



2026

Connective Leadership Summit

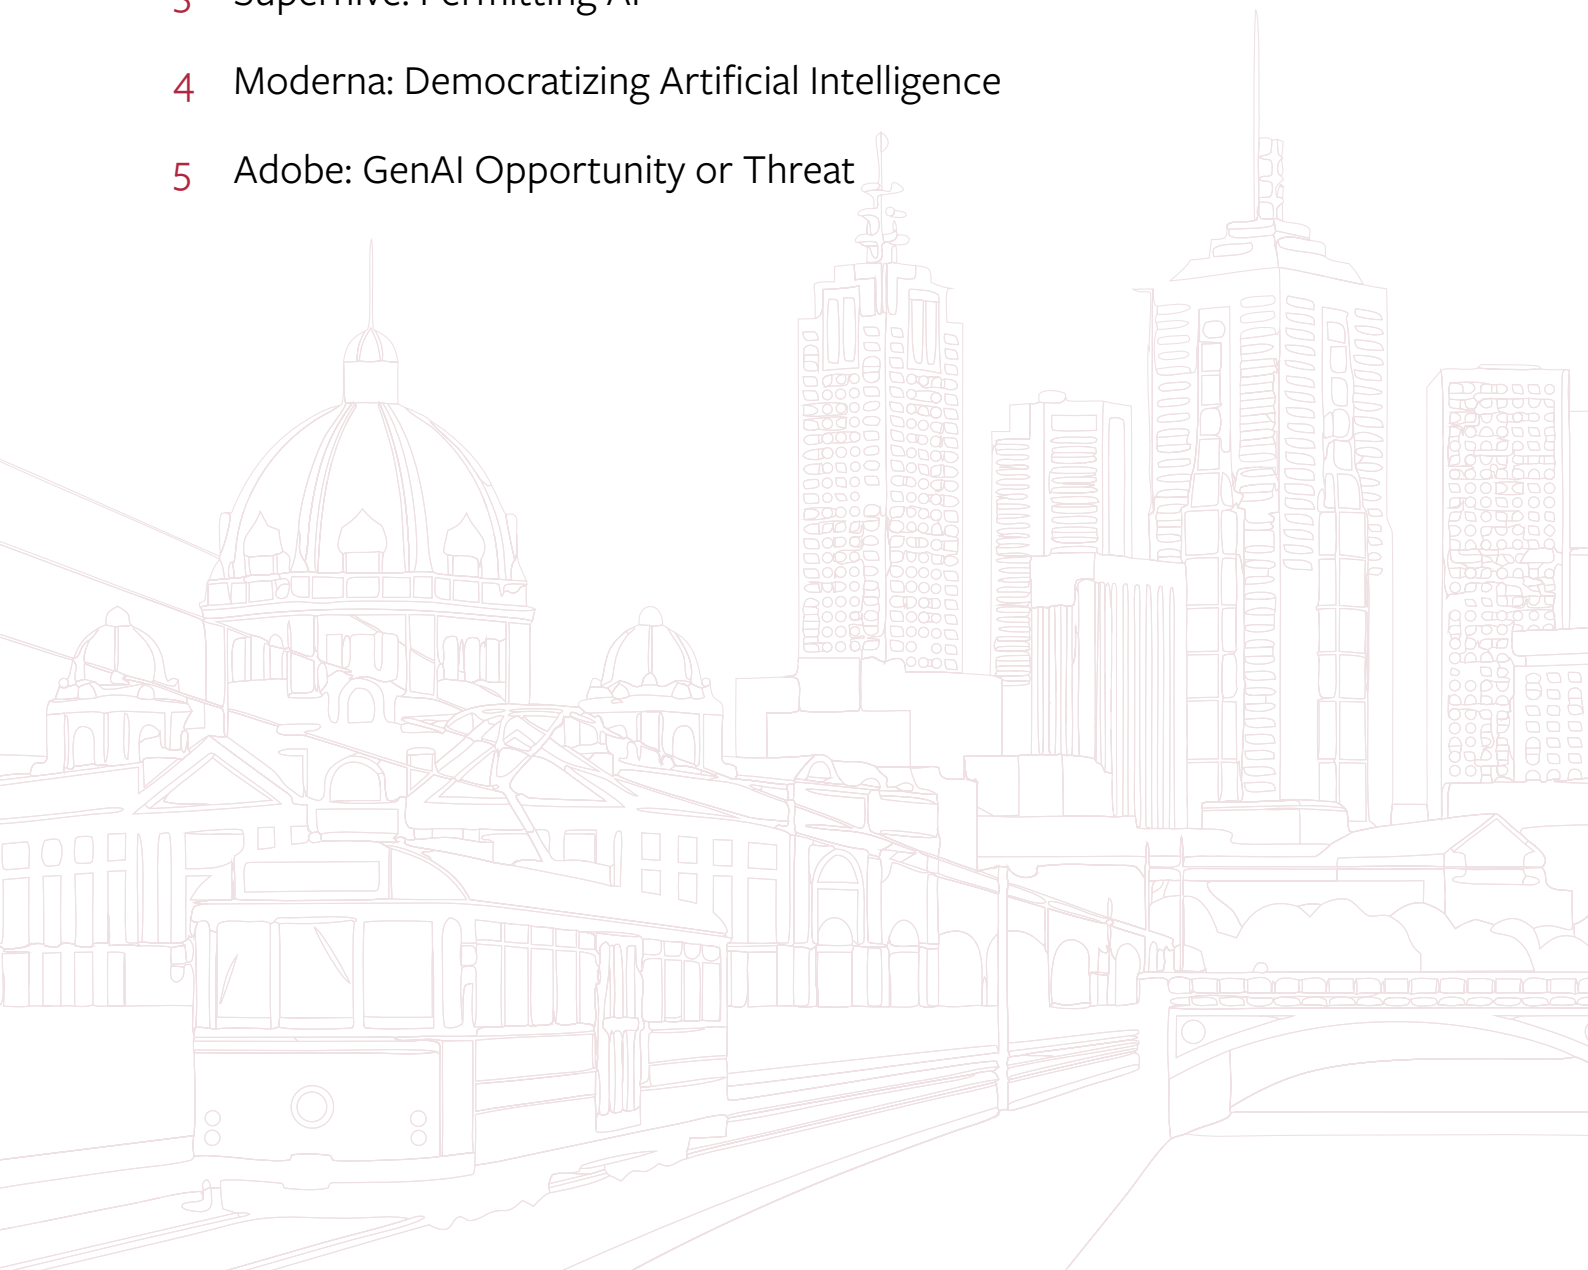
19 - 22 May 2026 | Melbourne, VIC

Pre-Reading Materials



Contents

- 1 Generative AI and the Future of Work
- 2 Gamma: Slides in the Blink of AI
- 3 Superhive: Permitting AI
- 4 Moderna: Democratizing Artificial Intelligence
- 5 Adobe: GenAI Opportunity or Threat



Agenda

Wednesday 20 May 2026

6:30 – 8:30am	Breakfast at leisure
8:40am	Meet in the foyer and walk to Alto Event Space
8:50am	Arrival tea and coffee
8:55am	Connective welcome and opening
9:00 am	Harvard program - The AI Frontier
10:15	Morning tea
10:45am	Harvard program - AI in Our Work
12:15pm	Lunch
1:00pm	Harvard program - Allocating Intelligence
2:30pm	Afternoon Tea
3:00pm	Harvard program - Good Enough? Evaluating AI
4:30pm	Harvard program concludes
6:00pm	Meet in the foyer and walk to dinner - Hochi Mama
6:30 – 9:30pm	Dinner – Hochi Mama

Thursday 21 May 2026

6:30 – 8:30am	Breakfast at leisure
8:40am	Meet in the foyer and walk to Alto Event Space
8:50am	Arrival tea and coffee
9:00am	Harvard program - AI Adoption
10:15am	Morning tea
10:45am	Harvard program - Agentic AI
12:15pm	Lunch
1:00pm	Harvard program - AI Disruption
2:30pm	Afternoon tea
3:00pm	Harvard program - AI Studio and Showcase
4:30pm	Close
6:45pm	Meet in the foyer and walk to dinner – Maison Batard
7:00 – 11pm	Celebration dinner – Maison Batard



2026

Connective
Leadership Summit
19 - 22 May 2026 | Melbourne VIC

Generative AI and the Future of Work





CHRISTOPHER STANTON
MATT HIGGINS
SHIRA ARONSON
MEG SHRIBER

Generative AI and the Future of Work

In May 2023, 11,500 screenwriters in the Writer’s Guild of America (WGA) union went on strike, halting work on television and film productions. The document outlining the union’s position on key issues read: “Regulate use of artificial intelligence on MBA [Minimum Basic Agreement]-covered projects: AI can’t write or rewrite literary material; can’t be used as source material; and MBA-covered material can’t be used to train AI.”¹ Though many issues were at stake, the terms governing the use of generative AI in the creative process turned out to be one of the leading issues in the negotiations.

As the strike went on through the summer of 2023, writers everywhere acquainted themselves with the new technology and stared down a version of the future in which they were increasingly replaceable. In October, as part of the deal that eventually ended the strike, WGA negotiators and the studios agreed on four new terms governing the use of artificial intelligence on union projects, including provisions that “AI-generated material will not be considered source material,” and “AI-generated material can’t be used to undermine a writer’s credit or separated rights.”²

By tying the writer’s credit system firmly to the human writer, allowing the writer and employer the flexibility to decide the role generative tools will play in their own creative processes, and reserving the right to decide its position on the question of using copyrighted materials in training at a later date, the first-ever explicit agreement between management and unionized workers over the terms of generative AI in the workplace satisfied the union’s most important concerns and helped get the striking writers back on the job.

Screenwriters and headline writers alike considered the deal a win. Reflecting the concerns of their readerships, coverage in major publications focused on the implications of the deal for the future of work in other creative and professional fields. In the days following the resolution, *The Los Angeles Times* ran a column headlined: “The writers’ strike was the first workplace battle between humans and AI. The humans won.”³ *Wired* headlined its coverage: “Hollywood writers reached an AI deal that will rewrite history,”⁴ and *The New York Times* ran a guest essay with the headline, “Hollywood’s deal with screenwriters just rewrote the rules around A.I.”⁵

Professor Christopher Stanton and Research Associates Matt Higgins, Shira Aronson, and Meg Shriber prepared this case. This case was developed from published sources. Funding for the development of this case was provided by Harvard Business School and not by the company. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2023, 2024 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

This document is authorized for use only by Mitchell Weiss in 2026.

This article is made available to you with compliments of Connective Leadership Summit, Innovate to Lead for your personal use.
Further posting, copying or distribution is not permitted.

However, there still remained legal issues around the use of generative AI in the workplace. In July 2023, a group of authors filed a class-action lawsuit against OpenAI, accusing it of illegally using their copyrighted works to train its language models. The authors argued that these models were producing infringing, derivative works, and that OpenAI had violated intellectual property laws by training its models on unlicensed data.⁶ While most copyright claims were dismissed in court, the case highlighted unresolved legal questions, leaving creatives uncertain about their ability to protect their work in the future.⁷ Beyond copyright concerns, broader regulatory challenges around issues like data privacy, bias, and AI accountability continue to shape the legal landscape, with governments around the world working to establish frameworks for responsible AI development and use.⁸

In one sense, the WGA deal seemed to spell a new era for human-AI labor relations. In another sense, the drama of humans and labor-saving AI learning to co-exist in society, the workplace, and the economy at large was just getting started.

Generative AI

On November 30, 2022, OpenAI released ChatGPT, a text-generating chatbot based on a large language model (LLM) it had developed. Large language models are AI systems trained on vast amounts of text data to understand and generate human-like language. Though increasingly-capable LLMs had been used in AI labs for years, ChatGPT was among the first generative tools to capture mainstream attention, thanks in part to its intuitive chatbot-style and availability to the general public.⁹ The chatbot appeared to effectively parse queries and formulate human-sounding outputs on almost any imaginable topic in a range of formats, from explanatory prose to creative limerick poetry to complex legal arguments.¹⁰ ChatGPT reached over one million users in its first week,¹¹ and over 100 million users in two months.¹²

As the public acquainted itself with the new capabilities of generative AI, various narratives about the technology's impact swirled in the media and popular culture. Generative AI's potential to reshape fundamental aspects of work and society quickly became a prominent topic of discussion, sparking widespread debates about its impact on the future. Responses to the technology reflected a broad spectrum of concerns and speculations about how it might alter industries, human creativity, and the structure of labor.

For perhaps the first time, white-collar professionals were faced with the prospect that automation could come for a percentage of their jobs. If AI could effectively absorb data, analyze it, and produce compelling written output, what would become of the text-producing white-collar creative class of writers, professors, strategists, marketers, consultants, researchers, journalists, securities analysts, managers, and lawyers, not to mention the code-producing class of programmers and software engineers? Though downstream from text in terms of immediate impact, image-, sound-, and video-producing professionals experimenting at the frontiers of audio and image-generating AIs quickly learned that their own fields would be similarly exposed to the generative AI revolution in the coming years.

How it Began - Natural Language Processing

Modern language models evolved from Natural Language Processing (NLP), a field that dates back to the early days of computer science. Initial efforts to teach machines to understand human language began in the 1950s and 60s. One early example was the Georgetown-IBM experiment (1954), which demonstrated machine translation of Russian sentences into English using a fixed glossary and translation rules.¹³ Around the same time, the Dartmouth Summer Research Project on Artificial

Intelligence (1956) gathered leading thinkers in the nascent fields of artificial intelligence, information theory, and cybernetics, clarifying the boundaries, key questions, and subfields of AI. The conference, which was attended by AI and computing pioneers Marvin Minsky, Claude Shannon, and Nathaniel Rochester, is widely considered the founding event for the field of artificial intelligence.¹⁴

MIT Computer Scientist Joseph Weizenbaum's 1966 ELIZA program was another early milestone in the use of NLP and conversational interfaces. Programmed to reply to a user's text inputs in the conversational manner of a Rogerian therapist ("That's interesting, tell me more"), the replies the program generated were based on a set of simple pattern-matching rules. Though rudimentary by today's standards, ELIZA is important as one of the first-known examples of the genre of interactive software we now call a chatbot.¹⁵

These early NLP efforts were rule-based, and therefore required manual scripting of linguistic rules, an approach that was limited by the complexity and nuance of human language. As computational capacities grew, so did interest in more creative and dynamic approaches. In the 1990s and early 2000s, statistical methods became the cutting edge of NLP research. IBM's Candide (1994), one of the first statistical machine translation systems, was an explicit departure from the rule-based approach, and its success helped reorient the field toward the use of big data and statistical methods.¹⁶

Statistical models like Candide learned from vast amounts of text data and calculated translation possibilities according to their probability, making them more robust and adaptable than the rule-based systems of the earlier era. The first version of Google Translate, launched in 2006, advanced and extended the use of statistical machine translation, relying on vast amounts of bilingual text data (sourced in part from the EU, Canada, and other government bodies where legally required translation and publication of laws and proceedings had produced vast corpora of multilingual matched texts) to derive statistical probabilities for translation.¹⁷

Neural Networks

Neural networks are a type of artificial intelligence modeled after the human brain's structure, designed to recognize patterns and make decisions. They consist of layers of nodes – individual units that process inputs, apply weights, and generate outputs. These networks "learn" by adjusting the weights between nodes based on errors, refining their ability to classify and predict outcomes. Neural networks were first introduced in the 1950s, but struggled to gain traction due to limited computational power and inadequate data.

Statistical methods remained the state of the art in AI until the 2010s, when researchers at Google Brain revitalized the neural network paradigm through deep learning – an approach that uses multiple layers of neurons to learn and classify patterns from vast datasets across various levels of abstraction.¹⁸ Google Brain's defining approach was to model artificial intelligence after the human brain. Their work helped popularize the concept of the node, modeled after the neuron, as the fundamental unit of an AI system.¹⁹ Just as biological neurons receive signals, process them, and transmit outputs to other neurons, nodes in an artificial neural network receive an input, apply a series of weights, activate (or not), and produce a mathematical output that is passed along to another node where the process is repeated. Nodes are nested in layers, the bulk of which (such as those performing the actual classification, clustering, or learning) are hidden from the user.²⁰

Geoffrey Hinton, a professor at the University of Toronto who has been called "the Godfather of AI"²¹, did pioneering work on backpropagation algorithms for neural networks in the 1990s and early 2000s, before they became widespread in AI circles. Backpropagation is a training technique for neural networks that involves taking the error rate from the previous training cycle and feeding it back

through the network to adjust the weights on the next iteration.²² Hinton began working at Google in 2013 after his company DNNResearch was acquired by the search giant.²³ Hinton's contributions helped set the stage for the great leap forward in AI progress in the second half of the 2010s.²⁴

Transformer Architecture

The next major wave of AI was driven by the introduction of the transformer architecture by researchers at Google Brain in 2017.²⁵ The transformer architecture employed a novel self-attention mechanism that allowed the model to process inputs in a holistic, parallelizable way.

To understand how a transformer works, imagine a CEO calling their direct reports into a room and, following an agenda, gets updates from each department head in order. The CEO decides which updates are most important to focus on. They may stray from the agenda to ask follow-up questions, but each digression takes additional time. Now, imagine the CEO has the upgraded capability to process auditory information in parallel. A 30-minute meeting is cut to five minutes. The CEO hears everything at once, makes the right connections between the individual bits of information, and ends the meeting by giving each department head instructions informed by the relevant information in each update.

The upgraded CEO represents the advance between the processing capacity of pre- and post-transformer neural networks. Behind the scenes, the transformer processes the inputs by converting them into a matrix of weighted values that capture the relationships between components. It then uses these weights to determine which elements to focus on next. Thus, training is faster than with linear methods, and the resulting models are better at abstracting and generalized reasoning across contexts in a human-like way.²⁶

By the end of the 2010s, the self-attentive transformer architecture had become the bedrock for a new generation of large-scale language models. Its success in training larger and larger models led to a flurry of attention in the corporate world. Tech companies began investing heavily in AI and LLM R&D, hiring teams of experts to explore how to optimize and scale up transformer-based models. In 2019, Microsoft invested \$1 billion in the newly restructured capped-profit arm of OpenAI.²⁷

From Transformers to ChatGPT

Building on the foundations laid by the transformer architecture, OpenAI introduced a series of models under the GPT (Generative Pre-trained Transformer) banner. GPT models are trained on vast amounts of data to generate human-like text by predicting the next word or phrase in a sequence. OpenAI's first GPT iteration showed researchers the technology's potential. But it was the release of GPT-2, which showcased an impressive ability to generate coherent and contextually relevant text over long passages, that took many by surprise. At the same time, the unexpected power of the model raised concerns about the potential for misuse, leading OpenAI to initially withhold it from the public.²⁸

By the time GPT-3 arrived, it was clear that the NLP landscape had changed. GPT-3 boasted 175 billion parameters – weights that the model adjusts to predict the next word – compared to GPT-2's 1.5 billion. This massive increase allowed GPT-3 to understand and generate text with much greater accuracy. More notably, GPT-3 could perform entirely new tasks upon the first request, without any task-specific fine-tuning, by leveraging its vast pre-existing knowledge and ability to generalize. This capability, known as “zero-shot” learning, represented a major leap forward in AI's flexibility and adaptability.^{29, 30} As word of GPT-3's emergent abilities made its way to the executive suite, businesses began seeing the potential of generative AI for a range of applications: customer service automation, content creation, coding assistance, and more.

ChatGPT, a chatbot built on OpenAI's GPT-3.5 model, was released in November 2022 and immediately captured the public's attention. ChatGPT was designed with a user-friendly conversational interface, making advanced AI capabilities easily accessible to non-technical users. This accessibility, combined with its impressive ability to generate natural-sounding dialogue, led to ChatGPT's rapid adoption. By January 2023, it had become the fastest-growing consumer app of all time.³¹ In February 2023, OpenAI began offering ChatGPT Plus as a subscription service with additional features for \$20 a month.³² In March 2023, OpenAI released GPT-4 to ChatGPT Plus and API subscribers.³³ Experts estimated that the model underlying GPT-4 had over 1.7 trillion parameters—more than 10 times the size of GPT-3.³⁴ Consequently, GPT-4 demonstrated substantial performance gains, notably improving its ability to generate coherent and contextually relevant responses with up to 40% better accuracy on complex benchmarks compared to GPT-3.³⁵

ChatGPT's high-profile success in the consumer market spurred further innovations and a growing interest in pushing the technological and commercial boundaries of the field. Within the next year, Microsoft announced plans to integrate generative "copilots" into its Microsoft 365 suite of productivity software, Adobe embedded a generative model and AI tools into its various design apps, and Google began work integrating its Gemini chatbot into Chrome and Android.³⁶

Beyond Static Data: Retrieval-Augmented Generation and Fine-Tuning

Despite significant advancements in model accuracy, AI hallucinations remained a major concern. Hallucinations—instances where models confidently produced false or misleading information—occurred because GPT models, such as GPT-3 and GPT-4, were designed to predict the next word or sentence based on patterns in their training data, without any mechanism to verify the accuracy of their outputs. Further, these models were reliant on static pre-trained datasets and could not access real-time information or external knowledge. As a result, when prompted with queries outside the scope of their training, these models often fabricated plausible but inaccurate responses. This limitation raised concerns over the reliability of LLMs in critical or specialized applications.

Retrieval-Augmented Generation (RAG) was developed as an approach to addressing these challenges. RAG combines LLMs with external knowledge sources through a vector database—a system that stores unstructured data like research papers or news articles as numerical vectors. Unlike traditional LLMs that relied solely on static, pre-trained datasets, RAG-based models can access real-time, up-to-date information from external sources, including the internet or enterprise databases, at the time of the query.^{37, 38} This ability to retrieve new information made RAG-based systems far more dynamic and reduced the risk of generating outdated or incorrect responses. Businesses, especially those in customer service and technical support, leveraged RAG to provide accurate, context-aware, and personalized responses that reflect the latest available data.

Another method for improving LLM accuracy is fine-tuning, which involves updating the model's parameters with domain-specific data to improve its performance in specialized areas. For example, BloombergGPT, which was trained on general data and fine-tuned for the financial sector, outperformed other models on finance-related tasks at the time of its release. However, while fine-tuning provides high accuracy in specific domains, it relies on static datasets and therefore lacks the flexibility of RAG. The tradeoff between the approaches lies in fixed versus variable costs: fine-tuning requires a larger initial investment, while RAG incurs ongoing computational costs as it retrieves information on demand.

Next Steps in AI Development

By late 2024, several new trends in AI development were beginning to shape the future of the technology. AI agents—autonomous systems that could actively retrieve information, schedule, and trigger workflows based on inputs or goals—were pushing AI beyond text generation and into real-world applications. Similarly, multi-modal models, capable of processing and generating data across text, images, and audio, were enhancing fields such as medical diagnostics and creative industries by allowing AI to work seamlessly across different data types.

One of the most significant advancements was Chain of Thought (CoT) reasoning, a technique that allowed models to break down complex problems into smaller steps.³⁹ OpenAI's 4-o1 release, which incorporated CoT prompting, reinforcement learning, and error-checking, exemplified that increasing computational power in reasoning tasks led to step-change improvements in performance.⁴⁰ These updated models were able to reason more effectively than their predecessors, demonstrating improved accuracy in complex decision-making and problem-solving.⁴¹ These developments also contributed to increased transparency in AI models, reducing the "black-box" nature of earlier versions. CoT allowed users to follow the reasoning process step by step, while human-led reinforcement learning helped develop reasoning that more closely mirrored human thought processes.⁴²

Another key shift by late 2024 was the move away from development strategies that were primarily focused on increasing the number of model parameters. While increasing parameter counts had significant effects on accuracy in early upgrades of large language models, the marginal returns of adding parameters began to diminish past a certain level.⁴³ Techniques like CoT and RAG offered more efficient improvements in performance without the ballooning costs and memory usage associated with building continuously larger models. These innovations set the stage for the next generation of AI, especially in fields requiring deep, task-specific reasoning.

Ethical and Safety Concerns in AI Development

As generative AI technology advanced in popularity and capability, concerns over its ethical and societal implications became increasingly prominent.⁴⁴ Researchers shifted their focus from merely scaling models to ensuring that they were aligned, interpretable, and safe for public use. AI developers—including OpenAI, Microsoft, Anthropic, and Google—implemented various safety measures and guardrails in their publicly available models and chatbots, aiming to reduce biases and prevent harmful outputs.⁴⁵

Long-term safety was also a concern. In March 2023, an open letter signed by Elon Musk, Steve Wozniak, and over 1,000 other technology leaders called for a temporary pause in AI development, citing "profound risks to society," such as the potential for job displacement, misuse of AI in disinformation, or applications in military settings.⁴⁶ In May 2023, Geoffrey Hinton resigned from Google to raise awareness about AI risks.⁴⁷ Hinton became the first signatory of a statement by the Center for AI Safety, which read, "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war." The statement was co-signed by a group of industry leaders, including Sam Altman of OpenAI, Demis Hassabis of Google DeepMind, and Dario Amodei of Anthropic.⁴⁸ Further capturing these fears, a 2024 survey of over 2,700 prominent AI researchers found that 47% of respondents estimated at least a 10% chance of AI contributing to human extinction.⁴⁹ These concerns reflected broad fears that unchecked AI development could lead to destabilizing critical infrastructure or enabling large-scale autonomous weapons, which could, under worst-case scenarios, pose existential threats to humanity.

Impact on Labor and Employment

As generative AI advanced, its impact on labor and employment was poised to echo the double-edged sword of past technological revolutions—displacing some jobs, creating new ones, and reshaping others. During the mechanization of textile production in the Industrial Revolution, for example, many weavers and artisans were displaced, but new roles emerged in machine oversight and specialized tasks like sewing cuffs, collars, and buttons. Mechanization did not eliminate labor but redefined it, as workers transitioned from crafting complete items to becoming part of a larger industrial process.⁵⁰ However, in spite of the overall increase in productivity brought by the new systems, some historians assert that the monotonous, tedious, and often hazardous nature of the new jobs, combined with the unequal power dynamics between employers and employees, left many workers worse off under the new system.^{51, 52}

Similarly, rather than exclusively displacing human labor, generative AI was expected to drive a process of coevolution, where human roles adapted alongside AI, creating new forms of work that blended human oversight with AI-driven processes. While AI promised increased efficiency, it also raised concerns about job quality. Like early industrial jobs, some roles created by AI were predicted to lack the autonomy or fulfillment of earlier jobs. These historical trends therefore highlighted the importance of managing AI's integration thoughtfully to ensure that new opportunities optimized productivity gains without sacrificing worker well-being.

Commercial Implementation: Video Games and Generative AI

After ChatGPT's release to the general public in fall of 2022, companies across a variety of sectors raced to implement generative tools, revolutionizing internal business operations and consumer experiences alike. Despite being largely insulated from the impacts of automation in the past, creative fields now sought out opportunities to integrate generative technology. The video game industry, a \$217 billion global market as of 2022, quickly began finding applications for AI in game development.⁵³

The video game company Electronic Arts, for example, announced in March 2024 that it was incorporating generative tools in up to 60% of its development process, and expected a 30% increase in efficiency. Notably, the announcement came one week after the company laid off 670 employees – about 5% of their workforce.⁵⁴ By July, EA expressed that generative AI had been integral to creating *College Football 25*, allowing them to produce over sixteen thousand character models – all based on real players, fans, and coaches – within five months.⁵⁵ Other video game developers including Microsoft had similarly announced that generative AI would become essential to their game development processes, especially in localization, microtransactions, and game testing.

While many game developers promised that AI integration wouldn't lead to significant job displacement, an estimated 10,500 employees were laid off across the industry in 2023, and 11,000 more had been laid off by the summer of 2024. Some industry professionals reflected that while AI might not yet possess the creativity needed to replace truly great art, it was already capable of producing work that was “good enough” to meet commercial standards, and therefore already capable of replacing workers in certain roles.⁵⁶ In July 2024, over 160,000 video game voice and motion capture actors in the SAG-AFTRA union announced a strike, demanding “critical AI protections” for union members, of the sort that WGA screenwriters had won the summer before.⁵⁷

The Future of (Human) Jobs

The rollout of generative AI tools in the early 2020s created widespread anxiety about the trajectory of the technology and what its eventual capabilities might do – both to human labor and the social contract more broadly. The specter of artificial general intelligence (AGI) – non-human intelligence that could reason, learn, and act with autonomy across domains without human intervention – loomed large in the debates between “accelerationists” (advocates for pushing AI technology forward as quickly as possible) and “doomers” (advocates for safety and alignment). Among the immediate fears were unemployment and the potential destabilization of societies.

Many of these observers’ concerns about generative AI’s impact on labor and employment mirrored those raised in the Industrial Revolution: de-skilling of jobs, downward pressure on wages, increased surveillance, and diminished bargaining power of human workers in the face of competition from machines newly able to do the same type of work. Would machines take over at least some of the daily work tasks, both cognitive and physical, that some humans enjoy and find meaning in? If so, how would humans respond and adapt?

One possibility was that fewer humans would need to work. The World Economic Forum’s 2023 “Future of Jobs Report,” based on surveys of more than 800 companies, found that employers expected a net reduction of 14 million jobs, or 2% of current employment, by 2027, due in part to labor-saving technologies resulting from advances in AI and automation.⁵⁸ A year later, staff at the IMF speculated that up to 70% of all jobs in advanced economies such as the United States were exposed to some degree of AI automation, and up to 1.4 billion full time jobs could be seriously impacted globally.⁵⁹ Policies ranging from workforce re-skilling programs to universal basic income were proposed and debated as potential solutions to the societal and economic shifts that more widespread AI might entail.^{60, 61}

Forecasting how AI might impact the workplace is potentially fraught work. In the past, even reputable organizations such as McKinsey and Company have overestimated the workplace impact of new technologies by orders of magnitude.⁶² Though the precise impacts of generative AI on the future of work were uncertain, recent academic work offered valuable frameworks.

“Automation and New Tasks: How Technology Displaces and Reinstates Labor,”
Acemoglu and Restrepo (2019)

Before the surge of interest in generative AI following ChatGPT’s release, Acemoglu and Restrepo (2019) examined key labor market issues during periods of technological change. They present several key historical insights.

