

The Transformation of OpenAI

Written by Louis Athmann & Simone Colavecchia

Part I

“From Vision To Velocity: How OpenAI Transformed Its Mission”

Since the recently expanded partnership between Microsoft and OpenAI was announced in late 2025, one question loomed large: how did a research lab founded on idealism become the engine of one of the most valuable AI companies in the world? The answer lies in a story of ambition, innovation, cost escalation, and structural transformation, all anchored in a bold mission.

MISSION & ORIGINS

OpenAI was founded in December 2015 with the aim of advancing digital intelligence in a way likely to benefit all of humanity, unconstrained by financial-return motives (**OpenAI**). Its founders, among them Sam Altman, Elon Musk, Greg Brockman, and Ilya Sutskever, envisioned a world where AI would enhance human potential rather than undermine it (**Montevirgen**). They built the organisation around two central pillars: open research and risk mitigation. As the introductory post of OpenAI put it, they wanted AI to be “an extension of individual human wills and, in the spirit of liberty, as broadly and evenly distributed as possible” (**OpenAI**). The non-profit status of the organisation reinforced the idea that profit should never eclipse public good.

From the outset, OpenAI posed a radical question: Can you build the most powerful systems in existence and still ensure they serve humanity, not just shareholders?

THE PRODUCT PATH:

Gpt-2 To Gpt-5 (And Beyond)

That question became more urgent as OpenAI’s models advanced. In 2019, the lab announced GPT-2, a natural-language model whose capacity to generate coherent text raised alarms about potential misuse (**OpenAI**). OpenAI, though, famously delayed its full release over safety concerns, without further study of potential misuse (**OpenAI**). Co-founder Jack Clark explained:

“Due to our concerns about malicious applications of the technology, we are not releasing the trained model,” reflecting the lab’s early emphasis on responsibility and transparency (**OpenAI**). Instead, they released a much smaller model “for researchers to experiment with, as well as a technical paper” (**OpenAI**).

The next leap came with GPT-3 in 2020, which opened the door to large-scale commercial applications: natural-language generation, translation, and even code synthesis. However, for much of this period, generative AI and OpenAI’s language models were primarily known within the tech community and remained largely out of public view. This limited visibility was intentional. In 2019, OpenAI announced that it would take a cautious and staged approach to releasing powerful models, citing the need to study potential risks and misuse before full deployment (**OpenAI**). The organisation maintained this strategy through GPT-3, offering access only via a restricted API to balance innovation with safety oversight. Similarly, in its 2018 Charter, OpenAI wrote: “We commit to freely share our research and data to the extent that we can do so safely and securely,” while simultaneously acknowledging that “safety and security concerns will reduce our traditional publishing in the short run” (**OpenAI**).

Yet this cautious obscurity did not last long. With the release of ChatGPT in 2022, generative AI was thrust into the mainstream, gaining traction on social media, particularly on TikTok, and finding its way into schools and universities, triggering a wave of innovation and competition (**Cefa**). I still remember the first time I encountered ChatGPT. A friend leaned over during a German class in late 2022 and typed a few prompts into the screen, and within seconds, paragraphs appeared: perfect translations, crisp summaries, even essays that could have passed for a student’s own. The precision and fluency were astonishing - it felt like witnessing a shift in what technology could do. At the time, few imagined it could advance much further. Yet only a few years later, it has become deeply woven into everyday life, shaping how we learn, communicate, and even think. Many students even admit they could hardly work without ChatGPT and would not know what to do without it (**Ettinghausen, The Guardian**). But this dependence carries risks. As educators have warned, overreliance on AI tools can keep students in their comfort zone, discouraging genuine skill development and critical thinking (**Schlott**).

Alongside this cultural shift came deeper questions about safety and responsibility. Reports emerged of users turning to ChatGPT in emotionally vulnerable moments and receiving inappropriate or harmful replies, sparking public debate about the psychological and ethical boundaries of AI interaction (**BBC**). At the same time, the benefits were undeniable: the technology revolutionised accessibility tools, research assistance, and creative collaboration, embedding itself in how millions of people interact with information (**Washington Post**).

