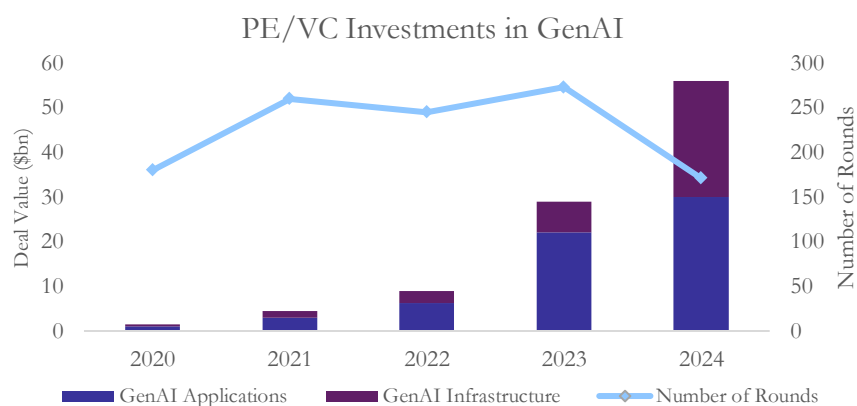


# Decoding Generative AI: Beyond the Hype

## Introduction

Generative AI has gone mainstream at a blistering pace worldwide. OpenAI's ChatGPT, launched in late 2022, reached 100m users in just two months and was valued at around \$29bn by January 2023. Now, on the verge of finalizing a new funding round, its valuation is expected to surge nearly tenfold to an estimated \$300bn. This remarkable growth has ignited a global AI investment wave, with total startup funding rising from \$29bn in 2023 to almost \$56bn in 2024. The viral uptake is mirrored in enterprise adoption: 72% of business decision-makers now report using generative AI at least weekly—a significant jump from 37% in 2023. Adoption, once limited to tech teams, is now spreading company-wide, with departments like Marketing, Operations, HR, and even Legal experimenting with AI for tasks such as drafting contracts. The initial awe and wariness are fading as executives become increasingly optimistic about AI's impact on the workplace.

Financially speaking, generative AI is evolving from a small niche into a colossal market. Industry analysts estimate the market's current size at roughly \$25.9bn in annual revenue (2024), with forecasts predicting growth to about \$174bn by 2028. Bloomberg analysts expect the broader GenAI (including hardware) market to reach around \$1.3tn by 2032. For context, McKinsey researchers estimate that generative AI could ultimately deliver \$2.6–4.4tn in annual economic value across industries—comparable to the productivity gains of an entire G7 economy. In short, if these projections hold, today's nascent market could become a cornerstone of the global economy within a decade.



Source: S&P Global

Investor enthusiasm is fueling this boom. GenAI's private equity and venture capital funding nearly doubled from about \$29bn in 2023 to \$56bn in 2024. Notably, this surge in capital came even as the number of funded startups fell, indicating that investors are concentrating on a smaller set of winners. The average funding round skyrocketed to around \$407m in 2024 from roughly \$130m the year before—a sign that big players are commanding mega-deals. Late 2024 saw enormous raises, including approximately \$10bn for Databricks, \$6.6bn for OpenAI, \$4bn for Anthropic, and \$6bn for Elon Musk's new AI venture, xAI. Moreover, about \$26bn of 2024's funding went into the infrastructure layer (e.g., cloud GPU providers and AI developer platforms), underscoring the critical demand for hardware and middleware to support AI at scale.

## What's Driving the Boom?

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

First, the technology itself has rapidly improved. Breakthrough models like OpenAI's GPT-4, Anthropic's Claude-3, and Google's Gemini-2 are far more capable than their predecessors—handling more complex tasks, larger contexts, and multiple modalities. For instance, in 2023, Anthropic scaled Claude's context window from roughly 9K tokens to 100K tokens in just two months. Today, standard models can handle 200K tokens, with Claude Enterprise supporting up to 500K tokens. These leaps have paved the way for new opportunities and intensified competition. Over the past year, many models have vied for market attention, including Meta's open-source Llama 3, Google's PaLM-2/Gemini-2, and emerging Chinese contenders like OpenSeek and Alibaba.

Second, access to AI has broadened dramatically. Thanks to cloud APIs and open-source releases, companies of all sizes can now experiment with top-tier AI without having to build everything from scratch. This increased accessibility—through AI-as-a-service offerings from Microsoft, Google, Amazon, and community-driven models on GitHub—has lowered barriers and accelerated adoption.