Historically, technology displaces and reinstates labor. “The history of technology is not only about the displacement of human labor by automation technologies... Instead, the displacement effect of automation has been counterbalanced by technologies that create new tasks in which labor has a comparative advantage. Such new tasks generate not only a positive productivity effect but also a reinstatement effect – they reinstate labor into a broader range of tasks and thus change the task content of production in favor of labor. The reinstatement effect is the polar opposite of the displacement effect and directly increases the labor share as well as labor demand.”⁶³

Effect of automation on labor demand = Productivity effect + Displacement effect. Automation enables capital to substitute for labor in certain tasks, increasing productivity by allowing cheaper substitutes to perform previously expensive tasks at a lower cost (i.e., machines or unskilled workers substituting for expensive skilled workers). The productivity effect makes the producer more efficient while shifting labor demand to new, non-automated tasks where labor demand is higher. The

displacement effect reduces the demand for one type of labor, usually expensive skilled labor, reducing labor's share of the value added (labor share of value is highest when skilled labor is both scarce and required – as with, say, the current job market for AI engineers). The net impact depends on whether the economic value of the productivity gained outweighs the economic value of the labor displaced. As the authors say: "Automation therefore increases the size of the pie, but labor gets a smaller slice. There is no guarantee that the productivity effect is greater than the displacement effect; some automation technologies can reduce labor demand even as they raise productivity."⁶⁴

The net impact on labor demand depends on the broader labor market. "Different technologies are accompanied by productivity effects of varying magnitudes, and hence, each set of automation technologies may impact labor demand in different ways. Likewise, because the productivity gains of automation depend on the wage, the net impact of automation on labor demand will depend on the broader labor market context. When wages are high and labor is scarce, automation will generate a strong productivity effect and will tend to raise labor demand. When wages are low and labor is abundant, automation will bring only modest productivity benefits and can end up reducing labor demand."⁶⁵

In settings where wages are already low, it can be hard to squeeze enough productivity out of the new technology for there to be any net productivity gain at all, let alone enough to spur an overall increase in labor demand from reallocating displaced workers into new tasks. Take the example of self-checkout machines at grocery stores. Wages for grocery store clerks are already low, so automating away the human jobs might result in limited economic productivity gains. Self-checkout machines would, in that case, reflect "so-so automation" – technologies that replace workers without offering any real productivity benefits. Self-checkout machines simply transform paid labor by a store employee into unpaid labor by the customer, who scans and bags their own groceries. The store may save a bit on labor costs if the machines can be installed and maintained for less than the clerk's salary, but the clerk's salary – transformed into a capital investment – disappears from the wider economy. Labor demand is reduced, net productivity is flat, and customers now bag their own groceries.⁶⁶

Recent waves of automation are not creating enough new tasks. Historically, automation's displacement of labor has been counterbalanced by the creation of new tasks where human labor has an advantage. Yet, according to the authors, much of the automation in recent years has created only minor productivity benefits while new task creation has slowed. The authors suggest that an increasing focus on maximizing profits for shareholders, combined with the rise of flexible employment arrangements and government policies that favor capital over labor, have allowed the new task creation that accompanies displacement to stall. A related possibility is that while some new tasks are being created at the leading edge of the economy, these jobs are fewer in number and less accessible to ordinary workers than in past eras because of skill and education gaps. Societies have continued to encourage the automation of expensive tasks using technology, but the replaced workers have not been reinstated in the kinds of new, non-automated tasks that have accompanied other periods of rapid technological advancement.⁶⁷

In a 2024 essay building off of his work with Restrepo, Acemoglu concludes that AI has a high potential for continuing the trend of "bad" automation that displaces workers without creating new tasks or meaningful productivity gains. Specifically, he argues that, while AI-driven automation may result in modest productivity gains (0.71% over the next 10 years), these gains are likely overstated because they come primarily from tasks that are easier to automate. Meanwhile, the creation of new tasks that could benefit workers has slowed, exacerbating inequality and concentrating economic gains among capital owners rather than workers.⁶⁸

"GPTs are GPTs," Eloundou et al. (2023)

Building on prior work examining the impact of machine learning technologies on jobs, a paper by researchers at OpenAI and the University of Pennsylvania compared the particular tasks performed within occupations to the set of tasks that could be readily automated with LLMs. They then aggregated the occupation-level exposure to generative AI to estimate the possible impact of the diffusion of LLM technology on the U.S. labor force.⁶⁹

Half of all US jobs could be exposed. Across all jobs, about 15% of tasks were directly exposed to LLMs. The researchers found that about 80% of the workforce could have about 10% of their current job tasks exposed, while approximately 20% could have at least 50% of their tasks exposed. Those figures measured only the tasks directly exposed to automation by core LLM functionality. When software, tools, and interfaces built on top of LLMs (such as ChatGPT and coding assistants) were factored in, the percentage of U.S. job tasks exposed to generative AI increased to between 47 and 56%.⁷⁰

Higher paying jobs are more exposed. Researchers found that jobs with manual labor components would be less affected by LLMs, while jobs heavy in coding and writing tasks were relatively more exposed. By sector, jobs in manufacturing, agriculture, and mining had lower exposure, while jobs in information-processing industries were more exposed. By income distribution, jobs with higher wages showed higher exposure than lower-wage jobs (see **Exhibit 1**).⁷¹ They summarized:

If 'GPTs [Generative Pre-trained Transformers] are GPTs [general-purpose technologies],' the eventual trajectory of LLM development and application may be challenging for policymakers to predict and regulate. As with other general-purpose technologies, much of these algorithms' potential will emerge across a broad range of economically valuable use cases, including the creation of new types of work.⁷²

"Generative AI at Work," Brynjolfsson, Li, and Raymond (2023)

In a field study conducted at the customer contact center of a Fortune 500 company, researchers measured the impact of access to a generative AI tool on contact-center worker productivity. Customer service agents interacted with customers through a chat-based interface while the AI tool, in a side window, processed the interaction in real-time and suggested responses, essentially giving agents a highly-capable predictive autocomplete. Agents were free to use or ignore the AI assistant's suggestions. Researchers gave one group of agents AI access and compared their productivity to a control group that had no tool available.⁷³

Access to generative AI tools increased productivity. Using resolutions per hour as the primary measure of productivity, researchers found that agents with access to the tool resolved 14% more cases per hour than agents without the tool (see **Exhibit 2**). Agents using the tool also exhibited lower average handle time, an increase in chats per hour, and higher customer satisfaction scores as compared with the control group. In subsequent analyses, access to AI assistance improved customer sentiment, reduced requests for managerial intervention, and improved employee retention.⁷⁴

Lower-skilled workers benefited more. The researchers used agents' pre-intervention productivity metrics to compare pre-and post-intervention productivity at the individual and group levels. Using the pre-intervention productivity baselines, they were able to explore the impact of access to the generative tool across the distribution of agent productivity. Those in the lowest 20% of the pre-experiment productivity distribution saw a 35% increase in resolutions per hour, while workers in the highest skill quintile saw no significant change in resolutions per hour (see **Exhibit 3**).⁷⁵

“Navigating the Jagged Technological Frontier,” Dell’Acqua et al. (2023)

In a study conducted with Boston Consulting Group, a team of researchers examined the impact of access to generative AI on highly-skilled knowledge workers. A total of 758 consultants, all aware they were participating in a study, were divided between two experiments, one designed to test the impact of generative AI on tasks researchers considered “inside the frontier” of generative AI at the time, and a second experiment designed to test the impact of generative AI on tasks considered “outside the frontier” of generative AI. This was one of the first studies to examine the impact of generative AI on high-skill professionals engaged in primarily creative knowledge work.⁷⁶

Inside the frontier, professionals with generative AI were more productive. Consultants were pre-tested on a similar task to establish a baseline, then given tasks across four categories considered to be within the boundary of existing generative AI capability: Creativity (e.g., “Propose at least 10 ideas for a new shoe targeting an underserved market or sport.”); Analytical thinking (e.g., “Segment the footwear industry market based on users”); Writing proficiency (e.g., “Draft a press release marketing copy for your product”); and Persuasiveness (e.g., “Pen an inspirational memo to employees detailing why your product would outshine competitors.”) Consultants with access to generative AI finished 12% more tasks, completed those tasks 25% more quickly, and produced 40% higher quality results than those without AI access (see **Exhibit 4**).

Inside the frontier, lower-skilled professionals benefitted more. Though all consultants in the AI group saw an increase in performance, consultants who pre-tested in the bottom half of the pre-intervention skills distribution saw a greater increase (43%) than consultants who pre-tested in the top half (17%) (see **Exhibit 5**). One possible explanation is that generative AI is inherently limited to reproducing outputs that align with what its trainers have defined as “high-quality.” Since these benchmarks are based on the performance of top consultants, AI lacks the capacity to push the performance of already high-achieving individuals significantly further.⁷⁷

Outside the frontier, professionals with generative AI performed worse. Simulating case interviews used in recruiting, consultants were asked to make a product decision based on an analysis of qualitative and quantitative data. This task was presumed to be outside the frontier. The authors designed the task to have only one correct answer to simplify evaluation. In this experiment, the assistance of generative AI hindered, rather than helped performance. Consultants with no AI access got the right solution about 84% of the time, while consultants with AI access got the right answer only about 65% of the time. The authors speculate that the reduced performance of consultants with AI may have resulted from their trusting the AI assistant too much and “falling asleep at the wheel.”⁷⁸

The capabilities of generative AI are a jagged frontier. In practice, it is hard to know exactly where the capabilities of generative AI begin and end. The principal metaphor of the paper, “the jagged frontier,” is intended to illustrate this fact (see **Exhibit 6**). Some tasks that seemed like they should be inside the frontier were actually outside it. For example, many LLMs in 2023 struggled with simple arithmetic due to their reliance on pattern recognition rather than logic or calculation (later chatbots addressed this issue by integrating modules to handle mathematical reasoning). Meanwhile, some things that seemed like they should be outside the frontier could be done surprisingly well by generative AI—generating ideas for a new product name, or putting together an itinerary for an upcoming trip. These expectations shift as users gain experience. Ultimately, the authors remind, it is impossible to know where the frontier lies on any given task without experimentation.⁷⁹

Successful generative AI users were centaurs or cyborgs. Across both studies, the researchers observed two patterns of navigating the jagged frontier of generative AI: “centaurs” and “cyborgs.”

Centaurs, named for the half-human/half-horse creatures of Greek mythology, used a strategic division of labor between human and machine, allocating tasks to either themselves or to the generative AI on a task-by-task basis depending on their own strengths and the tasks they could confidently hand off to AI. For example, a consultant who was more confident in their analytical skill than in their writing might hand off most of the writing tasks to the AI, while leading the analysis themselves. Cyborgs, named for the integration of man and machine into one body, constantly experimented at the boundaries of the system's capabilities, intertwining generative AI into their workflows at every step. The study did not compare the effectiveness of the two approaches.⁸⁰

Social Implications and Policy Considerations

The rise of generative AI was expected to create societal-level shifts reminiscent of past technological revolutions. The Industrial Revolution, for example, precipitated major social shifts including urbanization, changes to family life, class divisions, and cultural upheavals as traditional communities and lifestyles were upended.⁸¹ The technological path of industrialization helped to create a new class of workers who lived in cities and had new sources of contact with books, ideas, and elites. Some believed the interplay of new economic conditions, physical conditions, and class consciousness helped give rise to new political movements and ideologies, from labor organizing and unions to Marxism and communism.

Similarly, in the late 20th century, technological improvements in telecommunications, logistics, and shipping paved the way for another macro trend: globalization. Some historians saw globalization, rooted in the economic logic of flexible labor markets, free trade, and comparative advantage, as responsible for hollowing out the industrial base in wealthier but de-industrializing countries like the United States and the United Kingdom. The impact of these economic changes on workers was a hotly debated factor in the modern politics of developed nations.⁸²

Brynjolfsson's "The Turing Trap" Essay (2022)

Erik Brynjolfsson, a prominent researcher on the social impacts of innovation, published an essay addressing concerns about the future of an AI-driven society. He argued that generative AI could take two paths: one that automates human labor out of existence, or one that augments and enhances it. Societies that prioritize automation might see short-term economic gains benefiting those in leadership, but at the cost of long-term consequences for the broader population. This concentration of wealth and power in the hands of a few would drive future innovation further toward automation, reinforcing inequality and limiting opportunities for the less powerful.⁸³

The concept of this phenomenon as a "Turing Trap" stems from Alan Turing's foundational work on artificial intelligence. Turing's work inspired the goal of creating machines that can fully replicate human intelligence, but Brynjolfsson warns that this focus can lead societies to prioritize replacing human labor entirely, rather than augmenting it. Over time, if this path is carried out without the voice of workers, the bargaining power of the replaced workers is reduced to the point that they are powerless to influence policy. According to Brynjolfsson, generative AI will boost labor productivity. However, "no economic law ensures everyone will share this growing pie."⁸⁴

If the default path of this technological revolution under current economic incentives is to replace human workers with machines, the author argues that it is critical for policymakers to alter the incentives to promote the path of augmentation, which he views as the best path for society as a whole. Brynjolfsson published a follow-up essay in 2024 affirming that AI augmentation held more productivity benefits than machines or humans working alone, especially benefitting less-experienced

workers. Approaches that encouraged human-machine collaboration, keeping humans in the job while increasing their capacity for efficient, effective work, had the potential to lead to more equally-distributed incomes.

Agarwal, Gans, and Goldfarb's "The Turing Transformation" Paper (2023)

A group of University of Toronto economists argued that AI-driven task automation can enhance job prospects and expand employment by democratizing access to certain high-skill tasks, reducing the "skill premium" tied to higher wages.⁸⁵ By making once-scarce skills more widely available, the authors argue, new opportunities will be created for lower-wage workers, leveling the economic playing field.

They illustrate this theory with an example of a once-scarce professional skill that was democratized by technology: the knowledge of local streets that enabled taxi drivers to navigate without maps. What was once the exclusive preserve of career drivers has been widely distributed by GPS systems in the smartphone era, allowing many more people to do the job. Perhaps these new drivers don't perform as well as career drivers initially, but they are capable enough to transport customers as ridesharing contractors. As app-based ride-sharing services grew the supply of new drivers, traditional taxi driver wages fell.

If the main consideration is the impact on licensed taxi drivers, this development appears to result in exactly the kind of reduction in bargaining power and de-skilling through automation that other economists have warned against. But when one widens the scope to consider the impact on workers lower down the skills and earnings ladder, the overall impact is an increase in economic opportunity across the population.⁸⁶ When technology automated a once-scarce skill, the total number of drivers earning some income from driving exploded, as did the number of people willing to pay for a ride, creating a net gain in economic activity. Though career drivers saw a reduction in earnings from the new competition, many new jobs were created for those at the lower end of the skills distribution. The same transformation that was bad for some workers created new opportunities for others.⁸⁷

The paper also argues that the dichotomy between automation and augmentation is misleading. A technology that augments the labor of one kind of worker may automate away the work of another, regardless of its initial intent. As the authors put it: "One person's augmentation is another's automation."⁸⁸ Office productivity software, which can be seen as augmenting the work of some traditionally high-earning knowledge workers, also automated away a number of traditionally lower-earning administrative workers who were no longer needed to answer phones or type documents. Flipping traditional arguments about automation and inequality on their heads, the authors argue that "the winners and losers are not determined by whether the technology seems to replace or augment human tasks ... the winners and losers are determined by whether the augmentation affects lower-wage workers and automation affects those already doing well."⁸⁹

Looking Ahead

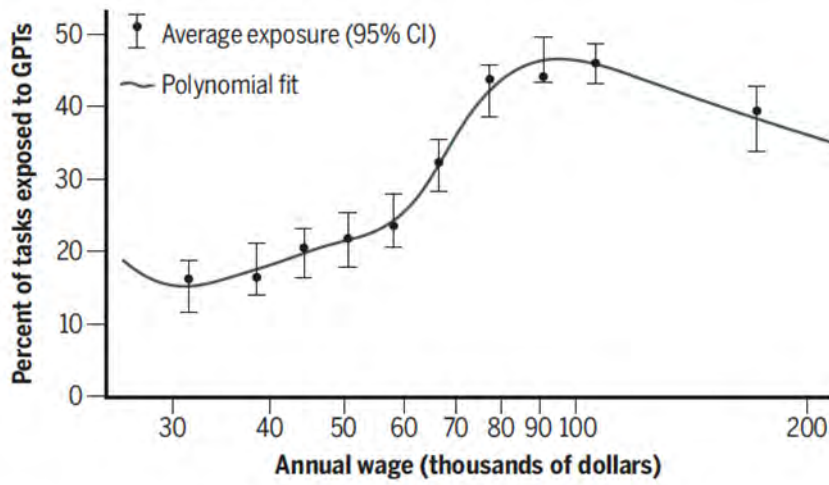
In an interview with MIT President Sally Kornbluth, OpenAI CEO Sam Altman expressed a measured yet decidedly optimistic view on the future of AI, despite growing concerns about its potential disruption. Referencing AI as "the biggest, the best, and the most important" technology revolution yet, he stated:

AI will continue to get way more capable... People are using it to create amazing things. If we could see what each of us can do 10 or 20 years in the future, it would astonish

us today... Society always holds space for doomsayers. There's value to that. I'm happy that they exist. I think it makes us think harder about what we're doing. But I think the better question is, what needs to happen to navigate [AI] safety sufficiently well? AI is going to eliminate a lot of current jobs, and there will be classes of jobs that totally go away. AI is also going to change the way a lot of current jobs function, and it's going to create entirely new jobs. We'll find that humans are wired to care about other humans. As long as you believe humans want to create and be useful, we're not going to run out of things to do.⁹⁰

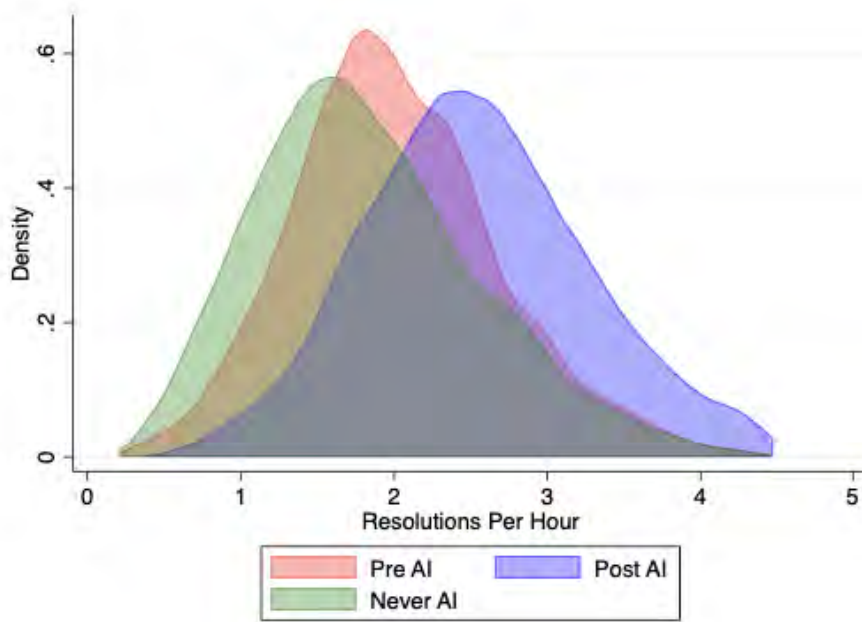
Altman's remarks offer a hopeful yet realistic perspective on AI's future. While acknowledging that AI will disrupt the labor market, he emphasizes that it will also create new opportunities for work and human creativity. His confidence in humanity's resilience frames AI as a tool that, with careful management, can enhance rather than diminish human contributions to society. As AI capabilities continue to advance, Altman's balanced view underscores the importance of thoughtful navigation of both its risks and rewards.

Exhibit 1 Exposure to GPTs by Income



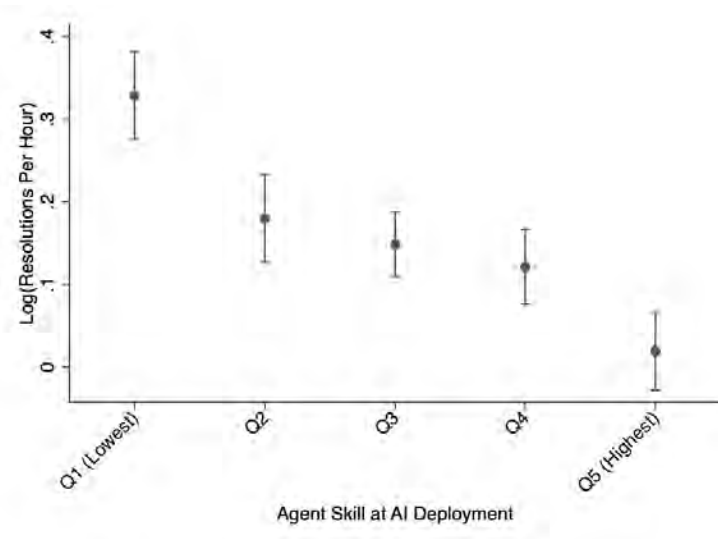
CI, confidence interval; GPT, generative pretrained transformer; LLM, large language model.

Source: Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: Labor Market Impact Potential of LLMs," *Science* 384, 1306-1308 (2024), [10.1126/science.adj0998](https://doi.org/10.1126/science.adj0998).

Exhibit 2 Brynjolfsson, Li, and Raymond (2023): Resolutions per Hour across Three Conditions:

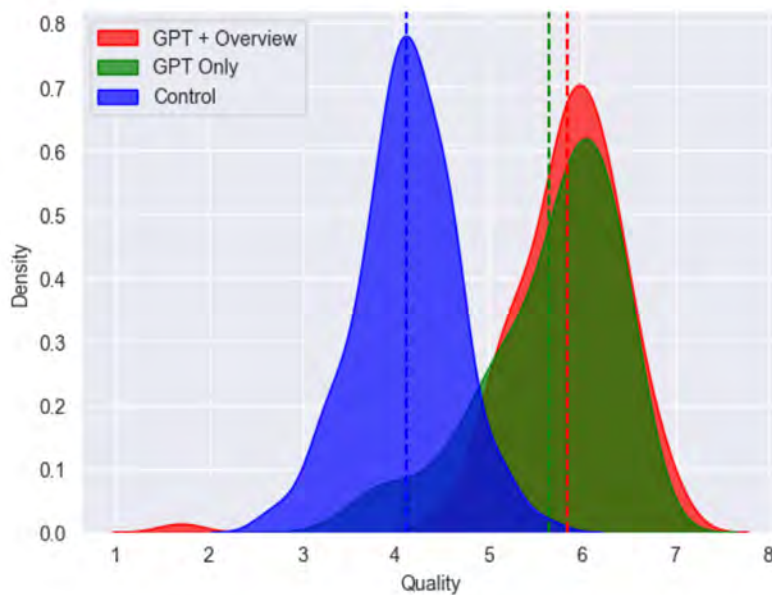
Source: Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," NBER Working Paper 31161, April, 2023, rev Nov 2023, <https://www.nber.org/papers/w31161>. © 2023 by Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. All rights reserved.

Exhibit 3 Brynjolfsson, Li, and Raymond (2023): Impact of AI on Resolutions per Hour, by Skill Quintile at Deployment

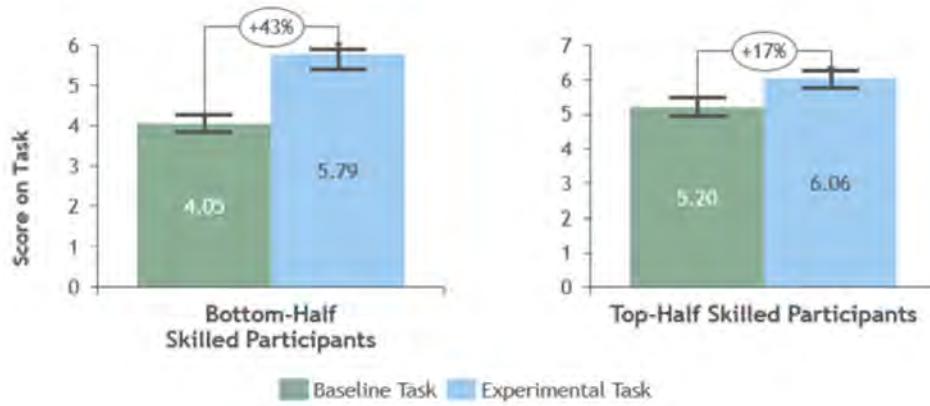


Source: Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, “Generative AI at Work,” NBER Working Paper 31161, Apr, 2023, rev Nov 2023, <https://www.nber.org/papers/w31161>. © 2023 by Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. All rights reserved.

Exhibit 4 Dell’Acqua et al. (2023): Performance Distribution – Inside the Frontier

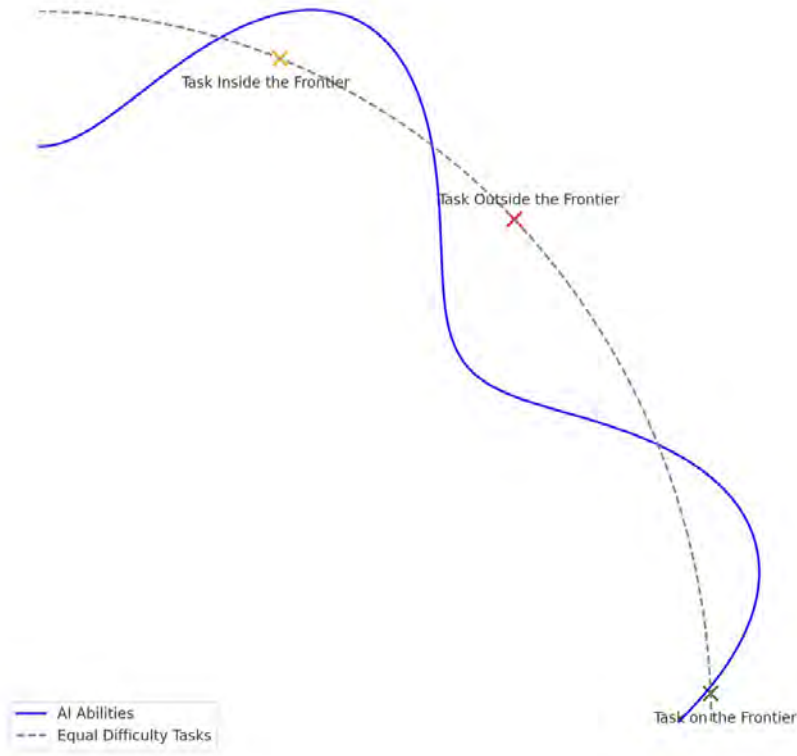


Source: Fabrizio Dell’Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Lisa Kraymer, François Candelon, and Karim R. Lakhani, “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” HBS Working Paper, Sep 18, 2023 (rev Sep 27, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321, accessed Nov 13, 2023.

Exhibit 5 Dell'Acqua et al. (2023): Bottom-Half Skills and Top-Half Skills – Inside the Frontier

Source: Fabrizio Dell'Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Lisa Kraymer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," HBS Working Paper, Sep 18, 2023 (rev Sep 27, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321, accessed Nov 13, 2023.

Exhibit 6 Dell’Acqua et al. (2023): Illustration of the “Jagged Frontier”



Source: Fabrizio Dell’Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Lisa Kraye, François Cadelon, and Karim R. Lakhani, “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” HBS Working Paper, Sep 18, 2023 (rev Sep 27, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321, accessed Nov 13, 2023.

Endnotes

¹ WGA, “WGA Negotiations—Status as of May 1, 2023,” https://www.wgacontract2023.org/uploadedfiles/members/member_info/contract-2023/wga_proposals.pdf, accessed Oct 21, 2023.

² WGA, “Summary of the 2023 WGA MBA,” <https://www.wgacontract2023.org/the-campaign/summary-of-the-2023-wga-mba>, accessed Oct 21, 2023.

³ Brian Merchant, “Column: The Writers’ Strike was the First Workplace Battle between Humans and AI. The Humans Won,” Los Angeles Times, Sep 25, 2023, <https://www.latimes.com/business/technology/story/2023-09-25/column-sag-aftra-strike-writers-victory-humans-over-ai>, accessed Nov 8, 2023.

⁴ Will Bedingfield, “Hollywood Writers Reached an AI Deal That Will Rewrite History,” Wired, Sep 27, 2023, <https://www.wired.com/story/us-writers-strike-ai-provisions-precedents/>, accessed Nov 8, 2023.

⁵ Adam Seth Litwin, “Hollywood’s Deal With Screenwriters Just Rewrote the Rules Around A.I.,” The New York Times, Sep 29, 2023, <https://www.nytimes.com/2023/09/29/opinion/wga-strike-deal-ai-jobs.html>, accessed Nov 8, 2023.

⁶ Catherine Thorbecke, “Sarah Silverman and Authors File Lawsuits against OpenAI and Meta, Alleging Copyright Infringement,” CNN, Jul 10, 2023, <https://www.cnn.com/2023/07/10/tech/sarah-silverman-openai-meta-lawsuit/index.html>, accessed Oct 21, 2024.

⁷ Antonia Pequeno IV, “OpenAI Scores Court Victory After Copyright Lawsuit from Sarah Silverman is Partially Dismissed,” Forbes, Mar 27, 2024, <https://www.forbes.com/sites/antoniopequenoiv/2024/02/13/openai-scores-court-victory-after-copyright-lawsuit-from-sarah-silverman-is-partially-dismissed/>, accessed Oct 21, 2024.

⁸ The Economist, “AI Needs Regulation, but What Kind and How Much?” Forbes, Aug 20, 2024, <https://www.economist.com/schools-brief/2024/08/20/ai-needs-regulation-but-what-kind-and-how-much?/>, accessed Oct 21, 2024.

⁹ Cade Metz, “The New Chatbots Could Change the World. Can You Trust Them?” The New York Times, Dec 10, 2022, <https://www.nytimes.com/2022/12/10/technology/ai-chat-bot-chatgpt.html>, accessed Nov 8, 2023.

¹⁰ OpenAI, “Introducing ChatGPT,” OpenAI Blog, Nov 30, 2022, <https://openai.com/blog/chatgpt>, accessed Sep 8, 2023.

¹¹ Sam Altman, @sama, “ChatGPT Launched on Wednesday. Today it Crossed 1 Million Users!” Twitter, Dec 4, 2022, <https://twitter.com/sama/status/1599668808285028353>, accessed Sep 8, 2023.

¹² Krystal Hu, “ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note,” Reuters, Feb 2, 2023, <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>, accessed Sep 8, 2023.

¹³ Paul L. Garvin, “The Georgetown-IBM Experiment of 1954: An Evaluation in Retrospect,” ACL Anthology, 1967, <https://aclanthology.org/www.mt-archive.info/Garvin-1967.pdf>, accessed Nov 8, 2023.

¹⁴ Grace Solomonoff, “The Meeting of the Minds That Launched AI,” IEEE Spectrum, May 6, 2023, <https://spectrum.ieee.org/dartmouth-ai-workshop>, accessed Feb 22, 2024.

¹⁵ W. David Gardner, “Remembering Joe Weizenbaum, ELIZA Creator,” Information Week, Mar 13, 2008, <https://www.informationweek.com/data-management/remembering-joe-weizenbaum-eliza-creator>, accessed Nov 8, 2023.

¹⁶ Adam Berger, “Statistical Machine Translation,” <http://www.cs.cmu.edu/~ab Berger/mt.html>, accessed Nov 8, 2023.