The release of GPT-4 in March 2023 marked another critical milestone. The model demonstrated multimodal capabilities, accepting both text and image inputs, and showed remarkable

improvement in reasoning and factual accuracy (**OpenAI**). GPT-4 also underpinned ChatGPT Plus, signalling OpenAI's shift toward premium, subscription-based access. With it came deep integration into Microsoft's ecosystem, powering Copilot features in Word, Excel, and Windows (**Microsoft**). This version was widely viewed as a turning point, bringing AI into everyday workflows and raising public awareness of both the potential and limits of the technology. However, GPT-4 also intensified debates about transparency, bias, and the environmental cost of AI training, setting the stage for the company's next leap.

By August 2025, OpenAI unveiled GPT-5, described as a “significant leap in intelligence” across writing, coding, reasoning, and vision (**OpenAI**). Early testers noted its striking ability to sustain logical reasoning across long prompts and generate multi-step plans, edging closer to what many call artificial general intelligence. Industry observers compared its coherence and adaptability to human-like problem-solving, while OpenAI claimed it “delivers higher quality answers” and “thinks more thoroughly about complex tasks”, achieving “expert-level results” (**OpenAI**).

The launch of GPT-5 coincided with a clear escalation in competition across the generative AI industry. Anthropic released Claude 3 in March 2024, positioning it as a rival model capable of human-level reasoning (**Reuters**). Google DeepMind refined Gemini 2 in December 2024, integrating advanced multimodal reasoning and coding capabilities (**Google**). Meta, meanwhile, expanded its open-source approach with the release of LLaMA 3.1 in mid-2024, aiming to provide a transparent alternative to proprietary systems (**The Guardian**). Within this context, GPT-5 reinforced OpenAI's role at the centre of the AI ecosystem, powering Microsoft's Copilot suite and OpenAI's own enterprise tools (**Microsoft**). Almost immediately, version 5.1 followed, enhancing speed, memory integration, and customization and creating “warmer” and more “personality” options: proof that iteration itself had become the defining rhythm of the AI race (**The Verge**).

Each generation not only increased capability; it also magnified cost, infrastructure demand, and the scale of what ‘safe, beneficial AGI’ might mean in practice.

THE TURN-FROM IDEALISM:

Why Structure Had To Change

As OpenAI's ambitions grew, so did the tension between its founding ethos and the practical realities of the AI arms race. To develop models like GPT-5 and deliver them globally requires massive computing power, proprietary data, and world-class talent: in short, huge capital.

In 2019, OpenAI acknowledged this when it launched OpenAI LP, describing it as “a hybrid of a for-profit and non-profit”, which they called a “capped-profit company” (**OpenAI**). The idea was to raise investment, attract top AI talent through equity participation, and scale rapidly while still

safeguarding the original mission by limiting investor returns to a fixed multiple: initially up to 100× their investment (**OpenAI**).

In essence, this “capped-profit” model was a deliberate compromise between the ethos of a non-profit and the financial logic of Silicon Valley. It offered investors meaningful upside without surrendering full control or mission integrity to pure market incentives. The concept sparked debate in both legal and ethical circles: critics argued that any profit mechanism might dilute OpenAI’s altruistic vision, while supporters saw it as a pragmatic innovation to finance AGI safely (**The Verge**). In practice, the model allowed OpenAI to raise billions in capital while claiming to retain moral oversight - a hybrid structure unprecedented in the technology industry and frequently described by scholars as a novel form of mission-focused corporate design” (**Collins-Burke**).

Why did this matter? Because the cost of inference alone (running models in production) reached billions. According to leaked figures from TechCrunch, OpenAI spent approximately \$8.65 billion in the first nine months of 2025 just on inference computing (**TechCrunch**). The company could no longer operate purely as a research non-profit.