Third, use cases are multiplying as firms discover real business value. Generative AI is now deployed for marketing copy generation, customer service chatbots, coding assistants, HR recruiting tools, R&D modeling, and beyond. McKinsey finds that about 75% of gen AI's potential value comes from applications in customer operations, marketing & sales, software engineering, and R&D alone, although promising pilots are underway in virtually every sector. Early successes are eroding skepticism and creating a fear of missing out. Many companies have formed dedicated AI teams or Centers of Excellence, with roles like Chief AI Officer (CAIO) emerging as key positions. Leadership is actively “rewiring” organizations to capture AI's value—redesigning workflows and training staff—which several studies suggest can boost productivity by 20–30%.

Yet, the initial euphoria is giving way to a more pragmatic approach. While about three in four companies plan to increase AI investment next year, many executives now emphasize measurable ROI and risk management. There is growing recognition that challenges like accuracy/bias, data privacy, and integration must be addressed for AI to deliver sustained impact. In short, after an explosive start, generative AI is entering a “prove-it” phase in business—similar to the early days of the iPhone when expectations were recalibrated. The consensus is clear: the AI revolution is real, but it is not an overnight event; companies are moving fast yet staying patient to ensure long-term success.

## Major Players and Models

Many firms have pioneered the generative AI sector due to its economic and societal potential. OpenAI, the most notorious, began developing its closed-source AI algorithm in 2015 and introduced its Generative Pre-trained Transformer (GPT) in 2018. In November 2022, ChatGPT was released, quickly positioning itself as the market leader. However, OpenAI's reliance on vast datasets has landed it in legal challenges with authors and publishers—such as a pending case brought by the New York Times—centered on alleged copyright violations. If OpenAI loses, fair use may no longer apply to the development of Large Language Models (LLMs), making training increasingly expensive and hampering performance.

From a financial perspective, optimism about OpenAI abounds. Microsoft [NASDAQ: MSFT] owns 49% of OpenAI, having invested around \$13bn over the last five years. Additional investments, such as a recent \$6.6bn funding round led by Thrive Capital (with a potential extra \$1bn contingent on certain criteria), have provided the capital needed to scale operations. SoftBank's pending \$40bn investment, which values OpenAI at \$260bn, is expected to fund infrastructure projects like Stargate—a \$500bn AI infrastructure initiative in the US led by a joint venture of OpenAI, SoftBank, Oracle, and MGX. Although still unprofitable (losing \$5bn on \$3.7bn of revenue in 2024), OpenAI is well positioned to grow.

---

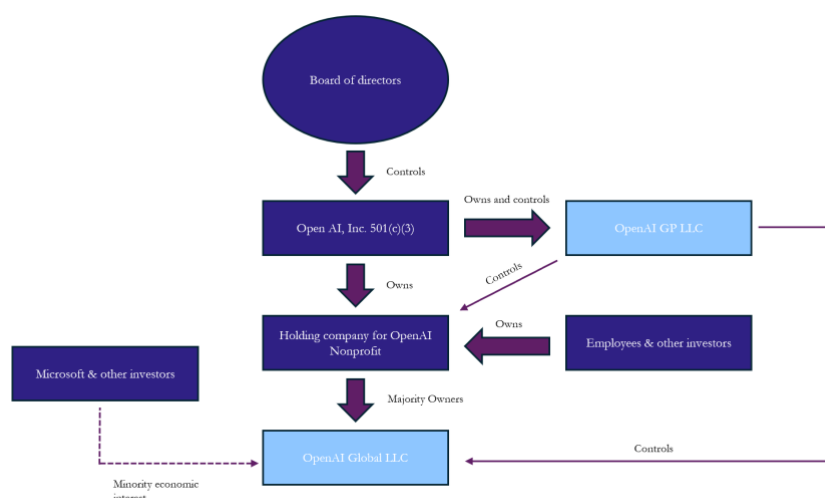
*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

OpenAI benefited from a significant head start. However, that gap is closing as well-funded technology giants invest heavily in AI. Google's [NASDAQ: GOOG] venture into AI began in 2001 with its search engine, and in 2023, Google DeepMind (a subsidiary of Google) launched Bard—later rebranded as Gemini. This closed-source AI, integrated into Google products like its search engine, leverages Alphabet's vast platforms (YouTube, Android, etc.) for rapid training and direct integrations. Despite risks such as potential FTC actions against Alphabet, DeepMind is on track to be a strong challenger.