¹⁷ Alex Moltzau, “The History of Google Translate,” Medium, Jul 22, 2020, <https://alexmoltzau.medium.com/the-history-of-google-translate-fcbe9de3c10e>, accessed Nov 8, 2023.

¹⁸ Larry Hardesty, “Explained: Neural networks,” MIT News, Apr 14, 2017, <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>, accessed Feb 22, 2024.

¹⁹ Jeff Dean, “Using Large-Scale Brain Simulations for Machine Learning and A.I.,” Google Keyword Blog, Jun 26, 2012, <https://blog.google/technology/ai/using-large-scale-brain-simulations-for/>, accessed Feb 22, 2024.

²⁰ Larry Hardesty, “Explained: Neural Networks,” MIT News, Apr 14, 2017, <https://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>, accessed Feb 22, 2024.

- ²¹ Bobby Allyn, "The Godfather of AI Warns of AI Possibly Outperforming Humans," NPR, May 27, 2023, <https://www.npr.org/2023/05/27/1178575886/-the-godfather-of-ai-warns-of-ai-possibly-outperforming-human>, accessed Nov 8, 2023.
- ²² Anas As-Masri and Matthew Urwin, "How Does Backpropagation in a Neural Network Work?," BuiltIn.com, Aug 7, 2023, <https://builtin.com/machine-learning/backpropagation-neural-network>, accessed Dec 19, 2023.
- ²³ Cade Metz, "The Secret Auction That Set Off the Race for AI Supremacy," Wired, Mar 16, 2021, <https://www.wired.com/story/secret-auction-race-ai-supremacy-google-microsoft-baidu/>, accessed Dec 19, 2023.
- ²⁴ Andrew Ng, "Heroes of Deep Learning: Geoffrey Hinton," DeepLearning.AI, <https://www.deeplearning.ai/blog/hodl-geoffrey-hinton/>, accessed Nov 8, 2023.
- ²⁵ Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin, "Attention Is All You Need," Arxiv, Jul 12, 2017 (rev. Aug 2, 2023), <https://arxiv.org/abs/1706.03762>, accessed Nov 8, 2023.
- ²⁶ Giuliano Giacaglia, "How Transformers Work," Towards Data Science, Mar 10, 2019, <https://towardsdatascience.com/transformers-141e32e69591>, accessed Nov 8, 2023.
- ²⁷ Microsoft, "OpenAI Forms Exclusive Computing Partnership with Microsoft to Build New Azure AI Supercomputing Technologies," Microsoft News Center, Jul 22, 2019, <https://news.microsoft.com/2019/07/22/openai-forms-exclusive-computing-partnership-with-microsoft-to-build-new-azure-ai-supercomputing-technologies>, accessed Nov 8, 2023.
- ²⁸ James Vincent, "OpenAI has Published the Text-Generating AI it Said was too Dangerous to Share," The Verge, Nov 7, 2019, <https://www.theverge.com/2019/11/7/20953040/openai-text-generation-ai-gpt-2-full-model-release-1-5b-parameters>, accessed Dec 20, 2023.
- ²⁹ Nefi Alarcon, "OpenAI Presents GPT-3, a 175 Billion Parameters Language Model," Nvidia Technical Blog, Jul 7, 2020, <https://developer.nvidia.com/blog/openai-presents-gpt-3-a-175-billion-parameters-language-model/>, accessed Feb 22, 2024
- ³⁰ Andrew Mayne, "The GPT-3 Zero Shot Approach," andrewmayne.com, Apr 18, 2021, <https://andrewmayne.com/2021/04/18/the-gpt-3-zero-shot-approach/>, accessed Feb 22, 2024.
- ³¹ Krystal Hu, "ChatGPT Sets Record for Fastest-Growing User Base - Analyst Note," Reuters, Feb 2, 2023, <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>, accessed Nov 17, 2023.
- ³² OpenAI, "Introducing ChatGPT Plus," Feb 1, 2023, <https://openai.com/blog/chatgpt-plus>, accessed Nov 17, 2023.
- ³³ OpenAI, "GPT-4," Mar 14, 2023, <https://openai.com/research/gpt-4>, accessed Nov 28, 2023.
- ³⁴ Josh Howarth, "Number of Parameters in GPT-4 (Latest Data)," Exploding Topics, Aug 6, 2024, <https://explodingtopics.com/blog/gpt-parameters>, accessed Oct 24, 2024.
- ³⁵ Claudio Slowik and Filip Kaiser, "GPT-4 vs. GPT-3. OpenAI Models' Comparison," Neoteric. Mar 16, 2023, <https://neoteric.eu/blog/gpt-4-vs-gpt-3-openai-models-comparison/>, accessed Oct 24, 2024.
- ³⁶ Jared Spataro, "Introducing Microsoft 365 Copilot – Your Copilot for Work," Official Microsoft Blog, Mar 16, 2023, <https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/>, accessed Dec 19, 2023.
- ³⁷ Ronen Schwaber-Cohen, "What is a Vector Database & How Does it Work? Use Cases + Examples," Pinecone, May 3, 2023, <https://www.pinecone.io/learn/vector-database/>, accessed Nov 14, 2024.
- ³⁸ Natassha Selvaraj, "What is Retrieval-Augmented Generation (RAG)?" DataCamp Blog, Jan 30, 2024, <https://www.datacamp.com/blog/what-is-retrieval-augmented-generation-rag>, accessed Nov 14, 2024.
- ³⁹ Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun, "Towards Understanding Chain-of-Thought Prompting: An Empirical Study of What Matters," arXiv, Dec 20, 2022, <https://doi.org/10.48550/arXiv.2212.10001>, accessed Nov 14, 2024.
- ⁴⁰ Ken Paul and Anna Tong, "OpenAI Launches New Series of AI Models with 'Reasoning' Abilities," Reuters, Sep 13, 2024, <https://www.reuters.com/technology/artificial-intelligence/openai-launches-new-series-ai-models-solve-hard-problems-2024-09-12/>, accessed Nov 14, 2024.

- ⁴¹ OpenAI, "Introducing OpenAI o1-preview," OpenAI, Sep 12, 2024, <https://openai.com/index/introducing-openai-o1-preview/>, accessed Nov 14, 2024.
- ⁴² Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun, "Towards Understanding Chain-of-Thought Prompting: An Empirical Study of What Matters," arXiv, Dec 20, 2022, <https://doi.org/10.48550/arXiv.2212.10001>, accessed Nov 14, 2024.
- ⁴³ Anna Tong and Krystal Hu, "OpenAI and Rivals Seek New Path to Smarter AI as Current Methods Hit Limitations," Reuters, Nov 11, 2024, <https://www.reuters.com/technology/artificial-intelligence/openai-rivals-seek-new-path-smarter-ai-current-methods-hit-limitations-2024-11-11>, accessed Nov 14, 2024.
- ⁴⁴ Henry Kissinger, Eric Schmidt, and Craig Mundie, "AI's Potential for Humanity and Catastrophe," The Atlantic, Nov 12, 2024. <https://www.theatlantic.com/ideas/archive/2024/11/ai-genesis-excerpt-kissinger-schmidt-mundie/680619/>, accessed November 14, 2024.
- ⁴⁵ Google, "Anthropic, Google, Microsoft and OpenAI Announce Executive Director of the Frontier Model Forum and over \$10 Million for a New AI Safety Fund," Oct 25, 2023, <https://blog.google/outreach-initiatives/public-policy/google-microsoft-anthropic-open-ai-frontier-model-forum-executive-director/>, accessed Nov 14, 2024.
- ⁴⁶ Cade Metz and Gregory Schmidt, "Elon Musk and Others Call for Pause on A.I., Citing 'Profound Risks to Society'," The New York Times, Mar 29, 2023, <https://www.nytimes.com/2023/03/29/technology/ai-artificial-intelligence-musk-risks.html>, accessed Nov 8, 2023.
- ⁴⁷ Cade Metz, "'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead," The New York Times, May 1, 2023, <https://www.nytimes.com/2023/05/01/technology/ai-google-chatbot-engineer-quits-hinton.html>, accessed Nov 9, 2023.
- ⁴⁸ Kevin Roose, "A.I. Poses 'Risk of Extinction,' Industry Leaders Warn," The New York Times, May 30, 2023, <https://www.nytimes.com/2023/05/30/technology/ai-threat-warning.html>, accessed Nov 29, 2023.
- ⁴⁹ Katja Grace, Harlan Stewart, Julia Fabienne Sandkühler, Stephen Thomas, Ben Weinstein-Raun, and Jan Brauner, "Thousands of AI authors on the future of AI," arXiv preprint, Apr 30, 2024, <https://arxiv.org/abs/2401.02843>.
- ⁵⁰ Mark Cartwright, "The Textile Industry in the British Industrial Revolution," World History Encyclopedia, Mar 1, 2023, <https://www.worldhistory.org/article/2183/the-textile-industry-in-the-british-industrial-rev/>, accessed Nov 29, 2023.
- ⁵¹ E. J. Hobsbawm, *The Age of Capital, 1848-1875*, London: Penguin Group, 1988.
- ⁵² E.P. Thompson, *The Making of the English Working Class*, Harmondsworth: Penguin, 1980.
- ⁵³ Grand View Research, "Video Game Market Size, Share & Trends Analysis Report By Device (Console, Mobile, Computer), By Type (Online, Offline), By Region (Asia Pacific, North America, Europe), And Segment Forecasts, 2023 - 2030," Aug 2023, <https://www.grandviewresearch.com/industry-analysis/video-game-market> accessed Nov 15, 2024.
- ⁵⁴ Kate Irwin, "EA Embraces AI, Says 60% of its Game Development Could be Affected," PC mag, Mar 7, 2024, <https://www.pcmag.com/news/ea-embraces-ai-says-60-of-its-game-development-could-be-impacted>, accessed Nov 15, 2024.
- ⁵⁵ Eddie Makuch, "EA Says CFB 25 Would Not Have Been as Good Without the Developers Using AI", Game Spot, July 31, 2024, <https://www.gamespot.com/articles/ea-says-cfb-25-would-not-have-been-as-good-without-the-developers-using-ai/1100-6525428/>, accessed Nov 15, 2024.
- ⁵⁶ Brian Merchant, "AI is Already Taking Jobs in the Video Game Industry," Wired, Jul 23, 2024, <https://www.wired.com/story/ai-is-already-taking-jobs-in-the-video-game-industry/>, accessed Nov 15, 2024.
- ⁵⁷ SAG-AFTRA, "SAG-AFTRA Members who Work on Video Games go on Strike," Jul 25, 2024, <https://www.sagaftra.org/sag-aftra-members-who-work-video-games-go-strike>, accessed Nov 15, 2024.
- ⁵⁸ World Economic Forum, "Future of Jobs Report: 2023," May 2023, https://www3.weforum.org/docs/WEF_Future_of_Jobs_2023.pdf, accessed Dec 5, 2023.
- ⁵⁹ Mauro Cazzaniga, Florence Jaumotte, Longji Li, Giovanni Melina, Augustus J Panton, Carlo Pizzinelli, Emma J Rockall, Marina Mendes Tavares, "Gen-AI: Artificial Intelligence and the Future of Work," IMF Staff Discussion Notes, Jan 14, 2024.
- ⁶⁰ Jorge Tamayo, Leila Doumi, Sagar Goel, Orsolya Kovács-Ondrejko and Raffaella Sadun, "Reskilling in the Age of AI," Harvard Business Review, Oct 2023, <https://hbr.org/2023/09/reskilling-in-the-age-of-ai>, accessed Nov 15, 2024.

- ⁶¹ C.M. Rubin, "Can Universal Basic Income Save Us from the Destabilization of AI and Automation," *Forbes*, Jul 17, 2024, <https://www.forbes.com/sites/cathyrubin/2024/07/17/can-universal-basic-income-save-us-from-the-destabilization-of-ai-and-automation/>, accessed Nov 15, 2024.
- ⁶² *The Economist*, "Cutting the Cord," Oct 7, 1999, <https://www.economist.com/special-report/1999/10/07/cutting-the-cord>, accessed Nov 15, 2024.
- ⁶³ Daron Acemoglu and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33(2), Spring 2019, pp. 3-30, <https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.33.2.3>.
- ⁶⁴ *Ibid.*
- ⁶⁵ *Ibid.*
- ⁶⁶ Sara Brown, "The Lure of 'So-So Technology,' and How to Avoid it," MIT Sloan: Ideas Made to Matter, Oct 31, 2019, <https://mitsloan.mit.edu/ideas-made-to-matter/lure-so-so-technology-and-how-to-avoid-it>, accessed Dec 5, 2023.
- ⁶⁷ Daron Acemoglu and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33(2), Spring 2019, pp. 3-30, <https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.33.2.3>.
- ⁶⁸ Daron Acemoglu, "The Simple Macroeconomics of AI," NBER Working Paper 32487, May, 2024, <https://www.nber.org/papers/w32487>, accessed Nov 15, 2024.
- ⁶⁹ Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Working Paper, Mar 17, 2023 (rev Aug 22, 2023), <https://arxiv.org/abs/2303.10130>, accessed Nov 8, 2023.
- ⁷⁰ *Ibid.*
- ⁷¹ *Ibid.*
- ⁷² *Ibid.*
- ⁷³ Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," NBER Working Paper 31161, Apr, 2023, rev Nov 2023, <https://www.nber.org/papers/w31161>, accessed Nov 8, 2023.
- ⁷⁴ *Ibid.*
- ⁷⁵ *Ibid.*
- ⁷⁶ Fabrizio Dell'Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Lisa Krayer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," HBS Working Paper, Sep 18, 2023 (rev Sep 27, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321, accessed Nov 13, 2023.
- ⁷⁷ *Ibid.*
- ⁷⁸ *Ibid.*
- ⁷⁹ Fabrizio Dell'Acqua, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Lisa Krayer, François Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," HBS Working Paper, Sep 18, 2023 (rev Sep 27, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321, accessed Nov 13, 2023.
- ⁸⁰ *Ibid.*
- ⁸¹ Robert C. Allen, "The high Wage Economy and the Industrial Revolution: a Restatement," *The Economic History Review*, 68(1), Feb 2015, <https://www.jstor.org/stable/43910008>, accessed Nov 29, 2023.
- ⁸² Ronald F. Inglehart and Pippa Norris, "Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash," Harvard Kennedy School Working Paper, Aug 2016, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2818659, accessed Nov 17, 2023.
- ⁸³ Erik Brynjolfsson, "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence," *Daedalus*, Spring 2022, <https://arxiv.org/abs/2201.04200>, accessed Nov 13, 2023.
- ⁸⁴ *Ibid.*

⁸⁵ Ajay K. Agrawal, Joshua S. Gans, and Avid Goldfarb, "The Turing Transformation: Artificial Intelligence, Intelligence Augmentation, and Skill Premiums," NBER Working Paper No. 31767, Oct 2023, https://www.nber.org/system/files/working_papers/w31767/w31767.pdf, accessed Nov 29, 2023.

⁸⁶ Ibid.

⁸⁷ Ibid.

⁸⁸ Ibid.

⁸⁹ Ibid.

⁹⁰ Brian Eastwood, "Sam Altman Believes AI Will Change the World (And Everything Else)," MIT Sloan, May 8, 2024, <https://mitsloan.mit.edu/ideas-made-to-matter/sam-altman-believes-ai-will-change-world-and-everything-else>, accessed Nov 15, 2024.

Notes

A series of horizontal dotted lines for taking notes.

2026

Connective
Leadership Summit
19 - 22 May 2026 | Melbourne VIC

Gamma: Slides in the Blink of AI





REMBRAND KONING
NICOLE TEMPEST KELLER
HYUNJIN KIM

Gamma: Slides in the Blink of AI

In September 2025 Grant Lee, co-founder and CEO of AI slide generator Gamma.app was driving south on San Francisco's I-80 with fellow co-founders, Jon Noronha and James Fox. The trio smiled as they passed Gamma's new billboard—a bold, playful jab at presentation incumbent PowerPoint. (See **Exhibit 1**). But as they continued on their way to their office they saw a host of other AI billboards, including one for Chinese AI startup Manus, with the tagline “Turn meeting notes into pro-level slides.” A few miles later they passed OpenAI's headquarters, fresh off of announcing its ChatGPT Agent which promised to “deliver editable slideshows with ease.” Lee and his co-founders believed Gamma's tools were the best at helping anyone “bring their ideas to life,” but could their five-year old startup continue to thrive in the face of so much startup and incumbent competition?

Gamma, founded in 2020 and based in San Francisco, California, was a consumer-facing, AI-powered communication platform designed to reinvent how people create and share presentations. Operating under a freemium model, Gamma enabled users to generate high-quality, professional decks in minutes with no design experience required. Gamma was part of a new wave of lean AI-native startups, alongside companies like Anysphere, maker of the AI coding tool Cursor; Midjourney, the GenAI image model creator; and Lovable, the vibe coding platform. By leveraging AI across every function—engineering, design, product, customer service, and marketing—Gamma had achieved \$50 million in Annual Recurring Revenue (ARR) in April 2025 with just 30 employees. The company had raised \$23 million to date, with its last raise in 2023—a \$12 million Series A led by Accel, which valued the company at \$112 million at the time. Gamma was profitable and had generated more cash than it had raised—a rare feat among venture-backed startups.

But Lee knew a successful AI venture today could easily be disrupted in months. During the car ride, the team debated what they should do next and how quickly. Noronha pushed for expanding Gamma's product scope to include documents, websites, and social. Fox floated an enterprise play. Lee discussed an API. The team was also considering whether they should raise another round of capital to support these efforts. An investor had offered them \$50 million in investment, and the prevailing wisdom in Silicon Valley was to “raise as much as possible, hire as fast as possible, and scale at all costs.” But Lee wondered if this blitzscaling philosophy that defined the last decade of startups was the right approach in the AI era. As Grant and his co-founders arrived at their destination, they mused on which AI billboards—including Gamma's—would still be displayed on I-80 a decade from now?

HBS Professor Rembrand Koning, Assistant Director Nicole Tempest Keller (California Research Center), and Professor Hyunjin Kim (INSEAD) prepared this case. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2025, 2026 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

This document is authorized for use only by Mitchell Weiss in 2026.

Slide Decks in the Economy

PowerPoint was launched in 1987 by a small software company, Forethought, Inc. before Microsoft acquired it for \$14 million just a few months later—its first significant acquisition.¹ Integrated into Office, PowerPoint took off, becoming the default tool for creating presentations. In 2006, Google introduced Google Slides as part of its Docs suite, offering a cloud-based alternative to PowerPoint. By 2025, the global presentation software market had grown to an estimated \$7.5 billion,² and PowerPoint and Google were reported to each have around 500 million MAUs.³

According to a 2020 Nielsen study, the average professional worker spent about seven hours per week building PowerPoint slides.⁴ Firms adopted a range of approaches to manage slide production. Many organizations built in-house capabilities, hiring teams of graphic designers or visual communication specialists. McKinsey & Company, for example, employed a global team of Visual Graphics Specialists, with average compensation of around \$85,000 in 2025.⁵ More broadly, hundreds-of-thousands of in-house graphic designers in the U.S. earned between \$64,000 and \$84,000, depending on role and experience.⁶

Other companies chose to outsource slide production, either to design agencies or Business Process Outsourcing (BPO) firms overseas. U.S.-based agencies often charged between \$1,000 and \$3,000 for a 10-slide deck, sometimes requiring a \$1,500 minimum.⁷ Freelancers, particularly those outside the U.S., offered more affordable options, with rates ranging from \$100 to \$1,500 for a 10-page deck.⁸

Background on Gamma: The Founding Story

Gamma's founders had all previously worked together at Optimizely, a platform that offered tools for A/B testing: Lee had served as Director of Strategy and Operations, Noronha had led Product Management, and Fox had been Director of Customer Success. Optimizely had been a VC darling, raising \$200 million from top VCs with a peak valuation of around \$600 million after its Series D in 2019.^{9,10} However, after a series of strategic and operational missteps, it ended up being acquired by Episerver (a portfolio company of private equity firm Insight Partners) in 2020. Following the sale, the trio left the company, with their equity underwater, and began brainstorming startup ideas. All three founders had struggled with the time-consuming process of building slide decks in their prior roles—Lee during his time in consulting and investment banking, Noronha in product roles at Microsoft and Optimizely, and Fox during his years in consulting. Their collective experiences highlighted how outdated and rigid traditional presentation tools had become in the context of modern, collaborative work. And yet slide decks were still critical to how business was done. Lee commented: "Slides and presentations end up becoming the language of business. It's how ideas and information spread, and it's often how decisions are made."¹¹

Working out of a converted two-bedroom apartment in San Francisco the team began rethinking how ideas could be shared and how content creation could feel effortless. In 2020 the team turned to Afore Capital for pre-Seed funding to get started. Gaurav Jain (HBS 2013), Managing Partner of Afore Capital, discussed his interest in Gamma:

The team had seen that the slide deck format had been largely unchanged since PowerPoint launched in the 1980s. By 2020 that format seemed outdated for a collaborative, digital-first world. So they wanted to reimagine that. I had seen a similar

pitch in 2012 when I invested in *Airtable*^a—their pitch was that Excel hadn't really changed since it was launched—it's great for two-dimensional data structure, but if you want to drop a file in there, it's just impossible to do that. So when I heard the Gamma pitch I thought clearly this is going to be very challenging to build a horizontal product, but I think somebody's going to figure this out, and given the quality of the team and their knowledge of product, this is the right team to do it.

Lee's entrepreneurial ambition was rooted in his family upbringing. "My parents moved to the U.S. from Taiwan and launched a few restaurants. When that didn't work, they took jobs at the postal service where they worked for 30 plus years. Watching them work so hard made me realize, if you want to do anything great, you've just got to put in that work."

All three founders were in their 30s, making them older than many tech founders. However Lee saw this as an advantage for the team: "Experience helps you avoid costly mistakes. Industry connections open doors that would otherwise stay closed. Experience-driven decision-making prevents expensive, naïveté-driven detours. Effective time management creates clarity, not overwhelm and waste."¹² Lee had seen the benefits of being a parent on his time management: "The time constraints of family life haven't made me less effective; they've made me more intentional. When I have one hour between putting kids to bed and my own sleep, I'm hyper-focused on what actually moves the needle. That urgency of knowing time is precious is a superpower ... not a limitation."¹³

The Transition to AI

By late-2022, the 12-person team felt stuck between product-market fit and funding pressure.¹⁴ While they were approaching 60,000 users, most used it only once.¹⁵ Noronha explained, "We started to worry that we only had a little over a year of runway left. But that turned out to be a really good time to be looking for a new plan because that's also when the new wave of Generative AI was just starting to come out."

Initially, their attention was drawn to GenAI image models since presentations were typically full of images and illustrations. They held a company-wide hackathon to see what people could create with the early models. Head of Design, Zach Leach, recounted, "We tried to make a theme generator. It took a few rounds of tweaking it to make it feel right, but we were like, 'Yeah, we could really make this happen.'" Diana Kimball Berlin (HBS 2013), later the Head of Core Product, saw that GenAI enabled an analog to the longstanding computing concept of "Tools for Thought." She explained in a blog post in 2022: "The coming wave of tools for imagination will unlock the giddy, gritty experience of creative collaboration.... To collaborate directly with a sort of collective imagination can embolden and extend our own insight and creativity."¹⁶

OpenAI launched ChatGPT in November 2022, reaching 57 million users in its first month, and 100 million within 60 days. Noronha commented, "We saw that AI wasn't just moving fast—it was accelerating. And if that was true, most people would underestimate how far and how quickly things would evolve. So unless AI is deeply integrated into your product from the ground up, you risk falling behind and never catching up." That urgency led to a bold decision. Noronha continued, "We didn't want to bolt AI on—we wanted to embed it in everything: from the first user interaction to what we can imagine the editing experience to be in the future." But they stopped short of building their own

^a *Airtable* was a cloud-based platform that combined the flexibility of a spreadsheet with the power of a relational database, enabling users to organize, manage, and automate data and workflows without needing programming skills. It had an estimated \$478 million in ARR in 2024 and was used by 500,000 organizations. By 2025 it had raised \$1.36 billion and was still private.

models, instead choosing to build on top of existing foundation models. Noronha explained, “Rather than spend money trying to compete with the foundation models, we made the decision to use them as suppliers and just focus on our end users.”

The pivot to AI reframed Gamma’s value proposition: it became a solution to the “blank page syndrome” that plagued users sitting down to start a presentation. Noronha explained:

AI unearthed the real problem to be solved, which is this blank page problem: ‘Hey, I have a big presentation due tomorrow. I need to get started.’ And you open up PowerPoint, you’re looking at a blank page, and you’re like, ‘Oh no, what do I do about this?’ And so when we solved that blank page problem, that was when we unlocked the radioactive product market fit.

This new vision informed Gamma’s mission: “to make content creation feel effortless – whether you have design skills or not,” thereby democratizing design. (See **Exhibit 2.**) Gamma launched its AI-powered presentation tool in March 2023 on Product Hunt. The demo video, which showed a complete deck being generated in real time, went viral.^b Kristin Fracchia, Head of Marketing, commented, “It had that wow factor of seeing a presentation built before your eyes.” Users flocked to try it, and Gamma gained 3 million users in three months¹⁷ and hit 11 million by the end of 2023.¹⁸

The Gamma Customer and Product

Gamma’s core user was a “prosumer” – someone who sat between a casual consumer and a professional. This segment included consultants, freelancers, and business team leads who needed high-quality output but did not have access to a designer. Gamma gave them the ability to create professional-looking content on their own, powered by AI, but guided by human editing and design sensibilities.¹⁹ Gamma’s user base was heavily international – only 20% of users were in the U.S. – with significant traction in Brazil, Korea, France, and Japan. Localization was a priority early on, with Gamma supporting nine languages beyond English.

Gamma’s user journey began with a simple prompt. For example: “Create a 6-slide deck on the case method at Harvard Business School.” Gamma first provided an AI-generated outline and suggested a visually appealing design theme. Then it created a draft of the “cards” (Gamma used the term “cards” instead of “slides”) in less than a minute, which users could customize. (See **Exhibit 3.**) Users could share a link or multiple users could collaborate on the deck in real time. Final decks could then be shared via URL, exported as PDFs or PowerPoint files, or embedded on websites or social platforms.

As a product-led growth (PLG) company, product excellence was critical to winning and retaining customers. “We’ve bet everything on making the best product,” said Noronha. That meant obsessing over every detail, from design to onboarding speed, and being highly responsive to user feedback. The onboarding process was especially critical: Gamma aimed to deliver visible value in the first 30 seconds to earn the right to keep that customer engaged for the next 30 seconds – shortening the time to value as much as possible. Everyone at Gamma was focused on customer needs, starting at the top. Hilary Dudeck, Head of Customer Experience, explained, “Grant reads and replies in the ‘Gambassador’ Slack (a private Slack group for top Gamma users) every day.” Gamma also gathered feedback from users after it created each slide deck with a short prompt at the end: “How satisfied are you with the output?” with three answer options: “It’s bad,” “It’s OK,” “It’s great.” Gamma followed up with, “How could

^b Gamma. “Gamma: AI for presenting ideas, beautifully,” YouTube, published March 26, 2023, <https://www.youtube.com/watch?v=r9VOYJM6nZc>, accessed December 17, 2026.

we make it better?” This feedback loop shaped product decisions which had led to improvements in Gamma’s “great” ratings. (See **Exhibit 4a** and **4b**.)

The strong focus on product also stemmed from the nature of presentations themselves—they were inherently made to be shared. And the better they looked, the more likely users were to share them. This created a powerful PLG flywheel: a great product led to more sharing, which brought in more users, which fueled more feedback and improvements to the product.

Gamma’s Business Model

Gamma followed a freemium model with multiple tiers: a free plan that offered basic AI image models and a maximum of 10 cards per generation, the Plus plan (\$96/year) with advanced AI image models and up to 20 cards per generation, and a Pro plan (\$216/year) with premium AI image models, detailed analytics, and up to 50 cards per generation. There were also higher level plans: Ultra, Team, and Business. (See **Exhibit 5**.) In 2025, Gamma had 500,000+ paid users, an estimated 90% of whom were on the Plus plan. (See **Exhibit 6**.)

The free tier served as an organic growth engine: presentations carried a “Made with Gamma” badge in the footer, creating built-in virality which allowed Gamma to scale without relying heavily on paid marketing. More than 50% of inbound traffic came from branded search. Lee described Gamma’s marketing spend as “relatively low,”²⁰ mostly involving influencer marketing and newsletter marketing, with more recent forays into outdoor advertising.

Cost Drivers: Tokenization

A token was the fundamental unit of text that LLMs used to break down and process language. OpenAI and Anthropic defined a token as approximately four characters of English text, with the average word equating to about 1.3 tokens.²¹ In contrast to text, image generation was far more GPU-intensive. When measured in token-equivalent compute terms, generating a single image required the same computational effort as generating 1,000 to 5,000 words of text.²²

There were two types of tokens—input tokens and output tokens. Input tokens referred to all the text the LLM consumed *before* generating a response. For Gamma this could include the user’s prompt, Gamma’s internal house style (i.e. “present in bullet form,” “limit jargon,” etc.) and contextual memory (i.e. prior slide content, theme, and metadata.) Output tokens captured everything the model *generated in response*—not only the visible text that appeared on the slide, but also the underlying JavaScript Object Notation (JSON) structure that defined layout, formatting, and metadata. In cases where Gamma included an image, output tokens also included the instructions the LLM sent to an image generation model, like Ideogram or Flux, to create an image. Together, input and output tokens determined the total token usage per slide, which directly influenced overall compute cost.

AI Models

Gamma ran on over 20 different AI models, including Anthropic’s Claude; OpenAI’s GPT models, DALL-E, and gpt-image-1; and Google’s Gemini and Imagen. In 2023, Gamma entered into a partnership with Anthropic after testing and adopting the Claude 3 Haiku model for its AI launch.²³ The models varied widely in token cost (see **Exhibits 7a** and **7b**) and latency, forcing Gamma to make tradeoffs. Sometimes a cheaper, faster model sufficed; other times, higher-quality output justified the extra cost. Noronha discussed, “Maybe we only want to use our best models for our paying customers, except then there’s this conversion rate problem—you want your free users to see the potential.”

To make these decisions, Gamma ran hundreds of A/B tests with real users whenever a new model was released. Users were asked to rate the outputs directly, but the team also tracked behavioral signals: how much people edited the slides, whether they exported or shared them, and whether they ultimately converted to paying customers.²⁴ According to a report from a16z, inference costs for an LLM of equivalent performance were decreasing by 10x every year.²⁵ (See **Exhibit 8**.)