STRUCTURE TO MATCH THE MISSION (AND MARKET)

To balance mission and market, OpenAI embraced a novel legal design. The non-profit parent would control the for-profit subsidiary, enforcing alignment. On its official website, OpenAI explains how in October 2025 it reorganised: the non-profit became the OpenAI Foundation and the for-profit became an OpenAI Group PBC: a public-benefit corporation required to consider broader stakeholder interests beyond shareholders (**OpenAI**). Equity stakes saw the Foundation holding 26 percent of OpenAI Group, Microsoft 27 percent, and the remaining 47% distributed among employees and investors (**OpenAI**).

In other words: the company moved from pure idealism to idealism matched with industrial scale. The hybrid model offered a legal chassis enabling both large investment and declared mission orientation.

THE TENSIONAL UNDERCURRENT:

Open Research Vs. Commercialisation

This structure reveals a deeper tension: how do you remain transparent, open, and public-centric while building one of the world’s most powerful commercial AI systems? OpenAI’s evolution mirrors that of the broader AI sector: an early culture of openness giving way to selective publication, API-only access, and proprietary models. GPT-2’s partial withholding, GPT-3’s API-only release, and GPT-5’s rapid commercial rollout each highlight those trade-offs.

Moreover, OpenAI's partnership with Microsoft, which began in 2019 and deepened through multiple investment rounds epitomises this shift (**OpenAI**). The "Next Chapter" agreement of 2025 emphasised commercial scale and integration, binding the company's research direction to a major corporate partner (**OpenAI**).

WHY IT MATTERS:

A Prelude To Structural Governance

The evolution from mission-lab to market leader is more than a neat chronology; it forces hard questions about accountability, purpose drift, and who OpenAI ultimately answers to. When a research organisation becomes a platform provider embedded in operating systems, classrooms, and workplaces, what mechanisms ensure the founding promise, benefiting all of humanity, still constrains day-to-day commercial choices? Who adjudicates trade-offs between rapid deployment and safety, or between shareholder value and the public interest?

Those questions now sit at the heart of OpenAI's hybrid chassis: a mission-driven foundation alongside a profit-seeking group company, linked by caps on investor returns and by formal oversight rights. They are not abstract. They colour debates about data access and transparency, pricing and availability (consumer vs. enterprise), safety disclosures, and the concentration of power created by deep partnerships with hyperscalers.

Following this, we will unpack the legal wiring behind these tensions: how the nonprofit parent, OpenAI Foundation, relates to OpenAI Group PBC; how the capped-profit logic operates inside OpenAI LP and its successors; what Microsoft's minority stake actually confers; and how board oversight evolved around the 2023 governance crisis. Understanding that architecture is essential to evaluating the consequences that follow in law, tax, and ethics.

Part II

Governance & Legal structure of OpenAI

LEGAL STRUCTURE

OpenAi was founded in 2015, with the mission to achieve Artificial General Intelligence (AGI) as a no-profit entity, and this choice was genuine because OpenAi carried out research for the benefit of mankind. Later in 2019 due to devouring costs of R&D a profit subsidiary was born, OpenAi LP, this branch used a capped-profit model, so the investor could have a fixed maximum ROI, initially set at 100x, meaning that early investor could make at maximum profit of 100 dollar for each dollar invested; the cap was then gradually shrunk until reaching single digit

numbers; additionally the LP was strictly overseen by the non-profit in order to ensure coherence and attract investors while maintaining control upon the project (**OpenAi**). With the deployment of its biggest success ChatGPT 4.0, the company scaled up and the previous structure was deemed obsolete, thus in October 2025 Sam Altman announced a business reorganisation plan in order to simplify the corporate structure and to complete the recapitalisation. Therefore today the entity is splitted in 2 companies: OpenAi Foundation (non-profit) and OpenAi Group, a Public Benefit Corporation (PBC), which is the for-profit branch. OpenAi is not the only company that restructured using this strategy, in fact many other underwent similar transformation, some examples might include:

- ***Anthropic (Claude.ai)***
 - An Artificial Intelligence research company that was established by former OpenAI staff members, their organizational structure bears considerable resemblance to similar entities, with the primary distinction being that the Public Benefit Corporation (PBC) operates independently and is not subject to control by a non-profit organization or any other external entity (**Anthropic**).
- ***Patagonia Inc***
 - An outdoor apparel business with an embedded mission, they also have a double structure, more specifically they are a Certified B Corporation that also operates under the PBC in the US (**Patagonia**). A Certified B Corporation is a for-profit company with a verified social mission, meeting high standards of social and environmental performance, accountability, and transparency as verified by B Lab (**B-corporation**).
- ***Emerson Collective***
 - A hybrid organization that blends philanthropy, social activism, and private investment under one umbrella. Structured as a Limited Liability Company (LLC), a structure that combines the liability protection of a corporation with the tax flexibility and operational simplicity of a partnership. So, rather than a traditional non-profit, Emerson Collective operates with the agility of a venture firm while pursuing mission-driven objectives in education, immigration reform, climate action, and technology. This setup allows it to fund both for-profit startups and charitable initiatives, effectively merging impact investing with advocacy (**Emerson Collective**).

The foundation directly controls the PBC through a 26% stake in the company. Additionally, special voting rights and warranties give the foundation the power to appoint or remove OpenAI Group's board members at any time.

WHY RESTRUCTURING ?

As mentioned earlier, in late October 2025, OpenAI decided to restructure to simplify its corporate structure, but what does this mean exactly?

First of all, the establishment of a for-profit entity provides OpenAI with the capability to access significantly larger amounts of capital from a broader range of investors and funding sources, while simultaneously allowing the organization to escape the various constraints and limitations that are typically associated with non-profit status, thereby enabling much faster and more aggressive scaling of operations and development efforts. Furthermore, the implementation of this dual organizational structure serves to significantly reduce the company's dependency on Microsoft as a primary financial backer and strategic partner, while also providing enhanced structural clarity and transparency in terms of governance and operational frameworks, which is absolutely key to successfully attracting a more diverse pool of investors and stakeholders who are interested in participating in the company's growth trajectory.

However, the transformation allows for new tax burdens to arise. When such an entity is formed (Corporation or LLC) in the US, it doesn't directly qualify for federal tax exemptions governed by **IRS §501**. To gain such status, the corporation must apply to the Internal Revenue Service (IRS) using Form 1024, (due to the nature and scale of OpenAI) and pass strict operational and organisational assessments. On the other hand, if the entity is for-profit, the owner may prefer pass-through taxation without seeking nonprofit status, which can impose heavy constraints on governance. One legal workaround is to elect S-Corporation status (via IRS Form 2553) if eligible. An S-Corp is not tax-exempt, but it's a pass-through entity, meaning profits are taxed only once at the shareholder level, avoiding the double taxation of a C-Corp. S-Corp eligibility requires the for-profit to be a U.S. entity with fewer than 100 shareholders, all of whom must be individuals holding only one class of stock.

Some important definitions:

- An **S-Corp** is a corporation that passes income, losses, deductions, and credits directly to shareholders for tax purposes, while a **C-Corp** is a standard corporation taxed separately from its owners, subject to corporate income tax (**Wolters Kluwer**).
- **Pass-through taxation** means business profits and losses flow directly to the owners' personal tax returns, avoiding corporate income tax at the entity level (**LII**).

Ultimately, OpenAI must pursue aggressive growth strategies in order to survive and thrive in such a highly competitive and rapidly evolving industry. Without this sustained and substantial growth, achieving its fundamental mission, which is to develop Artificial General Intelligence (AGI) and make it widely available for the benefit of all humanity, would be nearly impossible to accomplish.

THE PBC

Essence and Functioning of a PBC

In this case, the for-profit subsidiary of OpenAI is a PBC registered in the State of Delaware on October 28, 2025. The entity is regulated by the Delaware General Corporation Law (DGCL), specifically Subchapter XV (§§ 361–368), which governs Public Benefit Corporations (PBCs). The definition provided by § 362.