Google has diversified its AI bets by acquiring a 14% stake in AI startup Anthropic for roughly \$3bn. Anthropic released its closed-source model Claude in March 2023 and has attracted an \$8bn investment from Amazon to integrate its technology within AWS (Amazon Web Services). Anthropic differentiates itself through its emphasis on constitutional AI—training models based on principles of safety and ethics, which enhances transparency. Other notable players include Grok (led by xAI/Elon Musk), Deepseek (owned by Liang Wenfang), Meta's Llama, and Mistral AI, with some using open-source approaches to challenge the current landscape.

## Funding Structures & Investment Trends

OpenAI's mission is to build safe AI for humanity, and its non-profit character reflects these values. Initially founded as a 501(c)(3) organization with public donations totaling \$130m, OpenAI soon realized that donations were insufficient for the capital-intensive nature of AI development. In 2019, OpenAI introduced a capped-profit structure, creating a for-profit subsidiary under the oversight of its non-profit board. This structure allows profits for investors to be capped at 100× the initial investment, ensuring the mission remains paramount. Shortly thereafter, Microsoft invested \$14bn for a 49% stake in the subsidiary—albeit with no voting rights—to integrate third-party AI models into products like Microsoft Office 365 Copilot.



Source: BSIC, OpenAI

In contrast, Anthropic operates as a Public Benefit Company, supervised by a Long-Term Benefit Trust (LTBT) managed by five trustees who act in the best interests of humanity. The LTBT holds a special class of shares with significant voting power, allowing it to influence board composition. In 2023, Alphabet invested over \$3bn in Anthropic, acquiring a 14% stake without board representation, while Amazon's \$8bn investment provides access to Anthropic's technology via AWS.

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

Furthermore, the shift towards mega-deals—with funding rounds averaging around \$407m—signals that investors are favoring a concentrated set of winners. This focus not only reflects the capital-intensive nature of AI development but also suggests that market consolidation may accelerate, potentially creating higher barriers to entry for new players.

## **Big Tech Investments & Integration**

Big tech companies are further integrating AI into their core operations. Google's early investment in DeepMind (back in 2014) set the stage for its aggressive push into AI. DeepMind's evolution—merging with Google Brain and forming Google DeepMind—demonstrates the drive to lead in AI integration across Google's services. Similarly, Meta uses AI both to generate content and to engage users, though this raises concerns about coordinated misinformation, prompting the need for robust security measures.

Elon Musk's xAI and various venture-backed startups are also emerging to challenge the established players. Funding rounds have been notable, with projects like Stargate—a joint venture led by SoftBank and backed by the Trump administration—aiming to build independent AI infrastructure and reduce reliance on partners like Microsoft. Meanwhile, startups such as Mistral AI have attracted investment from firms like Andreessen Horowitz (a16z) and Nvidia, though some face challenges from increasingly efficient competitors like the Chinese model Deepseek.

## **Regulatory Environment & Policy Trends**

Governments around the world are beginning to set the rules that will shape the future of AI. The European Union's AI Act, for example, categorizes AI applications into risk tiers—from minimal to high risk—requiring strict oversight for systems deemed high risk. In contrast, the United States has adopted a more fragmented approach with executive orders and state-level initiatives, leaving much of the regulation to existing agencies and voluntary guidelines. Meanwhile, China enforces a centralized and stringent regulatory framework that ties AI development closely to state priorities, requiring providers to adhere to strict censorship and security measures. These diverse regulatory environments will influence innovation, market consolidation, and the competitive dynamics among global AI players.

## **Open-Source vs. Closed-Source Approaches**

A vital facet in comparing various AI models is whether they are open-sourced or closed-sourced. Closed-source models keep the underlying code and weights proprietary, meaning external developers aren't allowed access without approval. This reduces the risk of the model being manipulated in ways that conflict with the creator's intent. It also helps maintain quality assurance and protects the model's competitive edge. However, this control comes at the cost of transparency, customizability, and broader development.

AI is inherently controversial due to ethical concerns around bias and discrimination. Since AI is trained on data, the consequences can be severe if that data—or its interpretation—is biased. For instance, an AI used in financial risk modeling might wrongly flag a particular race as less creditworthy. In a closed-source model, the code and training data are hidden, making it extremely difficult to audit for such equity issues.

Moreover, closed-source models limit customization, even under licensing agreements, reducing their effectiveness for specific consumer or developer needs. This lack of flexibility can stall external innovation and delay breakthroughs from the broader developer community. Revenue models for closed-source companies vary:

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

OpenAI and Anthropic offer paid subscriptions and tiered API access, while Google focuses on integrating Gemini into its existing suite of products. Ultimately, closed-source models allow firms to retain tighter control over development and monetization.