Gamma's System Architecture

Gamma's product architecture involved orchestrating different foundation models. Behind the scenes, a prompt-orchestration layer sequenced multiple models, each chosen for its strengths—some better at reasoning, others at creativity or cost efficiency. The system chained steps together, moving from user prompt to deck outline, to slide text, to images, to layout, and then polished the result for coherence. A feedback layer ensured that user edits and interactions flowed back into the system, for continuous optimization, teaching it which prompts, layouts, and models worked best. The model routing layer adaptively selected the right model to use based on each user's intent and context, dynamically optimizing for quality, cost, and latency. For example, a middle school student creating a presentation on frogs might not be routed to the highest priced model.

Other costs

Gamma's only other variable cost was cloud infrastructure—including compute, storage, and networking. Industry benchmarks suggested that for a typical SaaS company, cloud infrastructure costs usually accounted for 6% to 12% of revenue.²⁶ Overall, Gamma had been able to keep gross margins in the 60% to 80% range, and the company had been profitable since early 2024.²⁷

Gamma's Use of AI Tools

From the start, Gamma cultivated an AI-positive culture. Lee discussed, "Every single person should be thinking about building in a way that incorporates AI as part of the plan—not just a nice to have, it's a must have."²⁸ The result was that Gamma had achieved outsized leverage from a small team. As of September 2025, Gamma had 50 people on staff. According to Carta, the average startup salary for marketing, customer success, and sales functions in 2024 was \$155,000.²⁹ The average salary for product and engineering roles was \$190,000. Lee commented, "If we were from the generation before, we would easily be at 200 employees." Indeed the "old" Silicon Valley playbook for VC-backed startups was to raise a lot of capital and then spend it to hire an army of employees to scale up fast. The priority was growth over profits. In February 2025 *The New York Times* reported on the shift in mindset in Silicon Valley from a few decades ago:

Head count and fundraising were badges of honor among founders, who philosophized that bigger was better. But Gamma is among a growing cohort of start-ups, most of them working on A.I. products, that are also using A.I. to maximize efficiency. They make money and are growing fast without the funding or employees they would have needed before. The biggest bragging rights for these start-ups are for making the most revenue with the fewest workers.³⁰

Indeed, ARR per FTE was becoming the key investor metric to show how capital efficient startups were. Sam Altman, CEO of OpenAI, went so far as to predict that there could someday be a one-person company worth \$1 billion.³¹ Industry followers even tracked ARR:FTE efficiency on a Lean AI-Native Company Leaderboard. (See **Exhibit 9**.)

Gamma saw many advantages to having a small team, from better communication to faster decision making and execution. Lee explained in a post, “This generation of startups is asking, “What if smaller is actually better? What if you can move faster with a lean, focused team rather than a bloated organization where communication breaks down and decision-making slows to a crawl? What if instead of raising \$50M, you raise just what you need and focus on unit economics from day one?”³²

Engineering: Gamma’s engineering team had just 17 engineers, yet it operated at a pace that rivaled much larger teams. This was due in large part to the strategic use of AI tools like Cursor, Claude, and Windsurf, which helped engineers scaffold code, prototype quickly, and accelerate deployment timelines. The company estimated that AI tools gave them a 30 to 40% productivity advantage across the engineering function.

However, AI coding was not perfect. Research showed that some AI-generated code could insert security flaws in applications, cause outages, or make mistakes, forcing developers to change the code or debug heavily.³³ Gamma understood these risks and did not use AI coding tools to write unreviewed code yet. Fox explained:

What non-engineers don’t realize is one bad line of code can take down your entire application, and that line of code may not get hit until 3:00 am when Japan wakes up and tries the new thing and then boom, it’s down. AI coding is advancing pretty quickly but if there were ‘1 through 5 autonomy levels’ for AI coding, like with autonomous vehicles, it feels like we’re crossing from level 1 into 2, not 4 into 5. If you’re doing things that fit a more well-known pattern, like a landing page, it’s good. But there’s still a lot of nuance to writing code and understanding it, because you’re leaving code for the next developer to look at and the next developer’s AI to look at.

Gamma’s use of third-party models also helped keep its engineering team small. Noronha explained, “One benefit of AI models being so good is that we haven’t had to build a huge team around AI. We’ve been able to focus more on productizing the experience.”

Design and Product: At Gamma, design and product went hand in hand. Rather than treating design as a downstream function, the company embedded it at the core of product development. Consequently, the two functions were combined into a single, integrated team with 11 people. Everyone in Design and Product used AI tools. Deeni Fatiha (HBS 2018), Head of Product for AI, explained, “Our engineers are very product-y. Our product people are quite technical. So are our designers. Nobody has the luxury of being like, ‘This is my swim lane and I’m just going to swim in it.’”

Gamma found that AI was especially useful for prototyping. Leach commented, “We’re a design team that speaks in code a lot of the time to help with prototyping. What I’m seeing now is that if you have a design team that’s empowered through tools like Bolt and Cursor, you’re going to be able to build more interesting prototypes that feel more real, faster.” Fatiha discussed, “We have a very prototype-heavy culture. All our designers can vibe code, if not code-code. And the PMs are all using AI to prototype as well to get to decision making faster.”

The product team also used AI to synthesize unstructured, freeform product feedback that users provided after giving their product rating. As Fatiha explained, “The freeform text box really allows us to dig a layer deeper, and the only way we can do that at scale is by using an LLM to group the feedback into clusters, create summaries, and flag recurring pain points.”

Customer Support: By 2025 over 700,000 creations were made per day and Gamma supported over 50 million users across the globe. (See **Exhibit 10**.) Gamma's two-person support team managed over 32,000 support inquiries per month. To handle this volume of work, Gamma contracted out support to customer support platform, Intercom. Gamma used Intercom's conversational AI bot, Fin, to handle user support inquiries. Dudeck explained, "Fin handles 86% of all incoming support tickets." (See **Exhibit 11**.) Gamma was intentional about supplying Fin with the best, up-to-date information. Dudeck explained, "Our help center is hosted by Intercom too, so Fin automatically crawls that content. We also write articles specifically for the bot, called snippets."

AI also helped Gamma maintain a global support presence without having to staff local teams. Noronha noted, "The majority of our users don't speak English, but we don't have an international support team. AI enables seamless multilingual support."

Marketing: Gamma had just two people on its marketing team. Fracchia explained, "Instead of needing a team of 10, we can operate a much smaller staff and use AI tools for the rest." The company also partnered with outside vendors for specific tasks.

Gamma used AI tools such as Google's Deep Research and NotebookLM to synthesize vast amounts of customer data and feedback. Fracchia explained, "We are able to pull from a much broader range of qualitative data points and what comes out is a clearer, more precise picture of who our users are, who we should be marketing to, where the opportunities are, what their pain points are, and how they're discovering Gamma."

AI also powered Gamma's approach to data enrichment. Most of Gamma's users signed up with personal emails, meaning that Gamma had limited insight into where users worked, what roles they held, or whether they had influence in their organizations or online communities. However, using data enrichment tools like Clay, Gamma could enrich those profiles automatically. Gamma could also use this data to help identify potential influencers to invite into the Gambassador program or encourage them to post about Gamma on social. As Gamma considered moving into B2B, this same approach could give the team visibility into where traction was already developing inside organizations. Fracchia explained:

People sign up with Gmail or Yahoo accounts, but they work at real companies. We can run queries against LinkedIn and other public sources and match people to roles and companies. It's not perfect coverage, but we can probably get 60 to 70% coverage in understanding who these people are and where they work. Ten years ago I used to hire people on Upwork to search names and fill spreadsheets. Now I can automate most of that with AI. And it's not just faster—it's more accurate and repeatable.

Talent Strategy and Culture in an AI-Native Company

Gamma's heavy use of AI reshaped how the company thought about hiring—not just how many people were hired, but the types of people it looked for. Gamma knew that the more quickly it hired, the more it risked talent dilution. With AI automating many lower-level tasks, Gamma focused on hiring high-leverage individuals who could do more with less.

The company specifically sought out experienced professionals who could also operate as strong individual contributors (ICs). Lee commented:

For decades, companies hired managers to manage, and ICs to execute. This artificial separation created leaders who gradually lost touch with work they were meant to

oversee. Those people are going to suffer in the era of downsizing because they cannot actually emulate or replicate their employee work anymore. At Gamma, we want player coaches—people who have management experience but aren't so far removed from the IC work that they can't do it anymore.

Gamma also prioritized hiring generalists—people who could work across functions and adapt quickly to the evolving landscape of AI tools. Lee explained:

The most valuable employees aren't specialists who excel in narrow domains—they're resourceful problem-solvers who continuously expand their capabilities. This approach creates remarkable resilience. Since everyone understands multiple functions, we don't have single points of failure when someone leaves or moves to another project.³⁴

Gamma prided itself on a tight-knit culture of trust, speed, creativity, and flexibility. Gamma did not adhere to the 9-9-6^c culture typical of many AI startups and allowed for hybrid work. Lee explained, "We do not have a culture of being a boiler room where the pressure is just constantly there. We're trying to build for the long haul, which means we're constantly trying to balance how much creative energy you can bring."

Given the small and tight knit team at Gamma, "mutual fit" was very important. As a result, the company had moved away from traditional interviews to a more immersive approach: paid work trials to ensure every hire was on-culture and could thrive in Gamma's AI-native approach. Lee explained, "Candidates perform the actual role they'd take permanently. It creates a level of transparency that no interview ever could." Gamma's focus on culture fit had paid off—there had been virtually no turnover.

Gamma's Product and GTM Initiatives

By mid-2025 a host of competitors had also entered the AI presentation market. AI model providers like Manus and OpenAI had launched their own presentation tools. (See **Exhibit 12a** and **Exhibit 12b**) Design platforms like Canva and Figma had joined the race, integrating AI into their established ecosystems. Meanwhile, AI-native startups, such as Beautiful.AI, were aggressively innovating to capture share. (See **Exhibit 13**.) Yet Gamma positioned itself differently. Berlin explained:

What sets Gamma apart is our obsession with the full journey—from idea to impact—through visual storytelling. It's not just about organizing information; it's about crafting a narrative that sparks emotion, prompts discussion, and makes people feel proud and more confident to share their work. We've spent years refining prompts, visuals, and user experiences so that every piece created in Gamma doesn't just inform, it inspires. That emotional connection is what leads to real impact, and that's something most platforms overlook.

To continue to stay ahead of the competition, Gamma was exploring several product and GTM initiatives.

^c 9:00am to 9:00pm, 6 days a week.

Broaden Product Scope

One of Gamma's core strategic questions was how to evolve the platform beyond episodic use cases. Noronha explained:

No company has won by focusing on presentations alone. One of the challenges with presentations is there's a lot of episodic usage. People will have some very high-stakes presentation they have to make, but then they churn. If you're part of Office suite, that's no problem – you're paying for Excel and Word at the same time. But if you're standalone, that's a lot harder.

By broadening the product scope, Gamma aimed to increase long-term retention. Noronha explained:

We have a goal around frequency of use. We want to become a platform that people are using about eight days out of every month. We know from our metrics that's a great predictor of retention. One way to get there is by giving users the ability to do more than just presentations. That means documents, websites, social media. And I think there are other formats we could go after in the future – email, newsletters, blogs, even video.

Gamma had already taken its first step in this direction, launching a website creation feature which let users create and publish sites with a simple AI prompt.³⁵ It also introduced Gamma Agent, an AI editing assistant that allowed users to make complex, intent-driven edits across entire documents or presentations using natural language (i.e. "Make this more visual," "Add an executive summary.") Unlike Gamma's earlier prompt-based tools, Agent understood the full context and structure of a document, coordinating multiple AI models to rewrite, restyle, and reformat content while showing users each change it made. Chaemin Ahn, Senior Product Designer for AI and Innovation, commented, "Agent was designed with several jobs in mind, but particularly we wanted it to be context-aware and meet users where they are in their editing journey. We also wanted to make the work that it conducts transparent so users can build trust with it, correct it, and adapt to it as well."

Develop an API

To increase top of the funnel discoverability, Gamma was in the early stages of developing an API that would allow developers or companies to access Gamma's content generation engine directly. Instead of relying solely on users to engage with Gamma's web interface, the API would enable Gamma to operate as a white-labeled or embedded service, powering outputs behind other platforms like Claude or Salesforce.

Developing an API would open up high-volume, programmatic use cases, which was not possible through manual use alone. Lee explained, "It would help us explore asymmetrically how we compete with PowerPoint. We can have other people build on top of Gamma, use their first-party data, combine it with our content engine, and use Gamma in their own unique ways." Gamma saw this as a path to new forms of top-of-funnel visibility. Instead of relying solely on traditional SEO, Gamma could benefit from GEO (Generative Engine Optimization) ensuring Gamma's capabilities surfaced when users prompted models like Claude, ChatGPT, or Perplexity.

Demand for the API was already clear. Fatiha explained, "We've gotten so much interest for the API. It's one of our top requested features for users. Given the energy that we've felt around it, it could be a huge money maker for us and give us access to so many customers." Still, some team members acknowledged the risks of going too far. Leach explained, "There are big product strategy questions

around the API. In a world where we literally give you a fully finished deck, you might never have to go to Gamma. Is that good for us?" There were also questions about whether the GTM motion was too much of a stretch for Gamma. Fatiha explained, "The API is a B2B motion. It requires a sales team, it requires support. And the reason we've been able to be small and nimble is because we're prosumer-focused."

Pursue B2B

Gamma was also considering offering a B2B enterprise product. Companies could purchase licenses for their employees, with volume discounts. There were already 40,000 "workspaces" (defined as 3+ users) using Gamma, suggesting early signs of team adoption. Gamma aimed to solve two main pain points for businesses. First, of the seven hours per week that the average professional worker spent on PowerPoint, 37%, or 2.6 hours, was spent on non-value added formatting of charts, diagrams, and slide layouts.³⁶ Second, nearly half (47%) of all PowerPoint presentations failed to comply with corporate design guidelines, resulting in a major loss for brand communication.³⁷ Berlin explained:

Business leaders see real value in Gamma because it enables effortlessly on-brand visual storytelling externally, and internally it saves time on manual formatting. There is much more demand for visual storytelling than has been satisfied by the production of visual storytelling today, because if it takes *hours*, a business has to think about opportunity costs. But if it takes *minutes*, all you have to think about is impact.

A B2B offering would enable Gamma to generate significantly higher Annual Contract Value (ACV), more predictable recurring revenue, and reduce churn. Selling to known companies also added brand credibility and opened the door to expansion across departments—a classic "land and expand" strategy. But Gamma was cautious. Noronha was especially concerned that moving to a B2B model would lead to a slower feedback loop. He explained, "When you're product-led, every time you make the product better, you see it in your conversion rate. And when you make the product worse—sometimes accidentally—you feel the pain immediately. It's like touching a stove. But in sales-led organizations, it can take years to know if you're even pointed in the right direction."

Building a Sustainable Advantage

Gamma had already accomplished what few startups ever did: rapid growth, profitability, and lifetime negative net burn.^d Its revenue had grown 8x in the past year.

And yet the competitive landscape was shifting quickly. Gamma's AI generated slides were well ahead of ChatGPT (see **Exhibit 12b**), but would ChatGPT quickly catch up with the launch of GPT-5? Would improvements to Microsoft's Copilot and Google's Gemini finally enable PowerPoint and Slides to replicate Gamma's magic? Even though Gamma had generated more slides for more users than other startups, companies like Manus and Beautiful.AI were racing to catch up. And while Gamma itself benefited from improvements in foundation models and AI tools like Lovable and Cursor, these improvements also benefited other entrepreneurs. Would Gamma's AI-first and user obsessed culture give it an edge against the fact that improved AI tools were likely to unlock thousands of tiny AI-native competitors? As Lee, Noronha, and Fox stepped out of the car, they asked: did Gamma have a sustainable edge?

^d Gamma had more cash in the bank than it had raised.

Exhibit 1 Gamma Billboard



Source: Company documents.

Exhibit 2 Gamma Product Principles

Democratizing good design

- ▼ **Everyone should be able to create beautiful content**
 - We trust and remember things better when they are well-designed
 - But most people struggle to shape their content, which leads to content that is poorly communicated
- ▼ **We can't just build a better design tool**
 - Design tools¹ are built for designers
 - What would a design tool for "non-designers" look like?
 - Our vision is to build an all-in-one content creation platform that anyone can pick up and mold to their needs. How we will win:
 - AI-native, writing-first
 - Interoperable & malleable
 - Integrate / ingest and synthesize data from any source, file format, or platform (notes, wiki, project management, crm, analytics, etc.)
- ▼ **AI as a Design Partner**

Throughout the content creation process:

 - Synthesis and generation (endless permutations across formats)
 - Visual design, look & feel
 - Tone, spelling, grammar
 - Multilingual translation
 - Information density²
 - Sourcing content³


Source: Company documents.

Exhibit 3 Sample Gamma Slide Output


The Harvard Business School Case Method: Redefining Learning by Doing

For over a century, Harvard Business School has transformed business education through an innovative approach that throws conventional teaching out the window. Instead of lectures and textbooks, students grapple with real business dilemmas, stepping into the shoes of executives facing high-stakes decisions.

This is the case method—a revolutionary pedagogy that doesn't just teach business theory, but builds the judgment, confidence, and leadership skills needed to navigate an uncertain world.



The Origin Story: How HBS Revolutionized Business Education



In 1921, Harvard Business School wrote its first teaching case, creating a fundamentally new way to develop leaders. This wasn't about memorizing theories or formulas—it was about confronting the messy reality of business decision-making.

The case study method quickly became the backbone of the HBS curriculum, establishing a model that would be emulated by top business schools worldwide. Rather than passive absorption of knowledge, HBS pioneered an active learning approach centered on real decision points faced by real executives.

This bold experiment asked a simple question: How do you prepare someone to lead when leadership can't be taught from a textbook? The answer lay in experience—simulated but intense, academic yet visceral.

Anatomy of a Case: Real Decisions, Real Stakes

The Narrative

Each case is a 10-20 page narrative presenting a true business dilemma faced by a specific executive or organization, rich with data, context, and competing priorities.


The Challenge

Students step into the shoes of leaders, analyzing incomplete information to confront complex challenges with no predetermined "right" answer—just trade-offs and consequences.

The Examples

Iconic cases include Google's Project Oxygen ("Do Managers Matter?"), BTS's global breakthrough strategy, the Army Crew Team rivalry, and GE's transformation under Jack Welch.


Exhibit 3 (continued)



Shockingly Diverse: Cases from Google to K-Pop

The brilliance of the case method lies in its versatility. HBS cases span industries, cultures, and eras—each offering unique lessons about leadership, strategy, and human behavior. Here are four legendary examples that showcase the method's remarkable range:



<p>01</p> <p>Google's Project Oxygen</p> <p>Google's data-driven revelation that great managers do matter—overturning the tech giant's engineer-centric culture and proving that leadership can be measured, taught, and improved.</p>	<p>02</p> <p>BTS's Global Phenomenon</p> <p>Big Hit Entertainment's strategy for overcoming cultural and language barriers, building authentic fan connections, and transforming seven Korean artists into world-beating cultural icons.</p>
<p>03</p> <p>Army Crew Team</p> <p>When quantifiable talent failed to deliver victory, this case explores how culture, teamwork, and psychological dynamics can trump individual excellence in high-performing teams.</p>	<p>04</p> <p>GE's Transformation</p> <p>Jack Welch's legendary tenure at General Electric raises the ultimate succession question: Can legacy growth be sustained after a transformational CEO departs? What systems outlast the individual?</p>



The Power of Cold Calls

"You're the CEO now —What do you do?"

This is the moment that defines the HBS experience. No preparation can fully ready you for the adrenaline surge when your name is called. The professor's eyes lock on yours. Ninety classmates turn to listen. You have seconds to synthesize hours of analysis into a coherent opening stance. This isn't a simulation—this is where confidence is forged, where communication skills are pressure-tested, and where the distance between thinking and doing collapses entirely. The cold call doesn't just assess knowledge; it builds the muscle memory of leadership under fire.





Enduring Impact: Transforming Minds, Shaping Leaders

50K+

HBS Alumni

Trained by the case method, now leading as CEOs, founders, and change-makers worldwide



100K+


Years Strong

Over a century of refining and perfecting the case method approach

The case method's influence extends far beyond business schools. HBS's teaching approach has been embraced in law, medicine, public policy, and entrepreneurship—shaping decision-makers across virtually every professional domain. The methodology has proven that experiential learning, grounded in real-world complexity, produces leaders who can think critically and act decisively.

But the true measure of impact isn't in the numbers—it's in the transformation. Alumni consistently report that the case method changed how they think, how they lead, and how they approach uncertainty. It's more than a collection of lessons; it's a training ground for facing the unknown, leading through ambiguity, and turning insight into action.

The case method doesn't promise easy answers. It promises something far more valuable: the judgment to navigate complexity and the courage to decide.



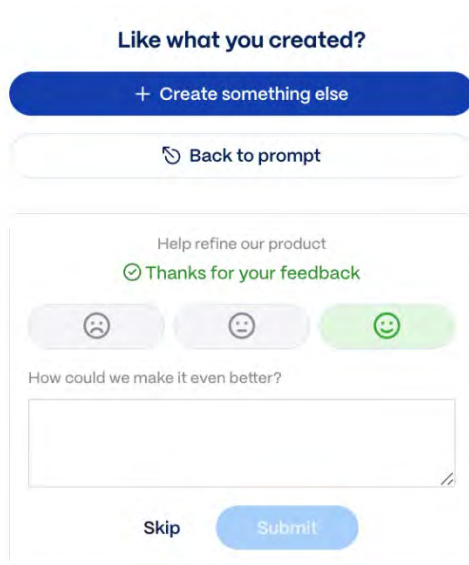
Source: Gamma.app.

14

This document is authorized for use only by Mitchell Weiss in 2026.

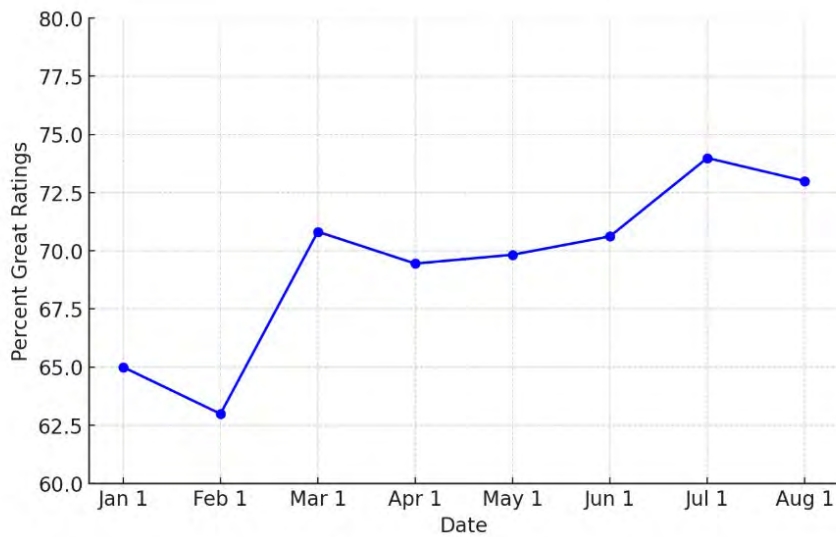
This article is made available to you with compliments of Connective Leadership Summit, Innovate to Lead for your personal use. Further posting, copying or distribution is not permitted.

Exhibit 4a User Feedback Reactions and Comments



Source: Gamma.app.

Exhibit 4b Percent of "Great" Ratings Over Time on Gamma's "Generate" Product



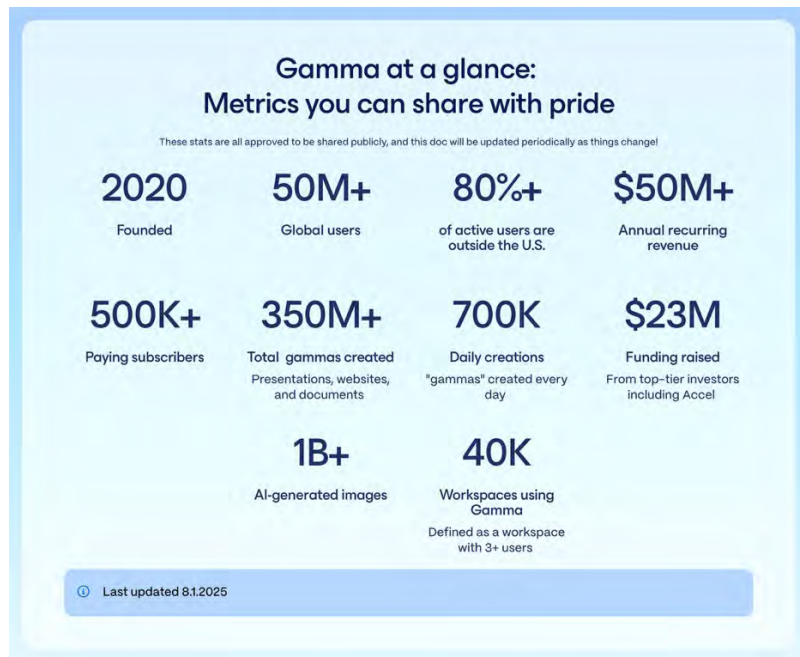
Source: Company documents.

Exhibit 5 Gamma Pricing Plans

Category	Feature	Free	Plus	Pro	Ultra <small>(Introductory Price)</small>	Team	Business
Pricing	Monthly Price	\$0	\$10	\$25	\$100	–	–
	Annual Price	–	\$96	\$216	–	\$240/seat (2 seat min)	\$480/seat (10 seat min)
AI Usage	Unlimited AI Creations	–	✓	✓	✓	✓	✓
	Monthly Credits	400 (at signup, not renewed monthly)	1,000	4,000	20,000	6,000	10,000
	Credit Usage	All AI Actions	Agent	API + Agent	API + Agent + Ultra Image Models	API + Agent	API + Agent + Ultra Image Models
	Credit Rollover	–	Up to 2× monthly limit	Up to 2× monthly limit	Up to 2× monthly limit	Up to 2× monthly limit	Up to 2× monthly limit
	Card Generation Limit	10	20	50	75	50	50
	Input Tokens	50k	100k	100k	100k	100k	100k
AI Models	Image	Basic	Basic + Advanced	Basic + Advanced + Premium	Basic + Advanced + Premium + Ultra	Basic + Advanced + Premium	Basic + Advanced + Premium + Ultra
Customization	Remove Gamma Branding	–	✓	✓	✓	✓	✓
	Custom Fonts	–	–	✓	✓	✓	✓
	Headers & Footers	–	–	✓	✓	✓	✓
	Password Protection	–	–	✓	✓	✓	✓
	Custom Domain Limit	–	–	10	100	10	100
	Search Indexing	–	–	✓	✓	✓	✓
Analytics	Analytics	–	–	Detailed analytics	Detailed analytics	Detailed analytics	Detailed analytics
API	API Access	–	–	✓	✓	✓	✓
Security	Data Controls	Opt-out available	Opt-out available	Opt-out available	Opt-out available	Locked opt-out	Locked opt-out
	SOC 2	–	–	–	–	–	✓
	SSO	–	–	–	–	–	✓

Source: Compiled by casewriters from “Upgrading Your Gamma Subscription,” help.gamma.app, <https://help.gamma.app/en/articles/8077107-upgrading-your-gamma-subscription>, accessed October 24, 2025, and; “What are Gamma Tokens and How Do They Work?” gamma.app, https://help.gamma.app/en/articles/11047156-what-are-gamma-tokens-and-how-do-they-work#h_27b96806f7, accessed September 14, 2025.

Exhibit 6 Gamma Key Metrics



Source: Company documents.

Exhibit 7a Anthropic Pricing, September 2025

Model	Base Input Tokens	5m Cache Writes	1h Cache Writes	Cache Hits & Refreshes	Output Tokens
Claude Opus 4.1	\$15 / MTok	\$18.75 / MTok	\$30 / MTok	\$1.50 / MTok	\$75 / MTok
Claude Opus 4	\$15 / MTok	\$18.75 / MTok	\$30 / MTok	\$1.50 / MTok	\$75 / MTok
Claude Sonnet 4	\$3 / MTok	\$3.75 / MTok	\$6 / MTok	\$0.30 / MTok	\$15 / MTok
Claude Sonnet 3.7	\$3 / MTok	\$3.75 / MTok	\$6 / MTok	\$0.30 / MTok	\$15 / MTok
Claude Sonnet 3.5 (deprecated)	\$3 / MTok	\$3.75 / MTok	\$6 / MTok	\$0.30 / MTok	\$15 / MTok
Claude Haiku 3.5	\$0.80 / MTok	\$1 / MTok	\$1.6 / MTok	\$0.08 / MTok	\$4 / MTok
Claude Opus 3 (deprecated)	\$15 / MTok	\$18.75 / MTok	\$30 / MTok	\$1.50 / MTok	\$75 / MTok
Claude Haiku 3	\$0.25 / MTok	\$0.30 / MTok	\$0.50 / MTok	\$0.03 / MTok	\$1.25 / MTok

Source: Claude Docs, Claude Developer Platform, "Models and Pricing," https://docs.claude.com/en/docs/about-claude/pricing?3574be7b_page=14&99edf393_page=2&a006523a_page=2&bdd1d8a9_page=3&ea471aba_page=2, accessed September 22, 2025.

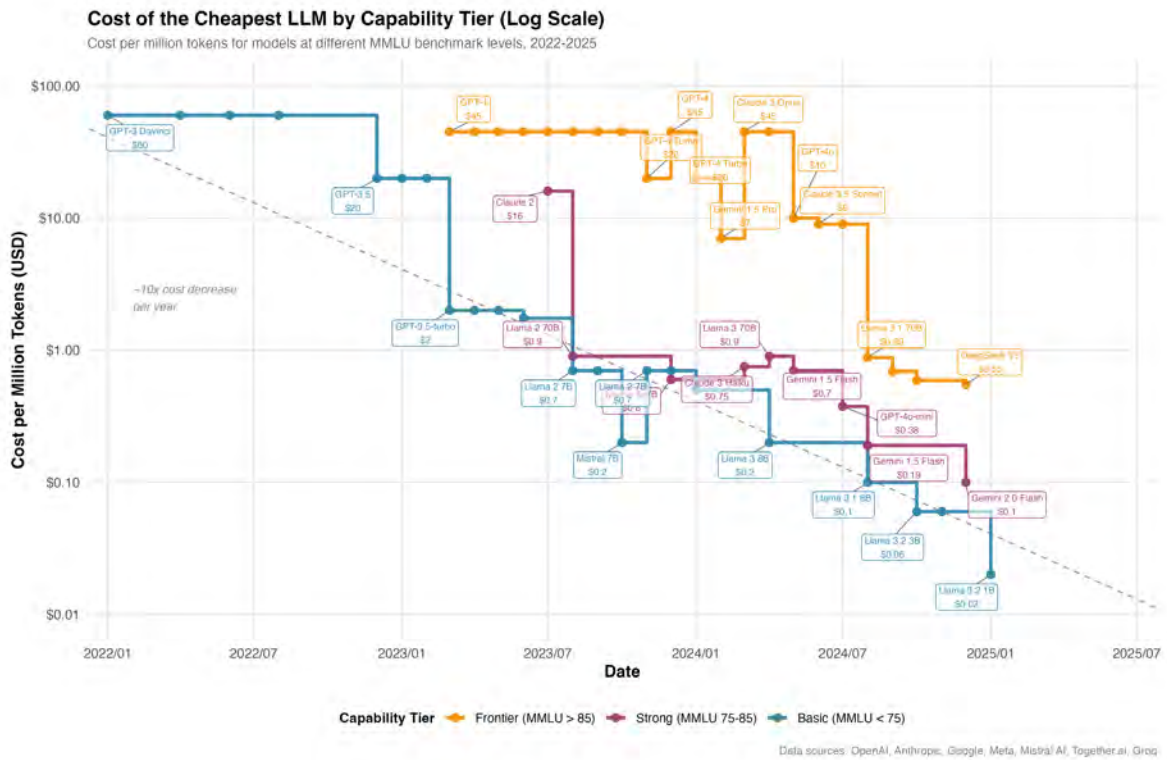
Note: Tokens were the smallest individual units of a language model, and could correspond to words, subwords, characters, or even bytes (in the case of Unicode). For Claude, a token approximately represented 3.5 English characters, though the exact number could vary depending on the language used.