Public benefit corporation defined; contents of certificate of incorporation.

(b) “Public benefit” means a positive effect (or reduction of negative effects) on 1 or more categories of persons, entities, communities or interests (other than stockholders in their capacities as stockholders) including, but not limited to, effects of an artistic, charitable, cultural, economic, educational, environmental, literary, medical, religious, scientific or technological nature. “Public benefit provisions” means the provisions of a certificate of incorporation contemplated by this subchapter.

The Public Benefit Corporation (PBC) is distinct from a traditional non-profit organization in that it does not exclusively or solely pursue what might be considered a purely "charitable" or philanthropic mission. Rather, the PBC is structured to carefully balance and integrate three distinct and important objectives: the financial interests and returns of its stockholders and shareholders, the broader interests and concerns of stakeholders who are materially affected by the corporation's activities and decisions, and the specifically identified public benefit or social purpose that the corporation has committed to advancing.

Why in Delaware

Incorporating in Delaware isn't accidental, it's rather a deliberate strategic move. In fact, over two-thirds of Fortune 500 companies are incorporated there, primarily for its flexible corporate law, specialized courts, and predictable legal environment (**Delaware Corp**).

More specifically, in the legal side, the Delaware Court of Chancery is a highly specialised business court with expert judges and no juries, with a rich case law record, therefore corporate disputes are resolved fast and predictably. Additionally, the presence of the business judgement rule, which protects corporate directors from liability for honest, informed, and good-faith decisions by presuming their actions serve the company's best interests unless proven otherwise, encourages risk taking.

On the taxation side, incorporating in Delaware means there is no state corporate income tax on activities outside Delaware. If your company operates elsewhere, you pay no Delaware income tax, though franchise tax and federal obligations still apply. Delaware also has no taxation on intangible assets (trademarks, patents, etc.) held by Delaware entities, which is a major advantage compared to the IP tax burdens in other regions like Europe. Finally, Delaware law offers more flexible tax elections for C-Corp, S-Corp, or LLC pass-through options depending on your structure and investor profile.

On the economical side it is worth noting that in Delaware, it's possible to have a one-person corporation. This feature can serve as an effective turnaround mechanism, as such an entity can, in turn, hold or control other corporations or businesses that employ their own workforce and directors. In this way, the controlling corporation avoids the need to maintain a broader internal management structure or payroll, thereby streamlining costs and operational complexity.

Delaware offers a corporate-friendly environment through three key features. First, its flexible capital structure supports multiple share classes, investor-preferred terms, and convertible instruments as standard practice. Second, it maintains non-public disclosure of directors and shareholders, in fact only the registered agent's information appears on formation documents. Third, it imposes minimal reporting requirements, which is just a short annual report and franchise tax filing, far less burdensome than European disclosure requirements. However there are some setbacks, for instance there is the franchise tax, which is not unique to Delaware law, but it stands out because it applies to nearly all corporations incorporated there, even if they don't operate in the state, and this flat fee made for small entities can reach up to \$200,000 for large authorized share counts. Lastly there is also the constraint of Dual compliance if the corporation operates elsewhere, qualifying there as a "foreign corporation," paying local taxes and fees too (DGCL).

DEVELOPMENT OF THE MISSION

Alongside its structural evolution, OpenAI's mission has crystallized around three strategic pillars: advancing the frontier of AI research (driving capability development), ensuring alignment and safety (so that AGI remains beneficial and controllable), and promoting broad, equitable access, following the steps of what happened with the World Wide Web, preventing thus any benefits from concentrating within a narrow elite. In practice, this means that OpenAI has consistently positioned itself not merely as a builder of increasingly powerful models, but also as a steward of technology, by making deliberate choices about what to create, how to build it, and when or whether to deploy it. Furthermore this entire process was guided by the broader interests of humanity rather than short-term commercial imperatives.

