reality. These grades are determined by validators of the subnet who are as well rewarded according to the correctness of their rating (determined by their difference with the mean of all the validators rating). Stakers inject the liquidity distributed with respect to this mechanism in the subnet. Later, if Zeus happens to convince more investors that it is useful and efficient, the value of the former stakers’ investment will grow.

3.2 Bittensor’s ability to addresses the risks

The transparency of the Bittensor’s subnets activity is ensured by its use of a blockchain mechanism (similarly to Bitcoin’s). The latter enables the formerly explained reward mechanism to be done using the cryptocurrency called TAO. As a result, everyone can see who contributed to a subnet and how. “Data Universe” perfectly illustrates the

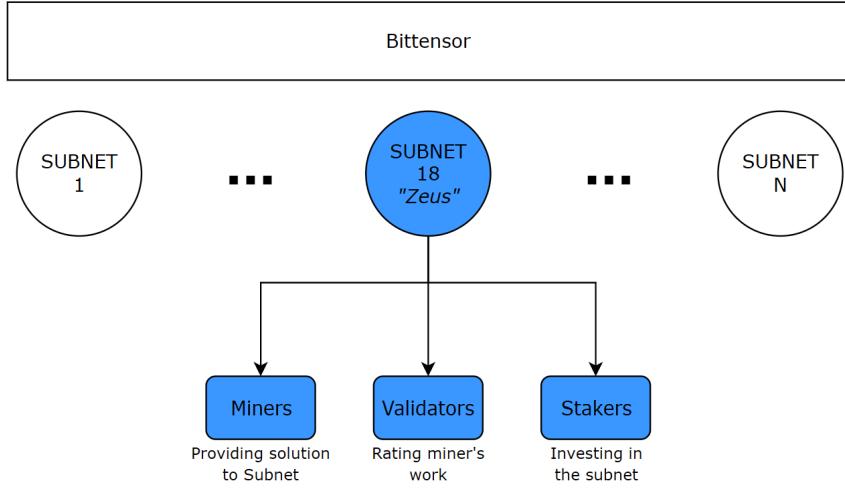


Figure 3.1: Bittensor subnets

importance of this point. This subnet collects and stores data being posted on social media in real time [6]. This data is mainly used to offer social insights fueling market analysis, political strategies, or AI chatbot context. Now, imagine an AI chatbot relying on this data to answer questions about people's opinions on political issues. It would clearly be of critical interest in a democratic system to be able to audit how the data is collected and managed. Well, Bittensor subnets like Data Universe make this possible in the most elegant way.

In addition to its transparency, Bittensor is by design decentralised. This is an advantage both on the user-side and builder-side of AI infrastructures. On the user-side, Bittensor subnets such as "Chutes" allow a user to become independent of the poor diversity of efficient AI-generative systems currently available on the market. On the builder-side, Bittensor is lowering entry barriers to contribute to AI-systems which therefore is fueling decentralisation. Indeed, some subnets allow engineers to propose solutions for which they can be gratefully paid. "Ridges" subnet illustrates this point particularly well: it lets programmers propose their solution for specialized AI systems writing code. One could argue that big corporations will always be in the best position to attract top-tier engineers. However, with hundreds of engineers competing to optimize their agentic workflows, the top miner are earning an estimated \$62,000 per day, demonstrating how decentralized competition can accelerate AI development. [7].

3.3 Limitations of Bittensor long term adoption and viability

One could argue that the decentralized computer architecture could never compete against the supercalculator and datacenters owned by a company such as Google. Nevertheless,

some evidence demonstrates this can be wrong with very popular networks. Indeed, Bitcoin's decentralized network is estimated to consume around 100 TWh/year (some sources report up to 240 TWh/year consumption) when the entire Google's infrastructure consumes 24 TWh/year according to their environmental report. [8]. With this in mind and the necessary decentralisation of intelligence Bittensor represents, this first argument does not undermine the viability of such a network to establish itself against the tech giants.

Now the ability for the companies that already have an oligopoly to take control of the majority of such a decentralised network should also be considered. In fact, in this case the decentralisation benefits would completely vanish. For example, if Amazon uses their massive cloud infrastructures they could become a massive minor on a subnet and dominate the competition. In reality, although sparks of centralisation cases have been reported, they are always limited to one subnet only and the resources needed to take control of a substantial part of the Bittensor network (on multiple subnets) remains very improbable.[9]

Chapter 4

Conclusion

The centralization of AI infrastructure concentrates unprecedented technological and societal power in the hands of a few corporations and governments, creating systemic risks that threaten democracy, education, and broader social issues. As shown throughout this essay, frontier generative AI models and recommendation algorithms amplify these risks by enabling opaque decision-making, monopolizing talent, and controlling the flow of critical information. Decentralized frameworks like Bittensor do not replace these centralized systems, nor do they eliminate all the risks. Nevertheless, they offer a credible alternative by redistributing power, incentivizing open participation, and providing transparency thanks to the blockchain backbone it is built on.

Bittensor's subnets illustrate how decentralized networks can enable innovation, lower entry barriers, and allow society to audit and influence AI development in ways that centralized monopolies cannot. Yet, some challenges remain. From potential dominance by large players on individual subnets to technological adoption barriers for small stakeholders. Acknowledging these limitations, the future of AI governance must combine decentralization with collaboration, and continuous experimentation. In this perspective, the current version of Bittensor may not be the unique solution, but it is definitely a model for the types of frameworks society needs.

Bibliography

- [1] AI-2027. *AI-2027*. <https://ai-2027.com/>.
- [2] Vellum. LLM Leaderboard. https://www.vellum.ai/llm-leaderboard?utm_source=perplexity&utm_medium=geo.
- [3] YouTube. (*Video*). <https://youtu.be/ETZfkkv6V7Y>.
- [4] Future of Life Institute. Pause Giant AI Experiments: An Open Letter. Wikipedia, https://en.wikipedia.org/wiki/Pause_Giant_AI_Experiments:_An_Open_Letter. :contentReference[oaicite:0]index=0
- [5] SubnetAlpha. Zeus Subnet. <https://subnetalpha.ai/subnet/zeus/>.
- [6] TensorPlex Labs. Bittensor Subnet 13 Data Universe — decentralised data scraping. Medium, <https://medium.com/@tensorplexlabs/bittensor-subnet-13-data-universe-decentralised-data-scraping-3787abfe2ae0>.
- [7] TensorPlex Labs. The Bittensor Subnet Dynasty: Decoding the Future of Decentralized AI Medium, <https://medium.com/@Defortuneteller/the-bittensor-subnet-dynasty-decoding-the-future-of-decentralized-ai-3306b8275a03>.
- [8] Google. 2024 Environmental Report. <https://blog.google/outreach-initiatives/sustainability/2024-environmental-report>.
- [9] Lui, Elizabeth, and Jiahao Sun. “Bittensor Protocol: The Bitcoin in Decentralized Artificial Intelligence? A Critical and Empirical Analysis.” *Preprint*, June 2025. DOI: 10.48550/arXiv.2507.02951. :contentReference[oaicite:1]index=1 Available via ResearchGate: https://www.researchgate.net/publication/393477381_Bittensor_Protocol_The_Bitcoin_in_Decentralized_Artificial_Intelligence_A_Critical_and_Empirical_Analysis