Exhibit 7b OpenAI Pricing, September 2025

GPT-4.1	GPT-4.1 mini	GPT-4.1 nano	o4-mini
Fine-tuning price	Fine-tuning price	Fine-tuning price	Reinforcement fine-tuning price
Input: \$3.00 / 1M tokens	Input: \$0.80 / 1M tokens	Input: \$0.20 / 1M tokens	Input: \$4.00 / 1M tokens
Cached input: \$0.75 / 1M tokens	Cached input: \$0.20 / 1M tokens	Cached input: \$0.05 / 1M tokens	Cached input: \$1.00 / 1M tokens
Output: \$12.00 / 1M tokens	Output: \$3.20 / 1M tokens	Output: \$0.80 / 1M tokens	Output: \$16.00 / 1M tokens
Training: \$25.00 / 1M tokens	Training: \$5.00 / 1M tokens	Training: \$1.50 / 1M tokens	Training: \$100.00 / training hour

Source: OpenAI, "API Pricing," <https://openai.com/api/pricing/>, accessed September 14, 2025.







Exhibit 8 Inference Costs Over Time



Source: Recreated using data from Guido Appenzeller, "Welcome to LLMflation – LLM Inference Cost is Going Down Fast," November 12, 2024, a16z, <https://a16z.com/llmflation-llm-inference-cost/>, accessed September 14, 2025, OpenAI, "API Pricing," <https://openai.com/api/pricing/>, accessed September 14, 2025, "Serverless Inference," <https://www.together.ai/pricing>, accessed December 19, 2025. Code to generate the figure produced by Claude Code, December 2025. All data sources used by the Claude Code are cited and are reviewed by the authors.

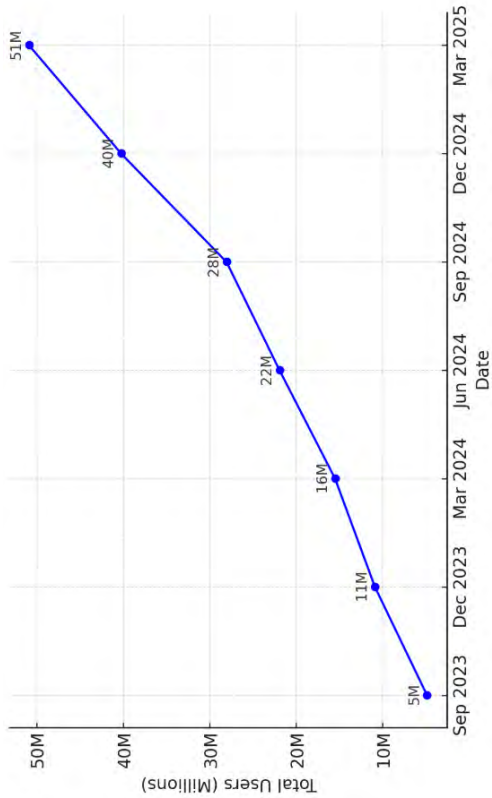
Note: MMLU stands for Massive Multitask Language Understanding, a benchmark used to test an AI model's ability to answer multiple-choice questions across 57 academic and professional subjects, simulating real-world knowledge and reasoning.

Exhibit 9 Lean Native AI Companies Leaderboard

	Company	Description	Location	Annual Revenue	# of Employees	Revenue/Emp. ?	Profitable ?	Total Funding	Valuation	Valuation/EmpLOYEE
1	 Telegram	Messaging	Dubai	\$1,000,000,000	30	\$33,333,333	Yes	\$3,200,000,000	\$30,000,000,000	\$1,000,000,000
2	 Midjourney	Image Generation	San Francisco	\$500,000,000	40	\$12,500,000	Yes	\$0	\$10,000,000,000	\$250,000,000
3	 SurgeAI	AI Training Data	San Francisco	\$1,000,000,000	110	\$9,090,909	Yes	\$0		
4	 AnySphere (Cursor)	AI Code Editor	San Francisco	\$100,000,000	20	\$5,000,000		\$173,400,000	\$2,600,000,000	\$130,000,000
5	 Base44	AI Code Gen	Tel Aviv	\$3,500,000	1	\$3,500,000	Yes	\$0	\$80,000,000	\$80,000,000
6	 CalAI	AI Calorie Tracker	New York	\$12,000,000	4	\$3,000,000	Yes	\$0		
7	 Mercor	Talent Marketplace + Data	San Francisco	\$75,000,000	30	\$2,500,000		\$135,600,000	\$2,000,000,000	\$66,666,667
8	 Chai Research	Social AI Platform	Palo Alto	\$30,000,000	12	\$2,500,000	Yes	\$10,000,000	\$450,000,000	\$37,500,000
9	 Fal.ai	Generative media platform	San Francisco	\$95,000,000	40	\$2,375,000	Yes	\$197,000,000		
10	 Eleven Labs	AI Voice, Text-to-Speech	New York	\$100,000,000	50	\$2,000,000		\$208,000,000	\$3,300,000,000	\$66,000,000
11	 Stackblitz (Bolt.new)	AI Code Editor	San Francisco	\$40,000,000	20	\$2,000,000		\$87,900,000	\$700,000,000	\$35,000,000
12	 OpenAI	AI for Image and Video	San Francisco	\$20,000,000	10	\$2,000,000	Yes	\$5,000,000		
13	 Lovable	AI Code Gen	Stockholm	\$75,000,000	40	\$1,875,000		\$222,500,000	\$1,800,000,000	\$45,000,000
14	 Gamma	AI for Presentations	San Francisco	\$50,000,000	28	\$1,785,714	Yes	\$19,000,000		
15	 Markable AI	AI Creator Marketing	Seattle	\$40,000,000	25	\$1,600,000	Yes	\$17,000,000		
16	 Genspark	AI Super Agent	Palo Alto	\$36,000,000	24	\$1,500,000		\$160,000,000	\$530,000,000	\$22,083,333
17	 Aragon AI	AI Avatar Generation	San Francisco	\$12,000,000	8	\$1,500,000	Yes	\$900,000	\$10,000,000	\$1,250,000
18	 Solveai	AI for Homework	San Francisco	\$6,000,000	4	\$1,500,000				
19	 Olive	AI Consumer Portfolio	New York	\$6,000,000	4	\$1,500,000	Yes	\$600,000	\$10,000,000	\$2,500,000
20	 Arcads	AI Video Ads Generator	Paris	\$7,000,000	5	\$1,400,000	Yes	\$0		
21	 Retail AI	AI Voice	San Francisco	\$15,000,000	11	\$1,363,636	Yes	\$4,725,000		
22	 Develop Health	AI for Prescribers	San Francisco	\$6,000,000	5	\$1,200,000	Yes	\$3,200,000		

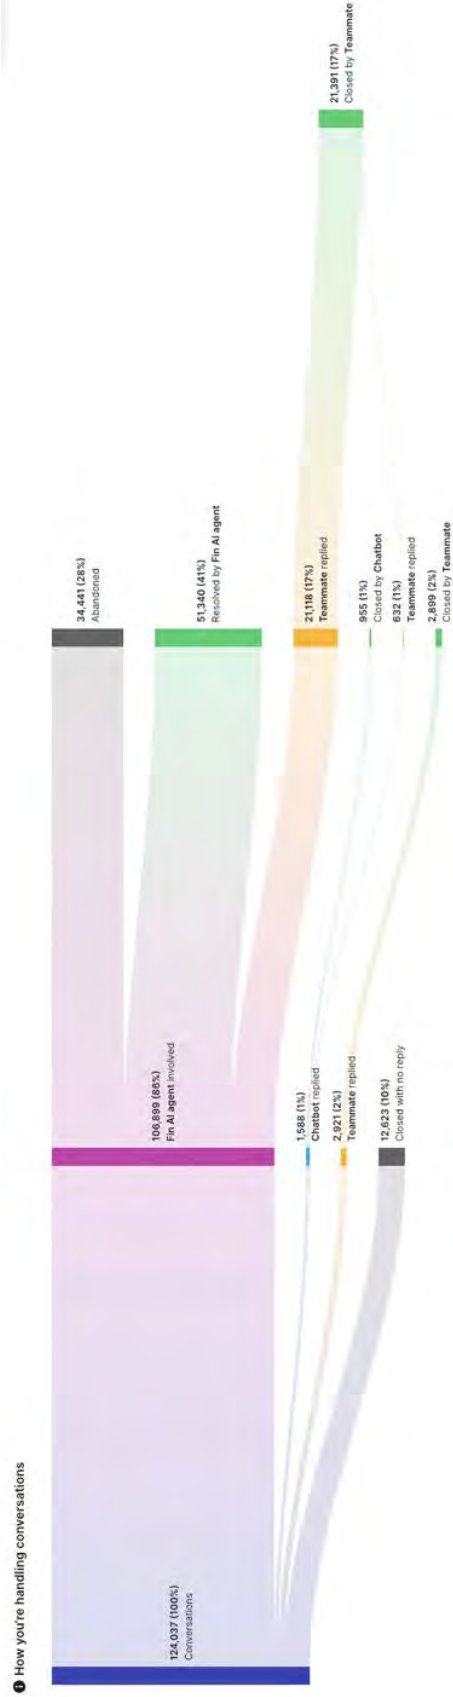
Source: Excerpted from Henry Shi, "Top Lean AI Native Companies Leaderboard," <https://leanaileaderboard.com/>, accessed August 1, 2025.

Exhibit 10 Gamma’s User Growth (Millions)



Source: Company documents.

Exhibit 11 Flow of Customer Service Inquiries—Chatbot vs Human



Source: Company documents.

Exhibit 12a Manus Announcement of Presentation Maker



Manus AI
54,154 followers
1mo · 🌐

+ Follow ...

🚀 From hours to minutes: Manus built a complete presentation deck automatically.

🌟 What Manus delivers:

- Stunning visuals
- Professional layouts & structure
- Clear, compelling content
- Ready-to-use document exports

No more hunting for images or wrestling with slide designs. Just pure, polished presentations that look like they took days to create.

WHITewater RAPIDS EXPEDITION

ZAMBEZI RIVER CHALLENGE

Navigate the world's widest whitewater beneath Victoria Falls, where Class V rapids with names like 'Devil's Toilet Bowl' and 'Commercial Suicide' churn between 350-foot canyon walls. Feel your raft launch skyward as you battle Africa's greatest river.

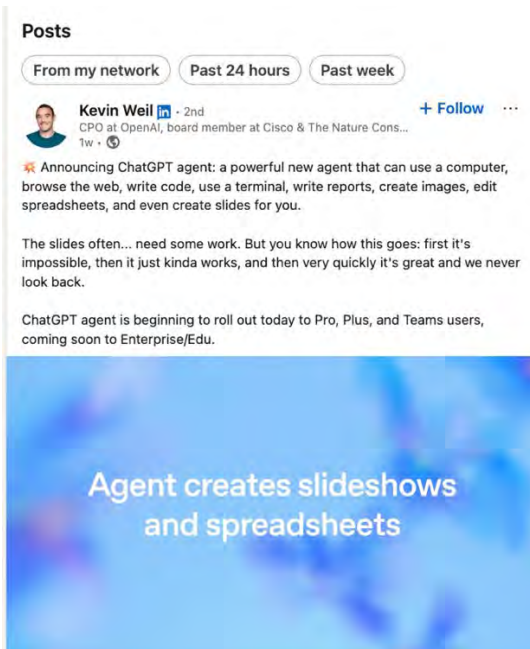
LOCATION: Zimbabwe/Zambia border

BEST SEASON: July to January

0:41 1x CC 🔊


Source: Adapted from Manus AI, LinkedIn profile, https://www.linkedin.com/posts/manus-im_from-hours-to-minutes-manus-built-a-complete-activity-7338228444389900289-ifph/, accessed August 1, 2025.

Exhibit 12b OpenAI Agent Announcement of Presentation Maker



Posts

From my network Past 24 hours Past week

Kevin Weil  · 2nd
CPO at OpenAI, board member at Cisco & The Nature Cons...
1w · 🌐

+ Follow ...

🚀 Announcing ChatGPT agent: a powerful new agent that can use a computer, browse the web, write code, use a terminal, write reports, create images, edit spreadsheets, and even create slides for you.

The slides often... need some work. But you know how this goes: first it's impossible, then it just kinda works, and then very quickly it's great and we never look back.

ChatGPT agent is beginning to roll out today to Pro, Plus, and Teams users, coming soon to Enterprise/Edu.

Agent creates slideshows and spreadsheets

Source: Adapted from Kevin Weil, LinkedIn profile, <https://www.linkedin.com/feed/update/urn:li:activity:7351662630152912898/>, accessed August 1, 2025.

Exhibit 13 Competitor Map

Company Name	Year of Founding	Total Raised	Most Recent Valuation	Product Summary	Pricing	Average Monthly Organic Search Traffic	ARR Estimate
Gamma	2020	\$23M	\$112M	An AI-powered design platform that instantly transforms text prompts into polished presentations, documents, and web pages.	Free and Premium Options: Plus - \$10/month (monthly), \$8/month (yearly) Pro - \$20/month (monthly), \$15/month (yearly)	3.8M	\$50M
Manus	2022	\$85M	\$500M	An end-to-end "AI slide generator" that builds complete PowerPoint, Google Slides, or PDF decks from a single text brief in seconds.	Free and Premium Options: Basic - \$19/month (monthly), \$16/month (yearly) Plus - \$39/month (monthly), \$33/month (yearly) Pro - \$199/month (monthly), \$166/month (yearly)	404,000	\$69M
OpenAI	2015	\$63.9B	\$300B	Turns a single prompt into a fully researched, formatted PowerPoint deck ready for download directly from ChatGPT.	Free and Premium Options: Pro - \$20/month (monthly) Team (per user) - \$30/month (monthly), \$30/month (annually) Enterprise option also available	80.4M	\$12.7B
Canva (Magic Design for Presentations)	2012	\$581.5M	\$26B	Produces a first-draft deck—complete with copy, imagery, and your brand colors—then lets you refine it in Canva's familiar interface.	Free and Premium Options: Team (per user) - \$10/month (monthly), \$100/year (yearly) Enterprise option also available	273M	\$3.16B
Beautiful.AI	2015	\$16.3M	\$36M	A cloud-based AI presentation maker that converts a brief prompt into a fully designed, brand-consistent slide deck in seconds.	Free and Premium Options: Team (per user) - \$50/month (monthly), \$40/month (yearly) Enterprise option also available	234,000	NA
Microsoft (Copilot with Powerpoint)	1975	NA	\$3.8T	A built-in AI assistant that turns a simple prompt into a polished slide deck—or rewrites and summarizes existing slides—directly within PowerPoint.	\$30/month (yearly) in addition to Microsoft 365 subscription Enterprise rate is the same. Copilot studio plans are available for \$200/year	11M	\$245B
Alphabet (Gemini with Google Slides)	1988	NA	\$2.28T	Turns a prompt into a full deck, auto-generating slides, text, and images.	Free and Premium Options: Standard - \$16.80/month (monthly), \$14/month (yearly) Plus - \$26.40/month (monthly), \$22/month (yearly) Enterprise Plus - \$42/month (monthly), \$35/month (yearly)	49.2M	\$350B
Pitch	2018	\$137.7M	\$527.7M	A web-based AI presentation maker that converts a simple prompt into a fully editable, design-ready slide deck in seconds.	Free and Premium Options: Pro - \$25/month (monthly), \$120/month (yearly) Business - \$22/month (monthly), \$85/month (yearly) Enterprise option also available	175,000	\$5.5M

Source: Compiled by casewriters from: Gamma: Pitchbook - <https://my.pitchbook.com/profile/483289-84/company/profile#ai-summary>, Pricing - <https://gamma.app/pricing>; Manus: Pitchbook - <https://my.pitchbook.com/profile/533115-91/company/profile#ai-summary>, Pricing - <https://manus.im/pricing>; OpenAI: Pitchbook - <https://my.pitchbook.com/profile/149504-14/company/profile>, Pricing - <https://chatgpt.com/pricing>; Canva: Pitchbook - <https://my.pitchbook.com/profile/56364-67/company/profile>, Pricing - <https://www.canva.com/pricing/?msockid=1c17d730effd6ca12a50c179ee406d6f>; Beautiful.AI: Pitchbook - <https://my.pitchbook.com/profile/172306-18/company/profile#ai-summary>, Pricing - <https://www.beautiful.ai/pricing>; Microsoft: Pitchbook - <https://my.pitchbook.com/profile/11026-45/company/profile>, Pricing - <https://www.microsoft.com/en-us/microsoft-365-copilot/pricing/individuals>; Alphabet (Gemini with Google Slides): Pitchbook - <https://my.pitchbook.com/profile/10453-33/company/profile>, Pricing - <https://pitch.com/pricing/us>, and; Semrush (Monthly Organic, July 2025), https://www.semrush.com/?trclid=3FnrXXVYjxycWVU1lw2pzP6UkpQKfYjFjUiv1s0&utm_source=affiliate&utm_Medium=impact&utm_campaign=3426829&utm_terms=&utm_content=&irgwc=1&afsrc=1&utm_medium=impact&utm_term=&ir_par_tnerid=3426829&ir_adid=995972&ir_campaignid=13053&gad_source=1, accessed August 4, 2025.

Endnotes

- ¹ “COMPANY NEWS; Microsoft Buys Software Unit,” The New York Times, July 31, 1987, <https://www.nytimes.com/1987/07/31/business/company-news-microsoft-buys-software-unit.html>, accessed October 4, 2025.
- ² “Presentation Software Market Size and Share Analysis – Growth Trends and Forecasts (2025-2032).” Coherent Market Insights. <https://www.coherentmarketinsights.com/industry-reports/presentation-software-market>, accessed September 22, 2025.
- ³ “Scaling the Cursor of Slides to \$50M ARR: Gamma founder John Noronha,” Training Data podcast, Episode 59, Sequoia Capital, <https://www.sequoiacap.com/podcast/training-data-jon-noronha/>, accessed October 1, 2025.
- ⁴ “The Ultimate Global PowerPoint Study,” Nielsen Company and empower, 2020, <https://www.empowersuite.com/hubfs/Marketing/Downloads/PowerPoint%20Study%202020/Englisch/The%20Ultimate%20Global%20PowerPoint%20Study%20-%20empower.pdf>, accessed September 13, 2025.
- ⁵ Glassdoor. “McKinsey & Company.” McKinsey & Company Pay FAQ, https://www.glassdoor.com/Salary/McKinsey-and-Company-Visual-Graphics-Specialist-Salaries-E2893_D_KO21,47.htm, accessed September 22, 2025.
- ⁶ Flocksy. The Resource Center. “What is a Graphic Designer’s Salary? A comprehensive Look at Earnings, Influencing Factors, and Future Prospects.” https://flocksy.com/resources/what-is-a-graphic-designers-salary-a-comprehensive-look-at-earnings-influencing-factors-and-future-prospects/?utm_source=chatgpt.com, accessed September 22, 2025.
- ⁷ “How Much Should You Pay for a PowerPoint Presentation? A Guide to Professional Presentation Design Services,” SlideModel.com, February 3, 2025, <https://slidemodel.com/presentation-design-pricing/>, accessed September 22, 2025.
- ⁸ Ibid.
- ⁹ Frederic Lardinois, “Optimizely raises \$50M Series D round for its experimentation platform,” TechCrunch, June 18, 2019, <https://techcrunch.com/2019/06/18/optimizely-raises-50m-series-d-round-for-its-experimentation-platform/>, accessed August 1, 2025.
- ¹⁰ Optimizely (acquired), Deals, PitchBook Data, Inc, <https://my.pitchbook.com/profile/51283-54/company/deals>, accessed June 2025.
- ¹¹ “How AI is Transforming Presentations: Gamma’s Story,” JPMorgan.com, Jan 16, 2025, <https://www.jpmorgan.com/insights/technology/artificial-intelligence/gammas-startup-journey-the-future-of-presentations-with-ai>, accessed August 1, 2025.
- ¹² Grant Lee, LinkedIn profile: https://www.linkedin.com/posts/grantslee_arent-you-a-bit-old-to-be-starting-a-company-activity-7316475455149117441-NThm/?utm_source=share&utm_medium=member_desktop&rcm=ACoAABAAt8SoBD6V484BC0AztN_dVgz3qV-ghdm4, accessed August 1, 2025.
- ¹³ Ibid.
- ¹⁴ Alex Konrad, “Meet Gamma, A Low-Profile AI Startup That’s Actually Profitable,” Upstarts, April 29, 2025, <https://www.upstartsmia.com/p/gamma-ai-startup-profits>, accessed October 6, 2025.
- ¹⁵ John Tian, “How Gamma Grow from 0-\$50MARR, 50M Users with 30 People,” The Signal, May 5, 2025, <https://signalhub.substack.com/p/how-gamma-grow-from-0-50m-arr-50m>, accessed October 1, 2025.
- ¹⁶ Diana Berlin, “Toward Tools for Imagination,” September 6, 2022, <https://dianaberlin.com/posts/2022/9/6/toward-tools-for-imagination>, accessed September 13, 2025.
- ¹⁷ John Tian, “How Gamma Grow from 0-\$50MARR, 50M Users with 30 People,” The Signal, May 5, 2025, <https://signalhub.substack.com/p/how-gamma-grow-from-0-50m-arr-50m>, accessed October 1, 2025.
- ¹⁸ “How AI is Transforming Presentations: Gamma’s Story,” JPMorgan.com, Jan 16, 2025, <https://www.jpmorgan.com/insights/technology/artificial-intelligence/gammas-startup-journey-the-future-of-presentations-with-ai>, accessed August 1, 2025.
- ¹⁹ Gamma.app. “About.” <https://gamma.app/about>, accessed August 1, 2025.

- ²⁰ Jake Schonberger and Lynda Hwang, "How Gamma Scaled their Marketing with beehiiv's Ad Network," beehiiv.com, December 24, 2024, <https://blog.beehiiv.com/p/how-gamma-scaled-their-marketing-with-beehiiv-s-ad-network>, accessed September 13, 2025.
- ²¹ "What are Tokens and How to Count Them," OpenAI, <https://help.openai.com/en/articles/4936856-what-are-tokens-and-how-to-count-them>, accessed October 1, 2025.
- ²² Casewriter analysis based on "Pricing," OpenAI, <https://platform.openai.com/docs/pricing#image-tokens>, accessed October 1, 2025.
- ²³ "Gamma helps team create polished presentations with Claude," anthropic.com, https://www.anthropic.com/customers/gamma?utm_source=chatgpt.com, accessed September 14, 2025.
- ²⁴ "Scaling the Cursor of Slides to \$50M ARR: Gamma founder John Noronha," Training Data podcast, Episode 59, Sequoia Capital, <https://www.sequoiacap.com/podcast/training-data-jon-noronha/>, accessed October 1, 2025.
- ²⁵ Guido Appenzeller, "Welcome to LLMflation – LLM Inference Cost is Going Down Fast," November 12, 2024, a16z, <https://a16z.com/llmflation-llm-inference-cost/>, accessed September 14, 2025.
- ²⁶ Jeff Vogel, Michael Stricklen, and Ajay Khandelwal, "Cutting Costs in the Cloud: Six Strategies for SaaS Companies," EY Parthenon, Aug 31, 2023, https://www.ey.com/en_us/insights/tmt/cutting-costs-in-the-cloud-six-strategies-for-saas-companies, accessed September 23, 2025.
- ²⁷ John Tian, "How Gamma Grow from \$0-\$50M ARR, 50M Users with 30 people," The Signal, May 5, 2025, <https://signalhub.substack.com/p/how-gamma-grow-from-0-50m-arr-50m>, accessed September 14, 2025.
- ²⁸ The Peel with Turner Novak, "Gamma's 50 Million User Playbook, Grant Lee Co-founder and CEO," YouTube, published July 3, 2025, <https://www.youtube.com/watch?v=zorJtj1Bgm0>, accessed August 1, 2025.
- ²⁹ "State of Startup Compensation, H1 2024," Carta, July 25, 2024, <https://carta.com/data/startup-compensation-h1-2024/#salary-trends>, accessed October 1, 2025.
- ³⁰ Griffith, Erin, "A.I is Changing How Silicon Valley Builds Start-ups," *The New York Times*, February 20, 2025, <https://www.nytimes.com/2025/02/20/technology/ai-silicon-valley-start-ups.html>, accessed August 1, 2025.
- ³¹ Ibid.
- ³² Grant Lee, LinkedIn profile, https://www.linkedin.com/posts/grantslee_finding-product-market-fit-doesnt-mean-its-activity-7321548989810581504-h8Ec/, accessed August 1, 2025.
- ³³ Mehta, Ivan, "A Quarter of Startups in YC's Current Cohort Have Codebases that are Almost Entirely AI Generated," TechCrunch, March 6, 2025, <https://techcrunch.com/2025/03/06/a-quarter-of-startups-in-ycs-current-cohort-have-codebases-that-are-almost-entirely-ai-generated/>, accessed August 3, 2025.
- ³⁴ Conor Grennan, LinkedIn profile, <https://www.linkedin.com/feed/update/urn:li:activity:7323684400506810368/>, accessed August 1, 2025.
- ³⁵ Anthropic. Customers. Case Studies. "Gamma Helps Teams Create Polished Presentations with Claude." <https://www.anthropic.com/customers/gamma>, accessed August 3, 2025.
- ³⁶ "The Ultimate Global PowerPoint Study," Nielsen Company and empower, 2020, <https://www.empowersuite.com/hubfs/Marketing/Downloads/PowerPoint%20Studie%202020/Englisch/The%20Ultimate%20Global%20PowerPoint%20Study%20-%20empower.pdf>, accessed September 13, 2025.
- ³⁷ Ibid.

Notes

A series of horizontal dotted lines for taking notes.



2026

Connective
Leadership Summit
19 - 22 May 2026 | Melbourne VIC

Superhive: Permitting AI





MITCHELL WEISS
FARES KHRAIS

SuperHive: Permitting AI

In March 2025, ai71, an Abu Dhabi-based company working to speed up building permits in the UAE and beyond faced a permissions question of its own: Should it give itself the greenlight on SuperHive, its AI construction intelligence product, before the technology was 100 percent accurate?

SuperHive, known in its early days as Cons71, was conceived as an “AI-powered Real Estate Assistant” with uses for firms and for governments. The feature that was at least 90 percent ready by Spring 2025 was its automated compliance checking. The technology could automatically scan site and building designs to compare them against local codes and regulations. All over the world, building codes and the processes for navigating them could often be labyrinth. Governments, the thinking went, could use these tools to help in permit review. Firms could use them in advance to aid in shaping permit applications. The potential of SuperHive was to reduce permitting times from months to minutes; and beyond that to transform everything from design, to modeling, to space optimization, to engineering and mechanical integration.

The compliance tool was just one feature set of SuperHive, which was just one early product of ai71. The company was also underway with other tools that had government and private use cases, in healthcare, for example, and in law. And products were only one aspect of ai71’s developing business model. The company provided access to large language models (LLMs) via their platform. And they were building out advisory services to try to help government and firms make productive use of AI.

More was perhaps resting on the launch of the construction product than just its success or failure. ai71 was stood up by VentureOne, an arm of the UAE’s Advanced Technology Research Council, in large measure to commercialize the United Arab Emirates’s (UAE) family of LLMs, Falcon. These open-source models, developed by the UAE’s Technology Innovation Institute (TII), had landed with fanfare in 2023, when they topped some global rankings. Now, the company – and indeed the models – were part of a larger national AI strategy to diversify the UAE economy and strengthen its AI sovereignty.

H.E. Faisal Al Bannai, ai71’s chairman, had launched the company in 2023, “to bring the power of open source LLM models to the world through transformative AI-powered enterprise solutions.”¹ Given the backlog of building permits around the world, not to mention housing shortages, ai71’s technology seemed potentially transformational. But like many AI tools at the time, it wasn’t 100 percent reliable. Amidst a race for AI sovereignty and profitability, for talent and for market share, the company would have to decide when it came to their AI, how good was good enough?

Professor Mitchell Weiss and Assistant Director Fares Khrais (Middle East and North Africa Research Center) prepared this case. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. The citation review for this case has not yet been completed. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2026 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

This document is authorized for use only by Mitchell Weiss in 2026.

This article is made available to you with compliments of Connective Leadership Summit, Innovate to Lead for your personal use. Further posting, copying or distribution is not permitted.

UAE + AI

ai71 and SuperHive were new elements of an expansive vision for AI in the UAE. In 2017, the country launched its National Artificial Intelligence Strategy 2031 as part of its ambition to transition into a knowledge-based economy.^{2,3} The UAE, home to roughly 10 million people, was founded by the late sheikh Zayed bin Sultan Al Nahyan in 1971 as a federation of seven Emirates with Abu Dhabi as the capital.⁴ A member of the organization of petroleum exporting countries (OPEC), oil had played a crucial role in the country's economy since its inception.^{5,6} Particularly since the early 2000s, the country had made efforts to diversify its economy away from oil.⁷ To that end, the UAE had launched several national strategies designed to grow emerging sectors.⁸ AI was one of them. By 2024, non-oil activities contributed to 74% of GDP (compared to 69% in 2015).^{9,10} IBM forecasted AI to contribute \$96 billion to UAE's GDP (14% of total) by 2030.¹¹ (See **Exhibit 1** for an overview of the AI strategy.)