In contrast, open-source approaches prioritize transparency. External developers can access source code, model weights, and sometimes even training data, allowing for greater scrutiny and faster iteration. This openness promotes innovation and customization, enabling developers to adapt models to specific use cases, often resulting in superior outcomes.

However, this transparency also opens the door to misuse. Malicious actors could manipulate models for harmful purposes, and quality control becomes harder when development is decentralized. Business models for open-source players typically rely on offering premium services, enterprise-level tools, or specialized applications built on top of freely available models.

Despite their risks, open-source models could ultimately close the gap—or even surpass—closed-source competitors. Their collaborative nature fuels innovation and can dramatically reduce costs. Governments may also lean toward open-source models due to their inherent transparency, especially as security, bias, and ethical concerns become more prominent in public discourse.

### **The Dual Engines Driving Innovation**

In the dynamic world of generative AI, breakthroughs rest on two crucial resources: massive amounts of data and enormous computing power. However, a closer look reveals an intricate interplay between data quality, user interaction, and compute that underpins the rapid evolution of AI.

Large models require high-quality, diverse datasets that capture the full depth of human knowledge. Yet, only about 5% of the world's data—the surface web—is easily accessible, while the remaining 95% lies behind paywalls or in private formats. As a result, AI developers must compete for a limited pool of clean, reliable data.

A significant emerging divide is between proprietary and public data. Tech giants like Google and Meta benefit from vast reservoirs of user-generated content—from search queries to social media interactions—that are simply out of reach for smaller players. In contrast, early AI labs like OpenAI relied heavily on public datasets such as Common Crawl, Wikipedia, and Reddit repositories. Now, as platforms like Reddit and Twitter begin charging for API access, obtaining quality data is becoming more complex and costly, potentially reshaping the competitive landscape.

Beyond static datasets, real-time user interaction data plays a critical role in refining AI models. Continuous feedback from millions of users—captured through techniques like reinforcement learning from human feedback (RLHF)—provides an invaluable, dynamic input. This first-mover advantage creates a formidable barrier for newcomers who lack the scale of engagement necessary to fine-tune their models effectively.

In addition, smaller startups are experimenting with synthetic data generation and federated learning as ways to supplement or collaborate on data collection. Although these approaches offer innovative alternatives, they also carry risks, such as introducing biases or errors that can propagate through the model over time.

If data is the fuel for generative AI, then compute is undeniably the engine. Training state-of-the-art models today requires vast amounts of computing power—often involving tens of thousands of GPUs running in parallel for weeks or months. For instance, GPT-4 was reportedly trained using around 25K Nvidia A100 GPUs, with the

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*

overall cost of such massive distributed computing reaching tens or even hundreds of millions of dollars due to both hardware expenses and energy demands.

Nvidia sits at the heart of this “GPU arms race,” commanding roughly 80% or more of the AI chip market as of 2025. Its CUDA software ecosystem has become the de facto standard for optimizing AI computations. While AMD, Intel, and several innovative startups are challenging Nvidia’s dominance, none have yet built an ecosystem as robust or garnered a developer base as extensive.

Moreover, cloud providers such as Microsoft, Google, and Amazon play a crucial role as AI gatekeepers. Their on-demand, scalable computing resources enable both startups and established companies to run AI models without the burden of massive upfront investments. Strategic partnerships—like Microsoft’s deep integration with OpenAI and AWS’s exclusive arrangement with Anthropic—demonstrate how these providers leverage their infrastructure to secure equity stakes and competitive advantage.

Together, these interconnected factors—data availability, user engagement, compute capacity, and cloud infrastructure—drive the relentless pace of innovation in generative AI. They not only define the technical and economic landscape but also create the feedback loops that continuously push the boundaries of what AI can achieve.

## Conclusion

Generative AI’s evolution from a nascent technology to an industry powerhouse is marked by rapid technological breakthroughs, unprecedented investments, and a shifting regulatory landscape. While challenges remain—from data scarcity and compute costs to ethical and legal concerns—the transformative potential of AI is undeniable. As established giants and nimble startups continue to innovate, the balance between innovation, risk management, and regulatory oversight will be crucial for the long-term success of the AI revolution.

TAGS: AI, OpenAI, Big Tech, Innovation, Regulation

---

*All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.*