This hybrid structure, blending nonprofit governance with a capped-profit subsidiary, stands in contrast to the approaches of traditional tech conglomerates such as Google and Meta, as well as mission-driven foundations like Wikimedia. Additionally with Microsoft holding a minority stake, OpenAI's model exemplifies a balance between capital-driven innovation and principled restraint. In a recent statement, the company underscored its broader vision by asserting that access to AI should be considered a universal right, not a privilege.

THE BIG TECH LOOPHOLE

As mentioned earlier, the restructuring of OpenAI enabled the corporation to reduce their dependency on Microsoft without losing capital, allowing thus other major players to acquire equity in the project. However, this is not an isolated strategy, matter of fact it's the most common way to survive in the BigTech industry, and that's because each major player relies on overlapping infrastructures, such as data pipelines and clouds, acquiring stake in the other players gives them preferential access to the fundamental technologies. Some examples of that could be Microsoft's multi-billion-dollar stake in OpenAI or Amazon's investment in Anthropic or even Nvidia's collaborations with all major hyperscalers (**CTech**).

This kind of vertical integration constitutes a self-reinforcing loop investment, where the profits are recycled into other players, consolidating thus the market toward an Oligopolistic market structure. Moreover, by structuring the investments as strategic partnership rather than M&A, the corporations avoid antitrust scrutiny, in fact by using financial instruments such as convertible equity deals, joint ventures, or licensing arrangements, they achieve quasi-control without triggering regulatory barriers. Finally by using a portfolio approach to the BigTech industry, by holding several stakes in the entire innovation spectrum, the corporations project confidence into the market forecast, creating attractive narratives that push investors to bet large amounts of capital in a seemingly risk-free industry.

Signs for a new bubble ?

The compelling narrative and investor overconfidence suggest a potential new bubble, similar to the dot-com bubble of the early 2000s. The industry is highly speculative and thus it's important to consider this threat. Those concerns arise from the financial forecasters, which based their concerns on several factors: rapidly rising prices, overvaluations, widespread hype, and elevated indicators like the P/E ratio.

The Price-to-Earnings (P/E) ratio measures how much investors pay for each dollar of company profit, indicating whether a stock is priced for growth or stability. In big tech, P/E ratios tend to be high because these companies combine growth, scalability, and monopolistic economics, which is a unique combination that distorts traditional valuation benchmarks. High P/Es in this industry reflect a market consensus that these firms are future infrastructure, not mere companies, so the investors aren't simply valuing current profits, they're however acquiring an oligopolistic position in the new digital economy.

These companies operate on high fixed costs but near-zero marginal costs, allowing exponential scalability once infrastructure is built. Add to that the surge in AI-driven growth expectations, vast intangible assets like data and algorithms, and massive institutional inflows treating these firms as quasi-safe assets, and you get valuations that far exceed traditional industries.

FINAL CONSIDERATIONS

Being a non-profit organization does not necessarily mean being a charity, Rolex, for instance, operates as a non-profit entity within a highly commercial and competitive environment. Therefore OpenAI's transition from a non-profit research laboratory to a Public Benefit Corporation should not be seen as a moral regression or ethical compromise, but rather as a structural adaptation to the unprecedented scale and substantial capital demands of contemporary artificial intelligence development. The organisational reorganisation reconciles mission and market imperatives, while preserving public-benefit oversight and accountability while simultaneously enabling sustainable financing mechanisms and industrial competitiveness in a rapidly evolving technological landscape. This strategic shift strengthens rather than weakens OpenAI's founding purpose and original mission. Building safe and broadly beneficial Artificial General Intelligence cannot realistically depend on traditional philanthropic models or donation-based funding alone, as exemplified by organizations like Wikimedia Foundation. It requires a sophisticated framework capable of mobilizing vast financial resources, substantial computational infrastructure, and extensive human capital while maintaining rigorous accountability to public-interest principles and societal benefit. The Public Benefit Corporation model, supervised and overseen by the OpenAI Foundation, represents precisely that necessary compromise, a legal vehicle through which idealism and industrial capacity can coexist productively and sustainably.