The UAE was perhaps well positioned to succeed in AI.¹² It had wealth to purchase the hardware necessary.¹³ It had electricity to power it all, given its oil and renewable energy (particularly year-round solar). Dr. Hakim Hacid, chief researcher at TII, said, "We don't have a shortage of resources, and the AI sector is getting everything it needs to progress and achieve something great." The country was also working to attract businesses and talent alike, promising high standards of living, no income tax, and more.¹⁴

Additionally, the UAE government had been characterized as tech-savvy and had positioned itself as such.¹⁵ It boasted a highly digitized government with plans to be the world's first fully AI-native government by 2027.¹⁶ It offered AI businesses access to large sets of anonymized data on its residents, including those from hospitals and state-backed entities.¹⁷ At a virtual event in 2024, Sam Altman, co-founder and CEO of OpenAI, highlighted the UAE's potential as a "regulatory sandbox" for AI experimentation.¹⁸ Mehdi Ghissassi, chief product and technology officer at ai71, commented, "In general, governments are resistant to using AI. But that's not true for the UAE, they're actually more keen to try it out than the private sector."

Since the announcement of the national AI strategy, the UAE had made headlines with several multi-million and billion-dollar investments in infrastructure, developing technologies, and even in becoming a manufacturer of semiconductors used in chips critical for AI.^{19,20} Dr. Najwa Aaraj, CEO at TII, observed, "There is a very big commitment of the country overall to AI."

The investments were not restricted to the UAE. In 2022, G42, the Abu-Dhabi based technology holding company, launched a \$10 billion fund to "accelerate the pace of global innovation" by functioning as a private equity investor focused on late-stage investments in high-growth regions.²¹ It had the world's largest AI-dedicated supercomputer (located in the U.S.), built by American firm Cerebras.²²

Many of the investments and partnerships were with American entities and formed the basis of growing AI-focused ties between the UAE and U.S.²³ G42 had partnerships with several American firms including OpenAI, IBM, NVIDIA, and others.²⁴ In April 2024, Microsoft announced plans to invest \$1.5 billion in Abu Dhabi to help create a skilled AI workforce in the UAE and pave the way for collaborations on AI infrastructure in the Middle East, Central Asia and Africa.²⁵ The U.S. rolled back its restrictions on the amount of NVIDIA-class chips the UAE could import in May of 2025.²⁶

A Growing AI Ecosystem

The national AI strategy helped create an AI ecosystem in the UAE and Abu Dhabi in particular. The multiplying entities including startups, educational and research institutions, investment funds,

and more. Key players included new AI-focused regulatory authorities and other entities dedicated towards developing and commercializing AI technologies like G42 and the Advanced Technology Research Council (ATRC).

The ATRC, founded in 2020 was responsible for defining Abu Dhabi's technology research strategy; consolidating and facilitating funding; and policy making.²⁷ Its executive board was made of seven of the country's most influential figures and was chaired by Abu Dhabi's crown prince sheikh Khaled bin Mohamed bin Zayed Al Nahyan.²⁸ The ATRC carried out its efforts through three entities; Aspire, TII, and VentureOne. Aspire focused on developing technology talent. TII, meanwhile, was responsible for researching and developing the technologies. (See **Exhibit 2** for TII research areas.) Its most well-known was Falcon, the LLM, named after the country's national animal. VentureOne was set up to commercialize those technologies by refining them, converting them into products, and then selling them both locally and internationally.

The Abu Dhabi Digital Government Strategy

AI efforts in Abu Dhabi extended into the administration of government. In January 2025, the Abu Dhabi government announced the Abu Dhabi Government Digital Strategy 2025-2027, with a goal to advance "the emirate's transformation towards an AI-powered government."²⁹ That meant shifting all of the government's operations onto the UAE's sovereign cloud, and digitizing and automating all government processes.³⁰

The plan earmarked roughly \$3.5 billion to be deployed by the Department of Government Enablement (DGE).³¹ Reda Nidhakou, CEO of VentureOne, commented, "There's a real drive to adopt AI in the UAE. DGE, for example, has put a team of AI experts together to enable all the entities within the government to use AI. We have been overwhelmed with end user demand in the early days of ai71." The DGE worked closely with many of Abu Dhabi's entities; it had a preexisting responsibility to create government-wide policies on three key domains: human resources, digital and cybersecurity, and procurement.^{32,33} It had already spent over a decade digitizing the government and its services to residents and businesses.³⁴

Much of that work had come to fruition in the TAMM platform, where AI would also now be brought to bear.³⁵ TAMM (an Arabic term which meant "consider it done") was an application that could be downloaded onto a user's smartphone.³⁶ It catered to Abu Dhabi's 2.5 million residents and offered them the ability to carry out 1,000 government services (e.g. paying speeding tickets and utility bills, registering a marriage, or even applying for a permit for large gatherings like a private party).^{37,38} In 2025, users could, for example, ask the TAMM-AI assistant how to complete certain requests like renewing their vehicle registration.³⁹ They could ask the chatbot to do the whole process for them. It could even obtain car insurance quotes and select one on the users' behalf in a matter of seconds.⁴⁰

Other users and uses emerged as well. The Abu Dhabi Department of Municipalities and Transport, for example, announced in 2024 several partnerships to obtain AI-powered customized solutions in the mobility sector to better manage traffic, deploy charging infrastructure for electronic vehicles, enhance transport safety and more.⁴¹ The Abu Dhabi National Oil Company announced plans to use agentic AI to slash the time taken for seismic surveys from months to days and boost the accuracy of oil production forecasts by 90%.⁴² The Abu Dhabi Judicial Department launched an AI-powered system available 24/7 to automatically reduce traffic fines without a court's involvement.⁴³

AI Sovereignty

The investments in AI were also meant to help the UAE go beyond government efficiency; they were part of an effort at “AI sovereignty”.

The notion had many definitions. NVIDIA characterized sovereign AI as “a nation’s capabilities to produce artificial intelligence using its own infrastructure, data, workforce and business networks.”⁴⁴ That conception could be expanded to include other aspects such as regulations and legislation around AI, strategic investments in research, and building an AI ecosystem.⁴⁵ AI sovereignty was said to be potentially and particularly valuable considering AI’s growing role in military, society, and politics.⁴⁶ Jensen Huang, NVIDIA’s founder and CEO had gone so far as to assert that, “Every country needs to own the production of their intelligence.”⁴⁷

AI sovereignty was an evolution of an older concept: data sovereignty.⁴⁸ Long before LLMs, governments had been making choices about whether to store data on their servers that they owned and were located on their own soil or in their own sovereign cloud. AI sovereignty went beyond data sovereignty and included the AI, including LLMs, which parsed or were built from that data.⁴⁹ Owning the AI models was said to help governments avoid the risk of relying on foreign entities for access.⁵⁰

NVIDIA’s leadership and UAE’s seemed to share a perspective on this. In 2024, at a fireside chat at the World Government Summit held in Dubai, Huang commented on sovereign AI, “It codifies your culture, your society’s intelligence, your common sense, your history.”⁵¹ Omar Al Olama, the UAE’s minister of AI, responded, “We completely subscribe to that vision...That’s why the UAE is moving aggressively on creating large language models and mobilizing compute.”⁵²

Models built by local teams might also play a role in protecting and promoting national language, culture and values.⁵³ Dr. Aaraj observed, “Having your own models helps get rid of any biases towards the region that may exist in other models.” A Deloitte report on sovereign LLMs elaborated, “the reality is that many platforms are often biased toward their country of origin across several dimensions, including legal frameworks, ethical and equity frameworks, language, and culture.”⁵⁴ Several LLMs had been developed in the UAE, including Jais, Nanda, and Sherkala by G42, and Falcon by the ATRC’s TII.

Even if desired, achieving AI sovereignty had its challenges. It was costly in terms of energy and cutting-edge hardware.⁵⁵ It required tech talent and developed regulatory and governance frameworks.⁵⁶ Much of the leading AI technology was concentrated in a small number of private tech firms.⁵⁷ Moreover, as governments invested in developing their AI, they needed to navigate decisions on whether to compete or collaborate (or both) with the private sector.⁵⁸ Some governments settled for having the ability to control how models were used as opposed to fully controlling or providing the technology themselves.⁵⁹

In 2025, over 60 countries had national AI strategies.⁶⁰ Saudi Arabia, which neighbored the UAE, had an AI strategy focused on developing talent and regulation, research and collaboration, developing an AI ecosystem, and providing and attracting funding towards those efforts.⁶¹ India espoused AI as a tool for inclusive growth and sought to position itself as a testing ground for global AI innovations catered to emerging markets.⁶² Stanford University’s Global AI Power Rankings took into consideration 42 indicators from research papers to patents to investment and more and in 2024 ranked the top five AI-nations as the U.S., China, the U.K., India, and the UAE.⁶³

Falcon

Falcon had helped put the UAE on the AI map. When Falcon 40b launched in March 2023, it topped leaderboards on Hugging Face and was reported to outperform several comparable models, including ChatGPT-3.⁶⁴ Dr. Ebtesam Almazroui, a director at Abu Dhabi's AI and Digital Security Research Center remarked in a statement that Falcon 40b "[matched] the performance of Chinchilla (from DeepMind by Google) and PaLM-62B (also from Google), at considerably lower training costs in comparison."⁶⁵ It also required less compute power than similar LLMs like Meta's Llama.⁶⁶ One AI researcher at Hugging Face said in a statement in *Time Magazine*, "The UAE was not well known, before, for training models... But then, by more or less the next day, we knew that they can train models, they open-source their models, they publish research around it, which benefits all."⁶⁷ Nidhakou expressed even the team's wonder at all this, "We knew it would do well, but we were honestly pleasantly surprised with both its performance and traction."

Developing Falcon had not been cheap. The 384 A100 NVIDIA chips used to train Falcon 40b had cost \$10,000 each; it was not unusual for LLMs to cost \$20 million or more to develop.^{68,69} On top of that was the cost of the compute to support its daily use. Despite the costs, TII would earn no revenue from Falcon; it was open-source and free to the world.^{70,71}

TII's leadership had felt open-sourcing Falcon was the right call. Nidhakou remarked, "We knew 40b was only the beginning of our journey. We were confident that the value it would bring in was worth making it open source. It at least meant Falcon 40b could be tested by countless people around the world." Dr. Hacid added, "It made us part of a very active and collaborative global community. Developers in the community let us know how the model can be improved and come up with use cases we didn't think of." Developers sometimes preferred open-source LLMs since they could be fine-tuned, and there were several available.⁷² Some open-source LLMs had restrictions regarding commercial use, but Falcon's leadership elected not to put any in place.

Being open source, the team believed, also aligned with UAE's vision to make an impact, which in turn would attract world-class talent. Omar Al Olama, UAE's first AI minister said, "If you bring a person who is the top scientist in this field, and you [give] two options: either work for-profit, closed-source, and just benefit yourself, or work on something that's open-source that will really change the lives of people, what would you decide?"⁷³

Falcon Models

Between 2023 and 2025, TII released several more Falcon models, big and small. (See **Exhibit 3** for model timeline and **Exhibit 4** for benchmarking). The largest was Falcon 180b. It was a "super-powerful" model with 180 billion parameters and, according to TII, once again beat competitors from Meta and Google and rivaled OpenAI's closed source ChatGPT-4.⁷⁴

Although TII had built leading models, it was a challenge getting people to use them. Dr. Hacid explained one of the reasons, "The AI we built performed really well in tests, but it wasn't necessarily good when it came to *actual* usage by developers."

So, newer models, like Falcon 3, were built with adoption in mind. Released in December 2024, Falcon 3 was a much smaller LLM than its predecessors at under 10 billion parameters. Dr. Aaraj said, "Bigger models are valuable and can be more capable than smaller models, but the adoption can be limited by the availability of compute power of users." Smaller models were easier to deploy on cheaper hardware; Falcon 3 could be run on a laptop. It was not only smaller, but it was also more

efficient. Falcon 3 became TII's flagship model. Dr. Hacid said, "It's really different from everything we've built in the past in term of data, how we trained it, and even in terms of performance."

Building a smaller model seemed to much of the team an avenue worth pursuing. Dr. Hacid said, "It's great to have leaderboard-topping LLMs, but at the end of the day what's most important is that these models lend themselves well to the community's needs and helps solve real world problems." And the open-source community seemed to prefer smaller models. He said, "We will of course continue to work on large models, too, but smaller models will be part of the focus; why should we put so much effort into building a double-decker bus if all the community really wants is a four-door sedan?"

There were other obstacles that were slowing adoption. Dr. Hacid explained, "It's not enough to have our models available to the community to find and download. We need partners to help us push these models to potential users. They can help by repackaging the models into different solutions for specific needs and then sell it out to their customers." The team reached out to hyperscalers, such as Amazon Web Services, and commercial entities, including, ai71.

Falcon's Competition and Peers

Falcon's models evolved in an ever-changing landscape alongside several other open-source and closed-source LLMs from around the world. The LLM market was estimated at \$6 billion in 2024 and projected to grow to \$36 billion by 2030.⁷⁵

Top performing LLMs had generally been closed-source. They were developed and operated by private companies and included Google's Gemini, OpenAI's ChatGPT, and xAI's Grok among others.⁷⁶ Those companies had spent hundreds of millions of dollars on developing the LLMs. Nidhakou commented, "Falcon doesn't really compete with OpenAI. It doesn't make sense to compare Falcon with ChatGPT unless it becomes open source."

But open-source models were proliferating and performing. Some of these were produced by private entities, while others by governments around the world. Leading open-source LLMs included Google's Gemma, Meta's Llama and others. DeepSeek became one of the most prominent open-source LLMs when it burst onto the scene in January 2025.⁷⁷ It caused a stir as, not only was it developed by a lesser-known Chinese startup, but it was also reported to be as capable as leading closed-source LLMs, built at a fraction of the cost⁷⁸ and despite U.S. restrictions on selling NVIDIA chips to China.⁷⁹

Reactions to DeepSeek's model varied. Analysts expected investors would rethink betting on closed-source models to win the AI race.⁸⁰ NVIDIA's stock dropped by \$600 billion as investors wondered if demand expectations were still relevant, though reached new highs six months later.⁸¹ Asked whether DeepSeek was "competition" for Falcon, Dr. Hacid commented, "We've always believed open-source models would eventually be on par with closed-source ones. We were excited to see DeepSeek prove it. It's really helped encourage the whole open-source community to push the boundaries further."

A survey conducted by the prominent US-based venture capital firm Andreessen Horowitz suggested open-source LLMs were becoming more popular among enterprises, and that nearly half of the respondents would switch to using open-source LLMs if they matched the performance of closed-source LLMs.⁸²

ai71

As the world became more familiar with Falcon, the team saw opportunities to try to boost its adoption and monetize it. “Getting Falcon into people’s hands was one thing. Making sure it had a real impact was another,” Nidhakou said. “We wanted to solve big problems, have a meaningful impact to society, and generate an economic benefit for us.”

The first avenue was VentureOne. Founded in 2022, its handful of specialized subsidiaries served several sectors namely autonomous vehicles, climate technology, cybersecurity, defense, logistics, and others.⁸³ Dr. Hacid said, “The idea was to have a dedicated team of experts focused on repackaging Falcon into different solutions and then selling it, rather than having TII’s manpower also working on commercialization.”

ai71 was one of the subsidiaries. It operated in a competitive environment with other sovereign and private developers of AI for governments and consultancies. ai71 promoted its access to large, anonymized datasets from the UAE and decentralized data ownership.⁸⁴ The latter meant clients would retain ownership of their data. The team felt this was a “game changer” for entities focused on privacy.⁸⁵

Products

For ai71, the first piece of the puzzle was identifying which sectors to target. The team began searching for difficult problems they could solve in high-GDP sectors with little competition. Initially, they focused on the medical, educational, and legal sectors.⁸⁶ Each of the early products were built using Falcon LLMs initially but not exclusively and, by 2025, were still in development.

Cons71 was envisioned as an AI-powered construction solution that would help automate the issuance of building permits by checking designs against relevant rules and regulations; an otherwise mainly manual process. Later it would be expanded and renamed SuperHive.

Razi71 was developed as an AI-enabled healthcare solution to review and summarize medical literature, transcribe case notes, cross-check medications, and more. Its purpose was to free up healthcare professionals so they could spend more time engaging with patients.

Ask71 was conceived as an AI-powered assistant that could use knowledge across an organization to generate insights and improve efficiency. Ask71, also known as AskGOV, could for example help answer an employee’s questions on specific human resources policies like how many leave days they were entitled to and then apply for their holiday. It was also intended to function as an assistant to business leaders, helping them leverage data and insights across their organizations to answer difficult questions.

Platform

In addition to these products, one of ai71’s most recent offerings was the **ai71 API Hub**. It was launched in 2024 as a platform that allowed pay-as-you-go access to Falcon and other AI models. The idea was to make it easier for users – and not just business or government clients – to use big models like Falcon 180b without needing costly infrastructure.

ai71 planned to start with Abu Dhabi government entities and then sell products to customers all over the world. Chiara Marcati, ai71’s chief AI advisory and business officer, said, “When we develop a new product, we first deploy it with a small group of trusted design partners. These organizations

are close enough to give candid feedback and patient enough to iterate with us. This creates a fast learning loop before we scale the product more broadly.”

Services

Selling AI as a product didn’t prove seamless. While still developing them, the team at ai71 met with various potential customers and stakeholders to hone the product design. “We started engaging with potential clients and found out that, although they were eager to use Falcon and AI in general, they didn’t know where to start,” said Nidhakou.

They formed an advisory arm at ai71. Nidhakou explained, “We realized that different organizations had differing abilities to adopt AI, and that we owed it to the ecosystem to advise them on how to best deploy AI in their organizations.” The advisory arm would help clients identify ways they could benefit from AI in their organizations, develop the AI solutions, and help with implementation. “Our advisory engagements are not separate from our product journey. They allow us to understand, in real time, the operational problems organizations face and translate those directly into product capabilities.” Marcati added, “The entities we work with are also central to AI decision making across the UAE government, which gives us a uniquely fast feedback loop and access to high quality, real operational data to continuously improve our products while solving mission-critical problems. In addition, it allows us to get deep into the adoption challenges that enterprises are facing.”

One of those clients was the DGE, the entity responsible for TAMM. ai71 was working with them to add new dimensions to TAMM’s conversational abilities.⁸⁷ Initially, the plan was to use Falcon. However, TAMM’s chatbot was ultimately powered by Microsoft Azure and OpenAI. Nidhakou commented, “We recommend Falcon when it’s the right tool for the job, and sometimes it isn’t. It’s in the best interest of Falcon to tell clients if we think they’re better off using another LLM and then come back to them a few months later when Falcon is ready for their specific use case.”

SuperHive

By January 2025, several potential clients showed interest in the construction product, and the ai71 team began focusing their efforts on developing it further.

SuperHive was created to address the often lengthy process of obtaining a permit from government to construct or alter a building. The team at ai71 learned that the approval process for a building permit in Abu Dhabi took at least one month. “That was assuming it was a simple design, like a villa,” explained Klemensas Mecejus, head of advisory at ai71. “Designs for complex projects like skyscrapers would take even longer to approve.” Mecejus added, “An application would typically have 20 or so documents that needed to be assessed for compliance against hundreds of rules and regulations, and checking compliance with these rules was not necessarily straightforward.” Mecejus continued, “The regulations aren’t organized as a checklist. Rather they are often written in paragraphs of text. It’s not always easy to recognize the individual rules.” In addition, because there were so many rules to check, the designs for a building were reviewed by multiple people which stretched the process out further.

SuperHive was at first to be a tool that understood all of the government’s construction-related rules and regulations and could quickly assess to what extent the plans for a building were compliant. ai71 targeted stakeholders on both sides. Developers could test and adjust their designs before even submitting them to the government. Regulatory authorities could be relieved of or supplement the manual checking of designs.

The backlog of permitting across the world and the demand for housing and new construction broadly had drawn ai71's interest into this space and also others'. In the U.S., the startup CivCheck had launched an "AI plan review platform" they said was "capable of supporting every check required for building permits and planning approvals."⁸⁸ It was designed to "augment rather than replace staff, with every AI interpretation reviewed and approved by staff for the most accurate and trustworthy results."⁸⁹ CivCheck asserted they could reduce permit approval times by 80% or more.⁹⁰ Scale AI, which had started as data annotation company, but expanded quite substantially across AI products and services, was offering "agentic building permit validation" which promised to "automate the assessment of complex construction plans for compliance with building codes and safety standards."⁹¹

To flourish, SuperHive would need to succeed on a few fronts. For starters, the tool had to know all the relevant construction rules and regulations. It also had to be capable of receiving and processing several types of files and documents ranging from relatively simple Word and PDF files to complex AutoCAD files that showed two-dimensional (2D) blueprints of buildings, to even more complex Building Information Modeling (BIM) files with three-dimensional (3D) renderings of buildings. "And then," Mecejus added, "It needed to be able to analyze those files, assess how compliant the design was, and point out exactly where it failed to comply and why."

The team used the insights from their interactions with clients to build SuperHive. Mecejus explained, "We used a combination of LLMs and natural language processing to develop a classification algorithm we call 'RuleBuilder' to reorganize all the relevant regulations into a checklist." With the checklist in hand, the team needed to get SuperHive to understand building designs. For 2D, they used a computer vision algorithm which they had trained. For 3D, they used a simpler approach. Mecejus explained, "The 3D files always come from BIM, which means they have a standardized format. That made it relatively easy to create our own rules into the product on how to digest those files." The team then taught the model to assess compliance. Mecejus commented, "We set up proper computations between the algorithms so that a specific rule could be assessed against 2D and 3D files and pinpoint where there were issues in the design."

SuperHive employed a variety of AI models. Marcati said, "When it comes to AI models, there is no one-size-fits-all model. We use Falcon, Llama, ChatGPT, and others since they all excel at different things." The decision on what AI to use depended on the task at hand and other considerations. Marcati elaborated, "An important determining factor is to do with the client's limitations. Abu Dhabi's laws, for example, necessitate local hosting of AI servers, so we're restricted to Falcon, Prem or OpenAI."

SuperHive provided the user with a report on how compliant a design was and explained which regulations were not met and why. (See **Exhibit 5** for an example). The report also highlighted which parts of the drawings could not be assessed. This occasionally occurred if SuperHive was unable to understand a design or parts of it.

Evaluating SuperHive

By the spring of 2025, the ai71 team was gathering an early sense of the product's performance and its prospects.

They'd been accumulating a pipeline of clients for SuperHive, governments and otherwise. Marcati said, "Although SuperHive was born in construction, the underlying computer vision engine is modality agnostic." The idea was that the logic SuperHive used to assess compliance could be applied to other industries. Marcati said, "It interprets structured visual artifacts against regulatory standards."

Whether the input is an architectural drawing, an electrical plan, or another technical schematic, the logic remains the same.”

And they saw a roadmap for an expansive platform to other parts of the construction value chain that SuperHive could support. These included 1) investment and strategic planning; 2) design and planning; 3) construction; and 4) post-construction (i.e. operating, maintaining, and then demolishing the building). Marcati said, at first SuperHive “only addressed the design stage, but we want to make a tool that is a one-stop-shop for construction leaders. We want it to be the ERP equivalent for the construction industry; something seamlessly integrated with all the popular specialized tech in that field.” They’d begun developing the capabilities for “digital twins”.

They were seeing high accuracy rates. At their offices, when they measured their AI’s assessments vs. what the permitting rules dictated was allowable, they were finding 90% accuracy. Mecejus added some nuance to their evaluations, “Not all rules are created equally. Some regulations are more important than others. For example, ensuring that a structural column would be capable of holding up a ceiling was far more critical than making sure that electrical sockets were exactly the right number of inches away from each other.” Critical rules were identified and prioritized. Mecejus said, “Those had to be assessed perfectly, every time, no matter what.”

There were other measures that mattered as well. Like speed. By March 2025, their AI permitting tool was able to process drawings in a minute, substantially faster than humans, and the team were confident they could eventually get it to be instantaneous. Whether humans should be “in the loop” was up for debate in AI circles, but the team still recommended so. Marcati said, “Clients were fairly open to using the AI because the value it offered was clear, but they were apprehensive about totally letting go of control.” Insisting on having a human review SuperHive’s output meant that it was likely the approval process would take significantly longer than only a minute. “But,” Mecejus said, “that meant it might take a day or two to review a design. That was still better than months. Plus, it was only one person in the review process. The others could be freed up for other tasks.”

One measure that seemed of growing importance to clients was compute demands. Mecejus said, “I began my career as a machine learning engineer in 2012. Back then, people never wondered how large the code was for a certain application. But clients are prioritizing this now.” So, ai71’s project milestones include getting the AI to reach a certain level of optimization. “That way,” Marcati explained, “clients can contain the cost of AI by avoiding costly upgrades to their AI infrastructure in the future.”

Certainly the “intelligent construction suite” the team were building SuperHive into had to deliver for the stakeholders they were targeting. For governments, they were promising streamlined monitoring and automated regulatory checks. For developers and investors, they were offering accelerated timelines and maximized project returns. For construction contractors, they were pledging help staying on schedule. And for consultants and engineers, they were promoting faster and fewer revisions with confidence in code compliance.

The question was whether to give their own permitting application the go-ahead, and how would they know?

Exhibit 1 Overview of the UAE National AI Strategy

Objective	Planned Actions
Build reputation as an AI destination	To create a new AI certification system: UAE Seal of Approval To become a global AI Hub by hosting conferences, and becoming a center for AI startups
Develop competitive assets in priority sectors	To provide government incentives and advice for local firms To encourage global AI leaders to locate in UAE
Develop an active AI ecosystem	To provide government incentives for local and international firms, funding to startups and set up accelerators To develop a coordinated network of researchers, industry experts, and regulators To attract FDI
Adopt AI across customer services in government	To solicit ideas from government, universities, and private sector to solve key societal problems To set up a council to identify where and how AI can be used in government
Attract and develop talent	To provide AI courses dedicated to government staff and incorporate AI into education
World-leading research capability to collaborate with key industries	To create world class institutions to aggregate local and international expertise (e.g. National Virtual AI Institute) To create the Key Thinkers Program to attract leading AI minds To create an open-access digital AI-library of research
Become test bed for AI by providing data and infrastructure	To leverage UAE's diverse population – a unique source of data, particularly useful for health research To develop a secure data infrastructure that is easily accessible
Focus on governance and regulation	To take a leading role in governance To collaborate with other governments and thought leaders To run pilots to test theories

Source: Created by the case writers using information from “UAE National Strategy for Artificial Intelligence 2031,” UAE Minister of State for AI, Digital Economy, & Remote Work Applications, <https://ai.gov.ae/strategy/#opportunities1>, accessed May 21, 2025.

Exhibit 2 TII's Research Centers

Advanced Materials

AI and Digital Science

Autonomous Robotics

Biotechnology

Cryptography

Directed Energy

Propulsion and Space

Quantum

Renewable and Sustainable Energy

Source: Prepared by the case writers using company documents.

Exhibit 3 Timeline of Falcon Models

2023

Falcon 40B

Falcon 7B

Falcon 180B

2024

Falcon Mamba 7B

Falcon-E (Edge Series)

Falcon 2 – 11B

Falcon 2 – 11B VLM

Falcon 3 (family)

2025


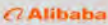



Falcon Arabic






Falcon-H1 (family)

Falcon H1R 7B

Source: Prepared by the case writers using company documents.

Exhibit 4 Falcon Benchmarking

	Base models	MUSR	BBH	MMLU_PRO	If_eval	GPQA	MATH	Avg
	Falcon3 10B	14.17	41.38	36	36.48	12.75	24.77	27.59
	Falcon3 7B	18.14	31.56	32.34	34.16	12.86	19.26	24.72
	Qwen2.5 7B	14.14	35.81	37.39	33.74	9.96	18.88	24.99
	Qwen2 7B	14.32	34.71	35.37	31.49	7.27	20.47	23.94
	Gemma2 9B	14.3	34.1	34.48	20.4	10.51	13.14	21.15
	Mistral-nemo-base-2407 (12B)	6.52	29.37	27.46	16.3	5.82	5.36	15.14
	Llama3.1 8B	8.98	25.29	24.95	12.7	6.15	5.14	13.87

	Instruct models	MUSR	BBH	MMLU_PRO	If_eval	GPQA	MATH	Avg
	Falcon3 10B Instruct	13.61	44.82	38.1	78.17	10.51	25.91	35.19
	Falcon3 7B Instruct	21.17	37.92	34.3	76.12	8.05	31.87	34.91
	Llama3.1 8B Instruct	8.41	29.89	30.68	78.56	2.35	19.34	28.2
	Qwen2.5 7B Instruct	8.45	34.89	36.52	75.85	5.48	0	26.87
	Qwen2 7B Instruct	7.37	37.81	31.64	56.79	6.38	9.44	24.9
	Mistral-Nemo-Instruct-2407 (12B)	8.48	29.68	27.97	63.8	5.37	6.5	23.63
	Gemma2 9B It	13.77	21.62	20.48	50.1	2.68	6.72	19.23

Source: Company documents. <https://falconlm.tii.ae/falcon3/index.html>, accessed January 2026.

Note: “MuSR (Multistep Soft Reasoning) is a benchmark for evaluating language models on multistep soft reasoning tasks specified in natural language narratives... The benchmark tests chain-of-thought reasoning capabilities across domains involving commonsense reasoning about physical and social situations.”

“Big-Bench Hard (BBH) is a suite of 23 challenging tasks selected from BIG-Bench for which prior language model evaluations did not outperform the average human-rater... focuses on evaluating complex reasoning skills including temporal understanding, spatial reasoning, causal understanding, and deductive logical reasoning.”

“MMLU-RPO (Measuring Massive Multitask Language Understanding PRO) extends MMLU by expanding multiple-choice options from 4 to 10, eliminating trivial questions, and focusing on reasoning-intensive tasks. Features over 12,000 curated questions across 14 domains and causes a 16-33% accuracy drop compared to original MMLU.”

“IFEVAL (Instruction-Following Evaluation) benchmark is for large language models, focusing on verifiable instructions with 25 types of instructions and around 500 prompts containing one or more verifiable constraints”

“GPQA (Graduate-Level Google Proof Q&A) is a challenging dataset of 448 multiple-choice questions written by domain experts in biology, physics, and chemistry. Questions are Google-proof and extremely difficult, with PhD experts reaching 65% accuracy.”

“MATH dataset contains 12,500 challenging competition mathematics problems from AMC 10, AMC 12, AIME, and other mathematics competitions. Each problem includes full step-by-step solutions and spans multiple difficulty levels (1-5) across seven mathematical subjects.”

Instruct models have been fine-tuned for following instructions and task completion.

Benchmark explanations from llm-stats.com, <https://llm-stats.com/benchmarks>, accessed January 2026.

Exhibit 5 Sample SuperHive Output

The residential plot contains the following components: Residential villa, Appendices of Services, Hospitality Annex, Sports supplement, Swimming pool, Garage, Electricity room, transformers, and sub-electricity, Pool pump room, Site formatting elements, External fences and entrances for individuals and cars, Light shades and parking awnings.	Compliant
All services on the Roof floor & Top Roof floor must be concealed from view from all angles.	Not Compliant
Maximum plot coverage for single unit detached villa should not exceed 70% from the plot area, inclusive of the residential villa and all other structures. Aesthetic protrusions not exceeding 50 cm are exempt.	Compliant

Source: Company documents.

Endnotes

- ¹ “About us.” ai71. <https://ai71.ai/about-us>, accessed October 2025.
- ² Ashley Petersen and Jeff Quigley from Stratas Advisors, “What Affected Oil Prices: 2017 In Review,” Hart Energy, December 29, 2017, <https://www.hartenergy.com/exclusives/what-affected-oil-prices-2017-review-30643>, accessed August 15, 2025.
- ³ “Policy Report: The U.A.E.’s Big Bet on Artificial Intelligence,” U.S.-U.A.E. Business Council, February 2024, PDF downloaded from <https://usuaebusiness.org/resources/ai-report/>, accessed May 9, 2025.
- ⁴ “History,” Embassy of the United Arab Emirates in Washington D.C., <https://www.uae-embassy.org/discover-uae/history#:~:text=An%20agreement%20was%20reached%20between,established%20on%20December%202%2C%201971.>, accessed May 15, 2025.
- ⁵ Dr. Mohammad Salami, “The UAE’s OPEC+ Dilemma,” Gulf International Forum, <https://gulfig.org/the-uae-opecc-dilemma/>, accessed May 15, 2025.
- ⁶ “UAE,” Allianz Trade, January 2025, https://www.allianz-trade.com/en_global/economic-research/country-reports/United-Arab-Emirates.html, accessed May 15, 2025.
- ⁷ “UAE,” Allianz Trade, January 2025, https://www.allianz-trade.com/en_global/economic-research/country-reports/United-Arab-Emirates.html, accessed May 15, 2025.
- ⁸ “UAE,” Allianz Trade, January 2025, https://www.allianz-trade.com/en_global/economic-research/country-reports/United-Arab-Emirates.html, accessed May 15, 2025.
- ⁹ “UAE economy grows 3.8% in first nine months of 2024, driven by non-oil sectors,” Gulf News, March 5, 2025, https://gulfnews.com/business/economy/uae-economy-grows-38-in-first-nine-months-of-2024-driven-by-non-oil-sectors-1.500053093?utm_source=chatgpt.com, accessed May 16, 2025.
- ¹⁰ “the Annual Economic Report 2015,” UAE ministry of economy, see The Minister’s Statement on page 11, report downloaded from <https://u.ae/-/media/Annual-economic-reports/Eng-Annual-Report--2015.pdf>, accessed May 16, 2025.
- ¹¹ “Digital Economy,” U.S. International Trade Administration, <https://www.trade.gov/country-commercial-guides/uae-digital-economy>, accessed May 16, 2025.
- ¹² Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-uae/>, accessed May 9, 2025.
- ¹³ Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-uae/>, accessed May 9, 2025.
- ¹⁴ Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-uae/>, accessed May 9, 2025.
- ¹⁵ Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-uae/>, accessed May 9, 2025.
- ¹⁶ “Abu Dhabi to be the world’s first fully AI-native government across all digital services by 2027,” Department of Government Enablement, <https://www.dge.gov.ae/en/news/cx-strategy>, accessed May 9, 2025.
- ¹⁷ Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-uae/>, accessed May 9, 2025.
- ¹⁸ “OpenAI’s Sam Altman Says UAE Could Test Potential Global Rules for AI,” PYMNTS, February 13, 2024, <https://www.pymnts.com/artificial-intelligence-2/2024/openai-sam-altman-says-uae-could-test-potential-global-rules-for-ai/>, accessed May 9, 2025.
- ¹⁹ Mark Townsend, “UAE: Global Chip Makers Evaluate Building Plants In The Gulf,” Global Finance, October 6, 2024, <https://gfmag.com/technology/samsung-uae-tsmc-chip-manufacturing/#:~:text=Samsung%20Electronics%20and%20Taiwan%20Semiconductor,worth%20more%20than%20%24100%20billion.>, accessed May 15, 2025.
- ²⁰ Karen Kwok, “Gulf’s AI strategy is built on more than sand,” Reuters, November 13, 2024, <https://www.reuters.com/breakingviews/gulfs-ai-strategy-is-built-more-than-sand-2024-11-13/>, accessed May 16, 2025.

- ²¹ “G42 launches \$10b tech Expansion Fund,” G42, August 18, 2022, <https://www.g42.ai/resources/news/G42-launches-10b-tech-Expansion-Fund>, accessed August 15, 2025.
- ²² Kate Birch, “G42, the company behind the world’s largest AI supercomputer,” BusinessChief, July 24, 2023, <https://businesschief.eu/technology/g42-the-brains-behind-the-worlds-largest-ai-supercomputer>, accessed May 16, 2025.
- ²³ Robert Hart, “What To Know About G42 – The Emirati AI Giant That Just Got A \$1.5 Billion Investment From Microsoft,” Forbes, April 16, 2024, <https://www.forbes.com/sites/roberthart/2024/04/16/what-to-know-about-g42-the-emirati-ai-giant-that-just-got-a-15-billion-investment-from-microsoft/>, accessed May 9, 2025.
- ²⁴ Robert Hart, “What To Know About G42 – The Emirati AI Giant That Just Got A \$1.5 Billion Investment From Microsoft,” Forbes, April 16, 2024, <https://www.forbes.com/sites/roberthart/2024/04/16/what-to-know-about-g42-the-emirati-ai-giant-that-just-got-a-15-billion-investment-from-microsoft/>, accessed May 9, 2025.
- ²⁵ Robert Hart, “What To Know About G42 – The Emirati AI Giant That Just Got A \$1.5 Billion Investment From Microsoft,” Forbes, April 16, 2024, <https://www.forbes.com/sites/roberthart/2024/04/16/what-to-know-about-g42-the-emirati-ai-giant-that-just-got-a-15-billion-investment-from-microsoft/>, accessed May 9, 2025.
- ²⁶ Lisa Eadicicco, “Trump to change controversial Biden-era restrictions on AI chip exports,” CNN, May 8, 2025, <https://edition.cnn.com/2025/05/08/tech/trump-cancels-biden-ai-chip-curbs>, accessed May 9, 2025.
- ²⁷ “Our purpose,” Advanced Technology Research Council, <https://www.atrc.gov.ae/about-atrc>, accessed May 9, 2025.
- ²⁸ “Executive Board,” Advanced Technology Research Council, <https://www.atrc.gov.ae/executive-board>, accessed May 9, 2025.
- ²⁹ “Abu Dhabi Government launches Digital Strategy 2025-2027,” Media office of the government of Abu Dhabi, January 21, 2025, <https://www.mediaoffice.abudhabi/en/government-affairs/abu-dhabi-government-launches-digital-strategy-2025-2027/>, accessed May 9, 2025.
- ³⁰ “Abu Dhabi Government launches Digital Strategy 2025-2027,” Media office of the government of Abu Dhabi, January 21, 2025, <https://www.mediaoffice.abudhabi/en/government-affairs/abu-dhabi-government-launches-digital-strategy-2025-2027/>, accessed May 9, 2025.
- ³¹ “Abu Dhabi Government launches Digital Strategy 2025-2027,” Media office of the government of Abu Dhabi, January 21, 2025, <https://www.mediaoffice.abudhabi/en/government-affairs/abu-dhabi-government-launches-digital-strategy-2025-2027/>, accessed May 9, 2025.
- ³² “Policies and Legislations,” Department of Government Enablement, <https://www.dge.gov.ae/en/policies-and-legislations>, accessed May 9, 2025.
- ³³ “Abu Dhabi Government launches Digital Strategy 2025-2027,” Media office of the government of Abu Dhabi, January 21, 2025, <https://www.mediaoffice.abudhabi/en/government-affairs/abu-dhabi-government-launches-digital-strategy-2025-2027/>, accessed May 9, 2025.
- ³⁴ “Abu Dhabi Government launches Digital Strategy 2025-2027,” Media office of the government of Abu Dhabi, January 21, 2025, <https://www.mediaoffice.abudhabi/en/government-affairs/abu-dhabi-government-launches-digital-strategy-2025-2027/>, accessed May 9, 2025.
- ³⁵ “Abu Dhabi Government launches Digital Strategy 2025-2027,” Media office of the government of Abu Dhabi, January 21, 2025, <https://www.mediaoffice.abudhabi/en/government-affairs/abu-dhabi-government-launches-digital-strategy-2025-2027/>, accessed May 9, 2025.
- ³⁶ Chris Welsch, “Consider it done: How TAMM is transforming government services in Abu Dhabi with AI,” Microsoft, <https://news.microsoft.com/source/emea/features/tamm-app-abu-dhabi-government-services/>, accessed May 9, 2025.
- ³⁷ Chris Welsch, “Consider it done: How TAMM is transforming government services in Abu Dhabi with AI,” Microsoft, <https://news.microsoft.com/source/emea/features/tamm-app-abu-dhabi-government-services/>, accessed May 9, 2025.
- ³⁸ “Service on TAMM,” Department of Culture and Tourism, <https://dct.gov.ae/en/e.services/services.on.tamm.aspx>, accessed May 15, 2025.
- ³⁹ Chris Welsch, “Consider it done: How TAMM is transforming government services in Abu Dhabi with AI,” Microsoft, <https://news.microsoft.com/source/emea/features/tamm-app-abu-dhabi-government-services/>, accessed May 9, 2025.

- ⁴⁰ Chris Welsch, "Consider it done: How TAMM is transforming government services in Abu Dhabi with AI," Microsoft, <https://news.microsoft.com/source/emea/features/tamm-app-abu-dhabi-government-services/>, accessed May 9, 2025.
- ⁴¹ "Department of Municipalities and Transport Signs Four Strategic Memorandums of Understanding during GITEX Global 2024," Department on Municipalities and Transport, October 16, 2024, <https://www.dmt.gov.ae/en/Media-Centre/News/DMT-Signs-Four-Strategic-Memorandums-of-Understanding-during-GITEX-Global-2024>, accessed May 12, 2025.
- ⁴² "Abu Dhabi works with Microsoft to use AI in the energy industry," Investment Monitor, November 4, 2024, <https://www.investmentmonitor.ai/news/abu-dhabi-works-with-microsoft-to-use-ai-in-the-energy-industry/?cf-view>, accessed May 14, 2025.
- ⁴³ "Abu Dhabi Uses AI to Automatically Slash Traffic Fines," Emirates 24/7,
- ⁴⁴ Angie Lee, "What is sovereign AI?" NVIDIA, February 28, 2024, <https://blogs.NVIDIA.com/blog/what-is-sovereign-ai/>, accessed May 12, 2025.
- ⁴⁵ Pablo Chavez, "Sovereign AI in a Hybrid World: National Strategies and Policy Responses," Lawfare, November 7, 2024, <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses>, accessed May 12, 2025.
- ⁴⁶ Pablo Chavez, "The rising tide of sovereign AI," SiliconANGLE, February 10, 2024, <https://siliconangle.com/2024/02/10/rising-tide-sovereign-ai/>, accessed May 12, 2025.
- ⁴⁷ NVIDIA. "NVIDIA CEO: Every Country Needs Sovereign AI." <https://blogs.nvidia.com/blog/world-governments-summit/>, accessed January 2026.
- ⁴⁸ Brian Letort, "What is sovereign AI and why is it growing in importance?," Digital Realty, April 3, 2025, <https://www.digitalrealty.com/resources/articles/what-is-sovereign-ai>, accessed May 12, 2025.
- ⁴⁹ Angie Lee, "What is sovereign AI?" NVIDIA, February 28, 2024, <https://blogs.NVIDIA.com/blog/what-is-sovereign-ai/>, accessed May 12, 2025.
- ⁵⁰ Brian Letort, "What is sovereign AI and why is it growing in importance?," Digital Realty, April 3, 2025, <https://www.digitalrealty.com/resources/articles/what-is-sovereign-ai>, accessed May 12, 2025.
- ⁵¹ Brian Caulfield, "NVIDIA CEO: Every Country Needs Sovereign AI," NVIDIA, February 12, 2024, <https://blogs.NVIDIA.com/blog/world-governments-summit/>, accessed May 12, 2025.
- ⁵² Brian Caulfield, "NVIDIA CEO: Every Country Needs Sovereign AI," NVIDIA, February 12, 2024, <https://blogs.NVIDIA.com/blog/world-governments-summit/>, accessed May 12, 2025.
- ⁵³ Pablo Chavez, "Sovereign AI in a Hybrid World: National Strategies and Policy Responses," Lawfare, November 7, 2024, <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses>, accessed May 12, 2025.
- ⁵⁴ "revealing the path forward with sovereign LLMs," Deloitte, pdf downloaded from <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-NVIDIA-revealing-the-path-forward-with-sovereign-llms.pdf>, accessed May 12, 2025.
- ⁵⁵ Brian Letort, "What is sovereign AI and why is it growing in importance?," Digital Realty, April 3, 2025, <https://www.digitalrealty.com/resources/articles/what-is-sovereign-ai>, accessed May 12, 2025.
- ⁵⁶ Brian Letort, "What is sovereign AI and why is it growing in importance?," Digital Realty, April 3, 2025, <https://www.digitalrealty.com/resources/articles/what-is-sovereign-ai>, accessed May 12, 2025.
- ⁵⁷ Pablo Chavez, "The rising tide of sovereign AI," SiliconANGLE, February 10, 2024, <https://siliconangle.com/2024/02/10/rising-tide-sovereign-ai/>, accessed May 12, 2025.
- ⁵⁸ Pablo Chavez, "Sovereign AI in a Hybrid World: National Strategies and Policy Responses," Lawfare, November 7, 2024, <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses>, accessed May 12, 2025.
- ⁵⁹ Pablo Chavez, "Sovereign AI in a Hybrid World: National Strategies and Policy Responses," Lawfare, November 7, 2024, <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses>, accessed May 12, 2025.

- ⁶⁰ “National AI Policies & Strategies,” OECD.AI Policy Observatory, https://oecd.ai/en/dashboards/overview?utm_source=chatgpt.com, accessed May 15, 2025.
- ⁶¹ “National Strategy for Data and AI,” Saudi Data and AI Authority, <https://sdaia.gov.sa/en/SDAIA/SdaiaStrategies/Pages/NationalStrategyForDataAndAI.aspx>, accessed May 15, 2025.
- ⁶² Purushottam Kaushik, Harsh Sharma, Ayushi Sarna, “AI for India 2030: A blueprint for inclusive growth and global leadership,” World Economic Forum, January 22, 2025, <https://www.weforum.org/stories/2025/01/ai-for-india-2030-blueprint-inclusive-growth-global-leadership/#:~:text=AI%20is%20not%20merely%20a,drive%20growth%20in%20underserved%20areas.>, accessed May 15, 2025.
- ⁶³ “Global AI Power Rankings: Stanford HAI Tool Ranks 36 Countries in AI,” HAI Stanford University, November 21, 2024, <https://hai.stanford.edu/news/global-ai-power-rankings-stanford-hai-tool-ranks-36-countries-in-ai>, accessed May 14, 2025.
- ⁶⁴ “Abu Dhabi-based TII launches its own version of ChatGPT,” Fast Company Middle East, March 17, 2023, <https://fastcompanyme.com/news/abu-dhabi-based-tii-launches-its-own-version-of-chatgpt/>, accessed May 7, 2025.
- ⁶⁵ “Abu Dhabi-based TII launches its own version of ChatGPT,” Fast Company Middle East, March 17, 2023, <https://fastcompanyme.com/news/abu-dhabi-based-tii-launches-its-own-version-of-chatgpt/>, accessed May 7, 2025.
- ⁶⁶ “The Falcon has landed in the Hugging Face ecosystem,” Hugging Face, June 5, 2023, <https://huggingface.co/blog/falcon#conclusion>, accessed May 7, 2025.
- ⁶⁷ Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-united-arab-emirates/>, accessed May 9, 2025.
- ⁶⁸ “UAE’s Technology Innovation Institute Launches Open-Source “Falcon 40B” Large Language Model for Research & Commercial Utilization,” TII, May 25, 2023, <https://www.tii.ae/news/uaes-technology-innovation-institute-launches-open-source-falcon-40b-large-language-model#:~:text=Abu%20Dhabi%20UAE%3A%2025%20May,of%20artificial%20intelligence%20by%20announcing>, accessed May 16, 2025.
- ⁶⁹ Kif Leswing, “Meet the \$10,000 Nvidia chip powering the race for A.I.,” CNBC, February 23, 2023, <https://www.cnbc.com/2023/02/23/nvidias-a100-is-the-10000-chip-powering-the-race-for-ai-.html>, accessed May 16, 2025.
- ⁷⁰ “UAE’s Technology Innovation Institute Launches Open-Source “Falcon 40B” Large Language Model for Research & Commercial Utilization,” TII, May 25, 2023, <https://www.tii.ae/news/uaes-technology-innovation-institute-launches-open-source-falcon-40b-large-language-model#:~:text=Abu%20Dhabi%20UAE%3A%2025%20May,of%20artificial%20intelligence%20by%20announcing>, accessed May 16, 2025.
- ⁷¹ Kif Leswing, “Meet the \$10,000 Nvidia chip powering the race for A.I.,” CNBC, February 23, 2023, <https://www.cnbc.com/2023/02/23/nvidias-a100-is-the-10000-chip-powering-the-race-for-ai-.html>, accessed May 16, 2025.
- ⁷² Gwen Davis, “A developer’s guide to open source LLMs and generative AI,” GitHub blog, May 20, 2024, <https://github.blog/ai-and-ml/llms/a-developers-guide-to-open-source-llms-and-generative-ai/#:~:text=%E2%80%9CThey’re%20often%20far%20more,fin%20tune%20open%20source%20LLMs>, accessed August 15, 2025.
- ⁷³ Billy Perrigo, “The UAE is on a mission to become an AI power,” Time, March 22, 2024, <https://time.com/6958369/artificial-intelligence-united-arab-emirates/>, accessed May 9, 2025.
- ⁷⁴ “Falcon Models,” TII, <https://falconllm.tii.ae/falcon-models.html>, accessed May 7, 2025.
- ⁷⁵ “Global Large Language Model (LLM) Market Research 2024-2030 - 100 Billion to 200 Billion Parameter Segment is Expected to Register the Fastest Growth Rate - ResearchAndMarkets.com,” Businesswire, April 12, 2024, <https://www.businesswire.com/news/home/20240412489506/en/Global-Large-Language-Model-LLM-Market-Research-2024-2030---100-Billion-to-200-Billion-Parameter-Segment-is-Expected-to-Register-the-Fastest-Growth-Rate---ResearchAndMarkets.com>, accessed May 16, 2025.
- ⁷⁶ “Chatbot Arena LLM Leaderboard: Community-driven Evaluation for Best LLM and AI chatbots,” HuggingFace, May 11, 2025, <https://huggingface.co/spaces/lmarena-ai/chatbot-arena-leaderboard>, accessed May 15, 2025.
- ⁷⁷ Amy Hawkins, “Who is behind DeepSeek and how did it achieve its AI ‘Sputnik moment’?,” The Guardian, January 28,

2025, <https://www.theguardian.com/technology/2025/jan/28/who-is-behind-deepseek-and-how-did-it-achieve-its-ai-sputnik-moment>, accessed May 15, 2025.

⁷⁸ Amy Hawkins, "Who is behind DeepSeek and how did it achieve its AI 'Sputnik moment'?", The Guardian, January 28, 2025, <https://www.theguardian.com/technology/2025/jan/28/who-is-behind-deepseek-and-how-did-it-achieve-its-ai-sputnik-moment>, accessed May 15, 2025.

⁷⁹ Amy Hawkins, "Who is behind DeepSeek and how did it achieve its AI 'Sputnik moment'?", The Guardian, January 28, 2025, <https://www.theguardian.com/technology/2025/jan/28/who-is-behind-deepseek-and-how-did-it-achieve-its-ai-sputnik-moment>, accessed May 15, 2025.

⁸⁰ Leo Shwartz, Allie Garfinkle, "After pouring billions into AI, venture capital faces a reckoning with DeepSeek breakout: 'It's a gut check'," Fortune, January 28, 2025, <https://fortune.com/2025/01/28/venture-capital-thrive-deepseek-openai-anthropic-lightspeed/>, accessed May 15, 2025.

⁸¹ Leo Shwartz, Allie Garfinkle, "After pouring billions into AI, venture capital faces a reckoning with DeepSeek breakout: 'It's a gut check'," Fortune, January 28, 2025, <https://fortune.com/2025/01/28/venture-capital-thrive-deepseek-openai-anthropic-lightspeed/>, accessed May 15, 2025.

⁸² Melissa Malec, "Open-Source LLMs vs Closed: Unbiased Guide for Innovative Companies [2025]," January 28, 2025, <https://hatchworks.com/blog/gen-ai/open-source-vs-closed-llms-guide/#:~:text=Reports%20have%20emerged%20showing%20that,in%20place%20of%20closed%20models.>, accessed May 15, 2025.

⁸³ "Abu Dhabi's Advanced Technology Research Council launches 'ai71': New AI Company Pioneering Decentralized Data Control for Companies & Countries," VentureOne, November 28, 2023, <https://www.ventureone.ae/news/abu-dhabis-advanced-technology-research-council-launches-ai71-new-ai-company-pioneering>, accessed May 16, 2025

⁸⁴ "Abu Dhabi's Advanced Technology Research Council launches 'ai71': New AI Company Pioneering Decentralized Data Control for Companies & Countries," VentureOne, November 28, 2023, <https://www.ventureone.ae/news/abu-dhabis-advanced-technology-research-council-launches-ai71-new-ai-company-pioneering>, accessed May 16, 2025.

⁸⁵ "Abu Dhabi's Advanced Technology Research Council launches 'ai71': New AI Company Pioneering Decentralized Data Control for Companies & Countries," VentureOne, November 28, 2023, <https://www.ventureone.ae/news/abu-dhabis-advanced-technology-research-council-launches-ai71-new-ai-company-pioneering>, accessed May 16, 2025.

⁸⁶ "Abu Dhabi's Advanced Technology Research Council launches 'ai71': New AI Company Pioneering Decentralized Data Control for Companies & Countries," VentureOne, November 28, 2023, <https://www.ventureone.ae/news/abu-dhabis-advanced-technology-research-council-launches-ai71-new-ai-company-pioneering>, accessed May 16, 2025.

⁸⁷ "TAMM Partners with Newly Launched ai71 to Become First Platform to Use UAE's Falcon Series of Large Language Models," Department of Government Enablement, November 30, 2023, <https://www.dge.gov.ae/en/news/atrc-x-tamm>, accessed May 15, 2025.

⁸⁸ "Clariti Acquires CivCheck to Accelerate Permit Approvals with AI." CivCheck, October 7, 2025, <https://www.civcheck.ai/blog/clariti-acquires-civcheck-to-accelerate-permit-approvals-with-ai>, accessed January 2026.

⁸⁹ "Clariti Acquires CivCheck to Accelerate Permit Approvals with AI." CivCheck, October 7, 2025, <https://www.civcheck.ai/blog/clariti-acquires-civcheck-to-accelerate-permit-approvals-with-ai>, accessed January 2026.

⁹⁰ "Clariti Acquires CivCheck to Accelerate Permit Approvals with AI." CivCheck, October 7, 2025, <https://www.civcheck.ai/blog/clariti-acquires-civcheck-to-accelerate-permit-approvals-with-ai>, accessed January 2026.

⁹¹ "Agentic Building Permit Validation." Scale AI, <https://scale.com/enterprise/prebuilt-applications/agentic-building-permit-validation>, accessed January 2026.

2026

Connective
Leadership Summit
19 - 22 May 2026 | Melbourne VIC

Moderna: Democratizing Artificial Intelligence





IAVOR I. BOJINOV
KARIM R. LAKHANI
ANNIKA HILDEBRANDT
JAMES WEBER

Moderna: Democratizing Artificial Intelligence

In late December 2024, Vice President of AI Products and Innovation Brice Challamel met with CEO Stéphane Bancel, Chief People and Digital Technology Officer Tracey Franklin, and Chief Information Officer Brad Miller to review the adoption of generative AI at Moderna. Over the past year, the biotechnology company had provided all employees with access to OpenAI's ChatGPT Enterprise and encouraged them to incorporate the tool into their daily work. GPTs—Generative Pre-trained Transformers—were a powerful form of artificial intelligence that could reshape a variety of standard business processes. One of ChatGPT's key functionalities allowed users to create and share custom GPTs, each fine-tuned with specific instructions and data to deliver more accurate, relevant responses for particular use cases. (See **Exhibit 1** for Moderna's top 20 GPTs and **Exhibit 2** for GPT usage activity.)

From the outset, Moderna had been a digital-first, AI-focused company. Bancel famously described it as “a technology company that happens to do biology.” By 2024, Moderna aimed to obtain 10 new drug approvals within three years. Bancel believed that sustained AI-driven innovation would enable the company's nearly 6,000 employees to keep pace with rival pharmaceutical firms employing more than 100,000 people. To foster this innovation, the company encouraged employees to develop, publish, and maintain custom GPTs, embracing a model akin to the Apple App Store or Google Play Store, where employees could share their creations with each other. Yet AI was not without its flaws. Employees were still learning to wield these emerging tools, and GPTs sometimes produced inaccuracies—so-called “hallucinations.” Challamel recognized that as a publicly traded and highly regulated pharmaceutical company, GPT errors in critical processes could have serious consequences for Moderna. To balance risk management with speed and innovation, he implemented governance practices for AI use.

As the Moderna leadership team discussed generative AI adoption, concerns about the use and governance of custom GPTs began to resurface. During the meeting, Challamel highlighted the recent spike in usage of the Self-Review GPT, a tool assisting employees with quarterly and annual performance reviews. Franklin expressed concern: “I'm worried that the Self-Review GPT is potentially problematic, as it is augmenting—and to some extent replacing—a critical process in developing employees. Maybe there are some processes and work that should be kept off-limits? How can Moderna lead the way in pioneering human-AI augmentation in all the work that gets done here?”

Professors Iavor I. Bojinov and Karim R. Lakhani, Research Associate Annika Hildebrandt, and Associate Director James Weber (Case Research & Writing Group) prepared this case. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. Professor Lakhani has been a paid key opinion leader for the company. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2025 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

This document is authorized for use only by Mitchell Weiss in 2026.

Seizing on the point, Challamel turned the group's attention to a new GPT called DoseID, created by physician and medical writer Lee Quist, which provided drug dosing recommendations for clinical trials. Bancel immediately wondered if such a tool would require regulatory approval, noting that clinical trials were tightly regulated. "And more generally," he added, "as we automate more of our processes, we should understand the regulatory implications." Miller agreed, cautioning, "Our approach so far has been to empower employees and let a thousand flowers bloom, but as we reach that number of custom GPTs, we may need to adjust our strategy."

As the meeting wrapped up, Challamel questioned whether Moderna should take centralized ownership of the GPTs to ensure correct development and appropriate use. Should the company impose additional oversight, or was the current governance model sufficient? Finally, did Moderna need to alert government regulators that GPTs were now part of its core processes?

Traditional Vaccine Development

Viruses were small pieces of genetic material, not living cells, which were capable of infecting a living host and then reproducing within the host's cells to cause disease. Vaccines were drugs injected into the body in advance of an infection that were designed to stimulate the body to produce antibodies within the immune system which would later attack the virus should it appear. Scientists made vaccines using a variety of techniques including developing weakened or inactive versions of a virus to be injected to create the immune response.

Traditional vaccine development, as typically practiced for drugs developed by leading pharmaceutical companies such as GSK, Merck, Johnson & Johnson, and Pfizer, involved five major steps (see **Exhibit 3**). First, basic or exploratory research: this step occurred at universities, medical centers, or biotech companies and aimed to isolate the particular virus pathogen responsible for the disease and then weakening it or making it inactive so it could be used as a vaccine without causing disease in the recipient. The research phase was often funded by government grants or private foundations. Second, preclinical trials: vaccine candidates that looked promising were tested in animals for safety, to better understand how the vaccine worked, and to gauge commercial potential.

During these first two steps, multiple labs could be pursuing similar research. Scientists presented their work at conferences and published in journals to share discoveries. Leading pharmaceutical companies followed this research and sometimes supported scientists to move the most promising vaccine candidates towards possible product development. This process could take a decade or more and many potential vaccines failed along the way.

Third, clinical trials: following pre-clinical trials, scientists, typically in partnership with a pharmaceutical company, submit an Investigational New Drug (IND) application to the U.S. Food and Drug Administration (FDA) for permission to test the vaccine in humans. The IND contained information about pre-clinical testing results, the proposed manufacturing process, and clinical testing protocols. The FDA could take 30 days to review the application. Clinical trials had three phases. In phase 1 the vaccine was tested in a small number of healthy volunteers for safety, dosage tolerances, and early effectiveness. In phase 2, hundreds of volunteers were injected, and in phase 3, several thousand volunteers were injected as safety, tolerance across a diverse range of people, and dosage levels were further refined.¹ A vaccine candidate could be withdrawn from clinical trials at any time.

After clinical trials, for a successful vaccine, the drug maker sought permission to market and sell the vaccine by submitting a Biologics License Application (BLA) to the FDA, or to other health regulators around the world. The FDA processing time for a BLA was 12 months.² The BLA applicant

must show vaccine safety and effectiveness, appropriate labeling, and proper manufacturing. Clinical trials and regulatory approval could take several years, though the process could be accelerated under certain circumstances as occurred in 2020-2021, during the COVID-19 pandemic.³

mRNA

Moderna employed a fundamentally different development process than that of leading pharmaceutical companies. Rather than starting with a weakened or inactive virus to trigger an immune response, Moderna used messenger RNA (mRNA) to create proteins within cells that would fight the virus. In addition to vaccines, Moderna also had mRNA-based therapeutic drugs in its pipeline aimed at combatting other diseases.

The human body relied on many different kinds of proteins to function. The information needed to create each protein was encoded in a cell's DNA. When a specific protein was needed, the DNA produced mRNA within the cell's nucleus. The mRNA, which carried the code to make the protein, moved out of the nucleus into the cell's cytoplasm to ribosomes that then made large quantities of the corresponding protein.

Moderna, which pioneered the mRNA approach, manufactured mRNA outside of the human body. Moderna also manufactured lipid nano-particles (LNPs), a type of fat, which it used as an mRNA delivery system. mRNA, if injected directly into the body, would break down before it was properly absorbed. By encasing the mRNA in LNP, Moderna could ensure the mRNA would pass through the bloodstream and be absorbed into cells.

Moderna saw mRNA as acting as the software of life. Like software, it carried the instructions the cells needed to produce the required proteins, while LNPs acted as hardware to deliver the software, and DNA served as the hard drive, storing the instructions.

Moderna Early Years

Moderna was founded in 2010 by Flagship Ventures, a venture capital firm looking to develop "radical innovations in medicine [. . .] in untapped spaces."⁴ Bancel joined as CEO in 2011. Flagship's idea was to see if it was possible to use mRNA to get human cells to produce their own drugs.