BIBLIOGRAPHY

BBC News, and Noel Titheradge. "I Wanted ChatGPT to Help Me. So Why Did It Advise Me How to Kill Myself?" BBC, 6 Nov. 2025, www.bbc.com/news/articles/cp3x71pv1qno.

Berrin Cefa, et al. "Responses to the Initial Hype: ChatGPT Supporting Teaching, Learning, and Scholarship?" *Open Praxis*, vol. 17, no. 2, 1 Jan. 2025, pp. 227–250, <https://doi.org/10.55982/openpraxis.17.2.872>.

Dastin, Jeffrey, and Reuters. "Anthropic Releases More Powerful Claude 3 AI as Tech Race Continues." Reuters, 4 Mar. 2024, www.reuters.com/technology/anthropic-releases-more-powerful-claude-3-ai-tech-race-continues-2024-03-04/.

Ettinghausen, Jeremy. "18 Months. 12,000 Questions. A Whole Lot of Anxiety. What I Learned from Reading Students' ChatGPT Logs." *The Guardian*, 27 July 2025,

www.theguardian.com/technology/2025/jul/27/it-wants-users-hooked-and-jonesing-for-their-next-fix-are-young-people-becoming-too-reliant-on-ai?.

Google. “Introducing Gemini 2.0: Our New AI Model for the Agentic Era.” Google, 11 Dec. 2024, blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/

Microsoft. “Introducing Microsoft 365 Copilot – Your Copilot for Work.” The Official Microsoft Blog, 16 Mar. 2023, blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/.

Montevirgen, Karl. “OpenAI.” Britannica, 10 Apr. 2024, www.britannica.com/money/OpenAI.

OpenAI. “Better Language Models and Their Implications.” OpenAI, 14 Feb. 2019, openai.com/index/better-language-models.

OpenAI. “Entdecke GPT-5.” OpenAI, 7 Aug. 2025, openai.com/de-DE/index/introducing-gpt-5/.

OpenAI. “GPT-4.” OpenAI, 14 Mar. 2024, openai.com/de-DE/index/gpt-4-research/.

OpenAI. “Introducing OpenAI.” OpenAI, 2015, openai.com/index/introducing-openai/.

OpenAI. “OpenAI Charter.” OpenAI, 2018, openai.com/charter/.

OpenAI. “OpenAI LP.” OpenAI, 2019, openai.com/index/openai-lp/.

OpenAI. “Our Structure.” OpenAI, 5 May 2025, openai.com/our-structure/.

OpenAI. “The Next Chapter of the Microsoft–OpenAI Partnership.” OpenAI, 28 Oct. 2025, openai.com/index/next-chapter-of-microsoft-openai-partnership/.

Schlott, Rikki. “Educators Warn That AI Shortcuts Are Already Making Kids Lazy: ‘Critical Thinking and Attention Spans Have Been Demolished.’” New York Post, 25 June 2025, nypost.com/2025/06/25/tech/educators-warn-that-ai-shortcuts-are-already-making-kids-lazy/.

TechCrunch. “Leaked Documents Shed Light into How Much OpenAI Pays Microsoft.” TechCrunch, 15 Nov. 2025, techcrunch.com/2025/11/14/leaked-documents-shed-light-into-how-much-openai-pays-microsoft/.

The Guardian. "Meta Launches Open-Source AI App 'Competitive' with Closed Rivals." The Guardian, 23 July 2024,
www.theguardian.com/technology/article/2024/jul/23/meta-launches-open-source-ai-app-competitive-with-closed-rivals.

The Verge. "OpenAI Says the Brand-New GPT-5.1 Is 'Warmer' and Has More 'Personality' Options." The Verge, 12 Nov. 2025,
www.theverge.com/news/802653/openai-gpt-5-1-upgrade-personality-presets.

The Washington Post, and Daniel Gilbert. "Despite Uncertain Risks, Many Turn to AI like ChatGPT for Mental Health." The Washington Post, 25 Oct. 2024,
www.washingtonpost.com/business/2024/10/25/ai-therapy-chatgpt-chatbots-mental-health/.