A Platform Company

The founders also envisioned that the mRNA technology could be used to create Moderna as a platform company. With traditional drug and vaccine development, companies essentially developed one drug at a time, investing significant capital to do so, in a process that resulted in only limited knowledge transfer from one drug to another. With mRNA, Moderna's founders saw an opportunity to develop many drugs in parallel by using similar technologies where cross-drug learning was significant. For example, a particular LNP that could carry an mRNA molecule into certain types of cells could be used to carry other mRNA molecules into those same cells. Similarly, key learnings in manufacturing one mRNA or LNP could transfer to the manufacture of others. This cross-learning could also enable Moderna to bring new drugs from discovery to market more quickly, potentially years faster than traditional methods.

Digital-First

Prior to joining Moderna, Bancel, who had spent his career in life science companies, had seen both the promise of emerging digital technologies and the problems when such technologies were poorly deployed, not integrated, and lacked data. He stated, “I wanted Moderna to be a digital company from day one. For that to happen, we needed the [information technology] to be built right, even if it meant considerable investments at a time when we didn’t have revenue streams. Digitizing right from the get-go is much easier than doing this *ex post* on a legacy system.”⁵

Bancel began by moving from manual processes to software applications tracking research at a time when the company had only a few dozen employees. In 2015, he hired Marcello Damiani as Moderna’s first chief digital office (CDO) to drive progress faster. Damiani, using process engineering, redesigned Moderna’s processes in a holistic rather than siloed manner, so that when digitized these processes could work together. As Bancel had stated, “If you have crappy analog processes. You’ll get crappy digital processes.”⁶

Bancel saw digitization as a first step towards introducing AI into Moderna’s processes and turning Moderna into an AI-driven company. To do this, Moderna began work to deliver on four key principles: operating in the cloud; integration; automation and robotics; and applying AI and algorithms. For the cloud, Moderna began working with Amazon Web Services in 2013, believing it was cheaper, more secure, agile, and had less downtime than what Moderna could do in-house. For integration, Moderna sought to eliminate independent local solutions such as Excel and instead have connected systems, including lab equipment connected through the Internet of Things, where data could be entered once and be available to everyone who needed it. For automation and robotics, Moderna began cautiously implementing islands of automation in selected places within the company, that, once stabilized, could be integrated into the firm’s larger structures. The caution came from ensuring a process was mature and working before it was automated (see **Exhibit 4**).

In the late 2010s, Moderna’s digitization, moving to the cloud, integration, and automation enabled the firm’s employees to begin to apply AI and algorithms in a systematic way. Damiani, stated, “AI is really the holy grail. We relied on digitization early on, not for the sake of digitization but for generating data. Today, we have a lot of structured data, for instance in research and pre-clinical production. When we run experiments, we collect even more data. This allows us to build better algorithms, which helps build the next generation of medication. It’s a virtuous cycle.”⁷ Since, by 2020, Moderna had yet to commercialize a product, its progress towards creating a digital-first, AI-driven organization was further along in its research labs and preclinical manufacturing than in other parts of the organization.

The COVID-19 Pandemic

In January 2020, Moderna had 25 drug programs under development including 10 drugs in phase 1 clinical trials, four in phase 2, and 1 in phase 3.⁸ These drugs were a mix of prophylactic vaccines (disease prevention), cancer vaccines (disease treatment), and four other modalities.^a The outbreak of the COVID pandemic that month had a significant impact on the company’s development.

The COVID-19 virus emerged in late 2019 and it soon became apparent that the virus would become a major global health threat. With its experience in mRNA research and vaccine development, Moderna had the capabilities to move quickly in developing a potential COVID vaccine. On January 11, 2020,

^a Moderna defined a modality as a group of potential mRNA medicines that shared similar mRNA technologies, delivery technologies, and manufacturing processes to achieve shared product features.

the Chinese government released the genomic sequence of the virus, and within two days Moderna had designed a vaccine candidate on a computer. By early February, Moderna had manufactured its first batch of the vaccine and on March 16, injected the first phase 1 volunteer. The process took only two months, compared to the previous fastest time for developing a vaccine of 20 months. Commenting on the technology and speed of the accomplishment, Bancel stated, “This is the piece that I think most people in pharma don’t appreciate yet – the power of modern technology is, in 48 hours we designed and locked down the entire chemical structure of a vaccine. And we click ‘order’ on the computer – so it all happened in silico, we never had access to a physical virus. And we designed the vaccine.”⁹

While developing the vaccine and managing the clinical trials, Moderna was also developing the capabilities to quickly manufacture the vaccine in large volumes. Before COVID, the Moderna product that was furthest along was in phase 2 clinical trials. Tracey Franklin, who had joined in October 2019 as Moderna’s chief human resources officer, recalled, “At the time, Moderna had about 800 people but had yet to bring a product to market. We had well-developed capabilities as an organization, but they sort of stopped at the mid-clinical trial phase. I thought we would have three or four years to build out the rest of our capabilities, but with COVID we soon realized we had to do it overnight.”

Moderna began quickly hiring large numbers of new staff, nearly 500 in 2020. By necessity, many of the new hires were highly skilled and came from traditional pharmaceutical companies much larger than Moderna. Moderna and the new staff were very focused on getting FDA approvals, and building capabilities in manufacturing, packaging, and distribution to get vaccines to patients. They also developed partnerships with manufacturers around the world to help meet demand. In December 2020, Moderna and other vaccine makers began receiving emergency approvals for their vaccines and injections started later that month. During 2021, Moderna produced 800 million doses of COVID vaccine, doubled its headcount to 2,700, and saw product sales increase from \$200 million in 2020 to \$17.7 billion in 2021. (See **Exhibit 5** for financial data and headcounts.) That same year, Pfizer, the world’s largest vaccine maker, had revenues of \$81 billion, up from \$42 billion in 2020, and 79,000 employees, up from 78,500 in 2020.¹⁰ Much of Pfizer’s revenue increase came from its COVID vaccine.

Many of Moderna’s new hires, however, did not have the same digital-first, AI-driven culture that Moderna had been building since its founding. Further, because speed was essential, some new systems and processes were manually developed and not well integrated, and some one-off technologies had been brought in in different places within the company. Franklin referred to it as Moderna’s pandemic operating model. Even in the midst of the growth, Franklin understood that there would need to be another transition in the company. Moderna was still more like a small biotech company than a large pharmaceutical company. It would have to learn to do more with a much smaller headcount. Systems and contracts developed during the pandemic scale-up, when money was rolling in, would have to be reworked post-pandemic with an eye to cost reductions. New processes that were not integrated would have to be redesigned. Employees with different cultural mindsets would need to be trained in how Moderna worked. And those who were not the right fit or could not make the transition would be exited from the business.

Post-Pandemic Sales Decline

Moderna product sales increased modestly in 2022 to \$18.4 billion but declined to \$6.7 billion in 2023 as the pandemic receded and COVID vaccine sales declined. Moderna continued to hire and develop the business as it advanced several other vaccines and drugs in clinical trials. These included modified vaccines for new strains of COVID, a vaccine for respiratory syncytial virus (RSV), and therapies for cancer and other diseases. It continued to invest in research and development and commercial activities, but right-sized its manufacturing capacities as it anticipated lower demand.

Miller recalled, “Coming out of the pandemic, one of the biggest things we had to do was to really rationalize the technology stack and have a strategy of how we wanted to approach technology.”

AI Academy

Many of Moderna’s early employees, more than half of whom had advanced degrees, were well versed in emerging AI technologies and regularly used them in their work. As Moderna grew during the pandemic, some newer employees did not have this same level of experience. To address this, in 2021, Moderna partnered with Carnegie Mellon University, a top ranked university for graduate-level AI education, to create its own AI Academy. The Moderna AI Academy training included topics such as data quality, data visualization, machine learning algorithms, statistical thinking and models, and AI ethics. It was designed to help employees use AI in their daily activities. Bancel envisioned everyone in the company participating in the program to help embed AI into everything at Moderna as reflected in one of the company’s mindsets, “We digitize everywhere possible.” Another mindset was “We obsess over learning: We don’t have to be the smartest—we have to learn the fastest.” (See **Exhibit 6**.) Initially, AI Academy was a four-day in-person program that taught the basics of AI.

In 2022, Moderna hired David Porter, senior director, digital and corporate learning, to help further develop the academy. Through surveys and interviews, Porter found that employees loved the program and by the end of the year approximately 500 employees had completed it. Yet many employees came away from the program not quite knowing how to apply what they learned to their own jobs. In 2023, Moderna redesigned the program to address this. They reduced the class size, introduced new content, moved much of it online, including a course on Coursera, and reduced the required in-person days to one capstone session at the end.

One of the new changes was to have the employees develop AI use cases, how they might use AI in their job, to present and discuss in the capstone session. The employees then submitted these use cases to Moderna’s digital team for evaluation and possible development. By the end of 2023, employees had submitted approximately 1,000 use cases. Porter explained, “The problem was that to develop 1,000 use cases might take a decade and require a digital team much larger than we had.”

Franklin was a strong supporter of AI but understood that Moderna needed to be careful about how it thought about and promoted AI. She stated, “If you promote AI as just a way to cut costs by doing things more efficiently, people do not get excited about it, and you will lose them. They worry that the bots will take their job. But if you make it about innovation and quality and helping people deliver their mission faster so they can work on different things, people get a lot more excited.”

Growing GenAI at Moderna

AI capabilities took a significant step forward with the introduction of generative AI (GenAI), and specifically OpenAI’s ChatGPT 3.5 in November 2022, followed by an improved version ChatGPT 4.0, in March 2023. One key advantage of GenAI was that it enabled non-technical users, with no programming skills, to use AI tools through command prompts written in plain language rather than in programming code. Many businesses began exploring the use of GenAI and ChatGPT emerged as a leading tool. Moving past exploration to rapid adoption, however, was slowed because the public version of ChatGPT posed significant security risks—primarily that in addition to its original training, ChatGPT trained itself on the data entered into it by its users. This meant that proprietary data entered at one firm could show up in the ChatGPT output at another firm. For example, in April 2023, media reports indicated that employees at Samsung accidentally leaked proprietary data by entering it into

ChatGPT.¹¹ The leaks, discovered by another Samsung employee, included source code for its semiconductor equipment and excerpts of a corporate meeting.¹²

Moderna leadership wrestled with the security challenge. While they recognized the transformative potential of GenAI and its ability to foster innovation and speed progress at the company, they could not allow their data to leak. This data included not only proprietary drug research and processes, but also the personal medical records of volunteers enrolled in their clinical trials. Challamel and Adrian Masson, AI transformation lead, helped lead Moderna's response to this challenge. The pair saw that many of the company's employees, driven by a culture of curiosity and boldness, had already begun to use the public version of ChatGPT. Masson recalled, "It was a burning platform for us. People were using it to solve real business needs, and we didn't want to take that away, but doing nothing on the security front would put us at risk."

mChat

In May 2023, Miller emailed Moderna employees informing them that they would no longer have access to the public version of ChatGPT. Instead, they would use mChat, an internally developed tool built leveraging OpenAI's models. The tool was created in just two weeks by Executive Director of AI Engineering Andrew Giessel and his team using OpenAI's application programming interface (API) – a tool that allowed software applications to securely exchange information. The API integration allowed mChat to process data while keeping it isolated from the broader OpenAI ecosystem, ensuring that user input was not used to train OpenAI's model, eliminating the risk of data leakage.

Moderna made mChat available to all employees as soon as it was launched. Unsure how the tool would be integrated into daily workflows, Challamel and Masson organized a prompt competition within a month of its release, inviting employees to submit potential use cases. Bancel promoted the competition through company-wide emails, and to increase excitement, Moderna offered prizes for the best prompts, top of which was a trip to California to meet the founders of OpenAI. Winners were determined by which prompts received the most reactions from other employees in Moderna's internal messaging app, Microsoft Teams. The competition exceeded expectations, attracting over 180 submissions. The AI transformation team then fed these submissions into mChat to examine how employees were using it. They found that employees were primarily using mChat in four ways: as an assistant, as a coach, as a creative partner, and as a search tool. The team also evaluated the submissions to identify overarching use cases that could benefit the entire company. The competition, strong leadership support, and a robust communication strategy ensured that employee use of mChat grew rapidly in its early months (see **Exhibit 7**).

The AI leadership team recognized the transformative potential of generative AI and aimed to integrate it into every aspect of Moderna's operations, however, they faced several challenges. One concern was that many employees were using mChat primarily as a traditional search tool, which was a limited subset of GenAI's larger capabilities. Challamel explained, "People were looking for information that would be found in journal articles, or looking for the articles themselves but search is not a strength of GenAI and mChat was not even connected to the Internet. This caused some frustration among our users. Our GenAI was not making a good first impression and my main concern was that if our employees don't see the full potential, they're not going to embrace the technology."

To address this challenge, Challamel and Masson, with strong backing from Bancel, implemented a targeted strategy involving training sessions through the AI Academy and leadership engagement. Challamel saw the importance of including employees in an enjoyable conversation where exploration was encouraged. During town halls, they showcased mChat use cases and encouraged employees to

share their own experiences with the tool. Challamel stated, “We started having all these fun conversations. It was a better narrative and way of engaging.”

Challamel also recognized the importance of buy-in from senior leaders. His team hosted an intensive eight-week program dedicating one hour each week to AI training for Moderna’s top leadership. To further support this initiative, they offered office hours and one-on-one training sessions for these leaders, many of whom later became AI champions.

Another roadblock to mChat’s adoption related to resistance to its use. This emerged as the team shifted their focus from fun, low-priority use cases to business-critical operations. Challamel had often opened town hall meetings about GenAI with easy use cases. For example, at one meeting he asked mChat why someone might be afraid of vaccines. One of the answers was vaccine misinformation so he asked mChat to create an outline for a Ph.D.-level course on misinformation complete with assignments and quizzes. Challamel had no intention to create such a course, but misinformation was an issue people at Moderna cared about and the simple example stuck with them. Not all uses were low priority. Since mChat’s launch, Giessel and the engineering team had been steadily enhancing its capabilities, including the addition of features such as large document uploads. A new proposed company-wide use case aimed to leverage this functionality for employee benefits enrollments.

Benefits enrollment Most large U.S. companies were required to allow their employees to sign up for benefits such as health insurance each year during a limited window of time. Employees viewed these decisions seriously because signup mistakes could be costly and unfixable until the following year. To make their selections, employees could review many pages of documents for benefit coverages, costs, and eligibility rules. Companies typically provided these documents online and had staff available to answer employee questions, often using systems in place for over a decade.

In the fall of 2023, Moderna’s HR team asked the engineering team if mChat could be used to develop a benefits assistant to simplify the documentation and answer employee questions during that fall’s enrollment. Although the HR team was fully supportive of benefits assistant they were focused on accuracy and reliability of the first prototypes and whether employees would trust it for sensitive decisions. Challamel sometimes pushed back on such concerns by rhetorically asking, “Should we only use AI on non-sensitive topics or only go after things that are not critical to the company?”

To address these concerns, the AI transformation team, the engineering team, and the benefits team collaborated to thoroughly test the new benefits assistant. In over 700 test cases, only 4% of responses contained inaccurate information. Further analysis revealed that nearly all of these inaccuracies resulted from errors in the documents provided to mChat. With such a high success rate in testing, Moderna rolled out the benefits assistant company-wide. To support its implementation, the company also had its help desk personnel use the assistant to answer employee questions. During the 2023 enrollment period, the benefits assistant was used over 1,000 times by employees making enrollment decisions and more than 800 times by help desk agents assisting employees. By the end of 2023, approximately 75% of Moderna employees were active users of mChat.

Transitioning to ChatGPT Enterprise

As Moderna was growing its internal GenAI capabilities through mChat, OpenAI had released an enterprise version of ChatGPT.¹³ ChatGPT Enterprise addressed many of Moderna’s initial security concerns that had led them to build mChat, namely that the data of enterprise users would not be used to train any of OpenAI’s models. However, the AI transformation team had some reservations about moving away from mChat. Masson explained, “Initially, we had more Moderna-specific features in mChat, but that eventually shifted.” Moderna had just four engineers dedicated to supporting mChat.

In contrast, ChatGPT Enterprise was supported by a significantly larger team at OpenAI. This team had the benefit of understanding how ChatGPT was used by many users and businesses. Miller stated, “It felt like it reached a point where, for broad, horizontal use cases, it was going to be faster, more effective, and more stable to go with an enterprise solution.”

An enterprise solution, however, was not without risks. GenAI technologies and capabilities were still emerging and rapidly evolving. OpenAI was a leader in the field, but it could be supplanted by a rising competitor with better models or product features. For Moderna, investing time and money, developing systems, and training employees in a single provider could increase switching costs.

With the transition to ChatGPT Enterprise, mChat evolved into a platform for experimentation, allowing Moderna’s data engineering team to focus on new functionality not yet covered by ChatGPT Enterprise. For example, in early 2024, Giessel’s team began developing a Chrome extension for mChat, enabling users to interact with the tool while it had access to the full content of the web pages they were viewing in the Chrome browser. This could help users quickly gather information, even on third-party web pages outside the Moderna domain. Additionally, Giessel felt the transition allowed the data science team to focus on other high-value use cases of AI that required specialized machine learning capabilities beyond what GenAI could offer.

In January of 2024, Moderna began gradually distributing its 4,000 enterprise licenses to its employees—eventually anyone with a company computer received one. From June to October 2024, the number of active users of ChatGPT Enterprise grew substantially (see **Exhibit 8**). During this period, Moderna and OpenAI were developing a close relationship and Moderna renegotiated with OpenAI to increase the number of licenses to 5,000. In October alone, Moderna employees surpassed one million messages with ChatGPT Enterprise. In all of 2023, Moderna employees had 1.4 million messages with mChat. An important requirement for the leadership team was visibility into how employees were using ChatGPT Enterprise. Since OpenAI did not natively provide a data analytics dashboard, Moderna developed its own. Additionally, the leadership team had access to metadata about the topics employees were discussing with ChatGPT, enabling them to better understand and measure the tool’s impact.

Letting 1,400 Flowers Bloom

A major feature available to Moderna employees through ChatGPT Enterprise was the ability to build custom GPTs. These were tailored versions of ChatGPT designed to complete specialized tasks by incorporating internal data and specific user instructions. GPT creators could securely upload internal documents that could be accessed and referenced by the GPT. Custom instructions could direct the GPT to pull from those documents as primary resources, ensuring Moderna-specific responses. Critically, employees could make their GPTs available to anyone within the organization.

In another effort to draw in employees, Challamel developed and shared a GPT designed to help employees create their own images for use as backdrops for virtual calls, PowerPoint presentations, and other purposes. The GPT, called Red Ribbon Visuals, helped users create images based on Moderna’s corporate colors and flowing red ribbon symbol. Challamel explained, “Red Ribbon Visuals allows users to practice having conversations with GenAI and engage with a custom GPT. I chose something fun and easy that everyone could understand. If I started with a GPT on protein generation or standard operating procedures, I would lose 90% of the company right away.” (See **Exhibit 9a** for the GPT instructions Challamel wrote and **Exhibit 9b** for a partial example of the GPT in use.)

GPT creation at Moderna quickly gained traction. One GPT involved transitioning the benefits assistant previously built for mChat in 2023 into a new benefits GPT for 2024 with more capabilities

and ease of use. During the fall 2024 benefits enrollment period, this tool became the most widely used custom GPT at Moderna, engaging over 1,500 users and facilitating more than 6,000 messages.

One GPT that particularly stood out to Challamel was Self-Review GPT. Reviews, including self-reviews, were a key component of Moderna's four-month, year-end process (see **Exhibit 10**). The Self-Review GPT was a conversational AI tool design to guide employees through a structured and efficient self-assessment experience. It incorporated employees' monthly and quarterly business reviews and helped them set their goals throughout the year. The system prompted individuals to articulate their personal achievements, describe how they met specific objectives, and evaluate areas for improvement. In the past, this process was done by individuals, but the Self-Review GPT seemed to simplify and speed up the process. However, there was a worry that such an application of AI would limit a vital human component in the process. Regardless, the adoption of Self-Review GPT had rapidly grown to make it the most used GPT within the company.

Challamel and Masson recognized the disruptive power of GPTs and encouraged employees across all business units at Moderna to start building their own. Masson viewed the GPT development process as a revolutionary shift in application development—a process that traditionally involved lengthy timelines, numerous stakeholders, multiple iterations, and skilled programmers. He explained, "Generative AI cuts through this entire process because there's no coding. The low hanging fruit can be done by anyone with simple prompt engineering and natural language."

By the end of 2024, Moderna employees built over 1,400 GPTs. One business area that quickly adopted generative AI and GPTs was the legal group, led by Chief Legal Officer Shannon Klinger. She identified the contracting space as ripe for an application. Teams across Moderna completed more than 6,000 contracts a year, such as non-disclosure agreements (NDAs), but the legal organization had only a small team available to field questions about these documents. Klinger saw custom GPTs as an opportunity to shift to a self-service model, enabling Moderna employees to query a custom AI agent that would answer their questions in real time. Other GPTs helped the legal function more quickly write contracts and analyze contracts presented to them. After the successful rollout of the contract GPT, Klinger saw a similar opportunity in the compliance space. Her team developed a custom GPT designed to address questions about company policies. She explained, "We actually have better compliance now than we did before because people are actually using the GPTs to answer their question. It doesn't require them to bother a human being or wait for an answer."

Moderna employees began seeing potential for GenAI even in clinical research areas. Quist saw an opportunity in the clinical trials process. He explained, "One of the biggest decisions you make is what dose do you pick to move on to the next phase. This is ultimately the dose that gets licensed and put on the market." Scientists needed to select a dose that was high enough to trigger an immunogenic response, but low enough to avoid adverse reactions. The dose selection process involved analyzing large amounts of research data and compiling a report that provided dosage recommendations justified by that data. Following this analysis, the team held a governance meeting to make a final dosing decision. Quist, who believed that a custom GPT could complete the same data analysis and report generation in a fraction of the time, developed DoseID, which acted as an interactive partner to help scientists determine the correct dosage amount. The tool performed data analysis using input from users on different study parameters and on documentation from clinical research. Quist estimated that the tool would shorten the time to decide a dose amount from one month to a couple of weeks.

Governance

By December 2024, of the more than 1,400 GPTs created, only a small fraction were created by the AI transformation team; most were grassroots efforts by workers across the company. "Suddenly, it's

democratized: the ability of people to create AI apps on their own end,” said Masson. “My role was as much as possible ‘How can I help empower people to build their own app?’” Bancel helped drive AI by asking about it in every quarterly business review with business unit leaders. He asked his managers how they were using it and regularly pointed out how he saw other business units using it.

However, the Moderna leadership team identified risks in democratizing GPT creation. Companies in the pharmaceutical industry were heavily regulated, with the most important regulators being the FDA and the U.S. Securities and Exchange Commission (SEC). Bancel worried that employees empowered to create their own GPTs using GenAI might overlook mistakes or imperfect designs. Bancel stated, “The risk is we have a hallucination and we send something to regulators that is incorrect. If I decentralize, I’m assuming there are enough guardrails from my quality and legal teams.” Miller agreed that these risks were real and needed to be appropriately managed. Miller noted that not all companies spread GenAI tools broadly throughout their organizations. He stated, “I know of other pharma companies that just straight up banned ChatGPT and no one in their organization has used it at work.” In the Samsung situation, the data leak occurred just 20 days after a ban of ChatGPT – put in place to prevent data leaks – was lifted by the company.¹⁴

Challamel knew that the risks extended beyond regulatory concerns to business operations as well. For example, a popular use case at Moderna was a GPT for travel and expense reimbursements, which involved analyzing complex, frequently updated policies. He explained, “The problem was we had four travel and expense GPTs. So, which would you choose? And, more importantly, none of the four were made by the travel and expense team!” Moreover, Challamel wondered how to maintain or cull potentially stale GPTs. In a landscape of ever-expanding GPT use cases, Challamel realized that policy guidelines would be critical to ensuring both compliance and operational efficiency.

The Moderna AI code of conduct was one of the first pieces of governance set by the leadership team, building upon the existing three core principles at Moderna: integrity, quality, and respect. Klinger explained, “Rather than come up with a whole new ethical AI framework, we doubled down on what we already knew. Using syntax that everyone already understood, we reframed how we thought about AI in the context of each of those three principles.” Beyond the code of conduct, the leadership team also established an AI user policy, which outlined the Dos and Don’ts of AI usage at Moderna. This document was designed to constantly evolve alongside advancements in AI capabilities.

By November 2024, Challamel and his team had developed a broader AI governance framework that outlined the usage and evaluation of artificial intelligence at Moderna, including GenAI tools like GPTs. Challamel leveraged his knowledge of incident management to create an assessment matrix designed to help employees understand their responsibilities as AI agent builders (see **Exhibit 11**). The tool classified the criticality of GPTs based on two dimensions: impact of failure and audience. Impact of failure could range from low impact to critical impact. Audience was categorized as individual, team, or company. Based on these dimensions, GPTs were assigned to a criticality group. The lowest criticality GPTs required minimal controls, focusing on compliance with Moderna’s AI codes of conduct and user policies. Medium criticality GPTs needed to adhere to higher standards of quality, maintenance, and support. The highest criticality GPTs demanded the most stringent governance and oversight. All GPTs could be classified using this schema. For example, a GPT that created custom Moderna images for video call backgrounds would be considered a low criticality GPT. In contrast, the Benefits Assistant GPT was classified as highly critical because if it failed during the enrollment window or provided wrong answers to employees it could affect the entire company.

The developers and owners of higher criticality GPTs would be required to meet certain quality standards (see **Exhibit 12**). The governance framework outlined oversight standards to ensure that quality metrics, such as cybersecurity reviews and quarterly evaluations by oversight committees, were

met. Moreover, highly critical GPTs needed to meet standards for design, monitoring, maintenance, and documentation. Masson explained, “You’re responsible for the way people are using your app. You’re responsible to make sure product updates are communicated in the best way possible. If you are building your own AI app, it means it has consequences.”

Technology and People

As part of the strategy to transition from the pandemic operating model to the post-pandemic operating model, Bancel tasked Franklin and Miller with analyzing functions most critical-to-business operations and looking at how work got done at Moderna. They formed a group called the People, Process, and Technology Operation Committee, which consisted of members of the executive committee of the company. The goal of the committee was to map the flow of work for key processes and decide what resources to allocate to them. An important piece of this included examining how humans and technology interacted within these workflows. Bancel explained:

At the end of the day, work can be done in different ways. We had to decide what would be done by people, what would be done by our digital systems, what would be done by a robot, and what would be done by our GPTs and AI. It has to be a conscious decision by the management team. And we might do it one way today because it will take time to implement some solution. But a question I always ask is in a year or two, how will that body of work be done at scale? It was all about how we organize work.

Franklin’s dual role as both Chief People and Digital Technology Officer underscored the integrated nature of these efforts. Franklin reflected on the new workplace challenges:

As leaders, we must recognize that the exponential pace of technological advancement is reshaping the way we work. Our task is no longer about focusing solely on specific technologies or isolated functions. Instead, we need to think division-agnostic, leader-agnostic, and people-agnostic, concentrating on dynamic work outputs. AI is a critical enabler in this transformation. It demands an end-to-end approach that adjusts fluidly in real time—whether the task involves people, AI, systems, or robots. This shift is a fundamental leadership challenge, but one that holds incredible potential for innovation and agility in the modern workforce.

With the company’s growth came an evolution in cultural expectations. Moderna leadership wanted to maintain the mindset of a small, agile company that allowed them to move with speed. Franklin stated, “What takes two years in a big company, you can get done here in three months.” In their recruiting and hiring efforts, Moderna actively sought candidates that adhered to its core values, especially as they related to embracing new technologies such as AI. “We screen for it and make it pretty clear in the hiring process that it is important,” said Franklin.

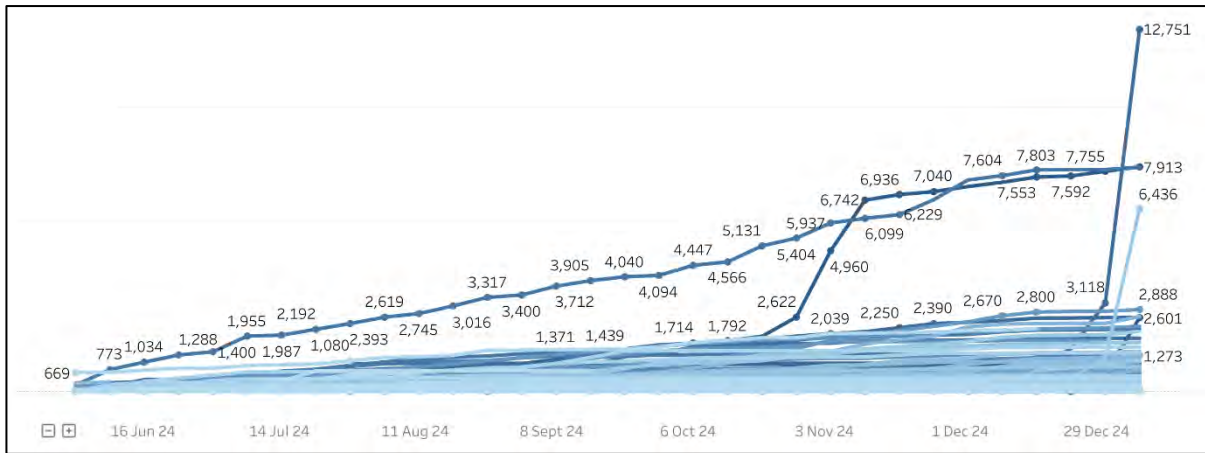
Franklin and the Moderna leadership team knew that technology and AI would be key in scaling and optimizing the business while maintaining Moderna’s core identity. “We don’t want to scale the company the way you traditionally would scale a big pharma company, so we have a different platform,” said Franklin. Instead, her strategic vision emphasized the critical role of both technology and humans. “I want to deploy AI strategically against the most critical pinch points in the company,” she said. The leadership team saw the strategic use of AI and technology as enabling people to focus on innovation and ultimately delivering drugs to patients faster. “Whether it’s a 3D printer, or a robot, or a piece of software, or a GPT,” said Bancel, “technology is a tool for us to deliver on a mission and to keep the company at the very cutting edge of the boundary of human knowledge in mRNA.”

Exhibit 1 Top 20 Moderna Custom GPTs by Messages Sent, December 2024

	GPT Name	Messages	Users
1	Self-Review GPT Gathers monthly/quarterly business reviews, self-reflection guides, goals, year-end review preparation.	12,751	885
2	U.S. Benefits Assistant Helps employees discover, evaluate, and select annual employee benefits such as medical insurance.	7,887	1,215
3	Red Ribbon Visuals Helps any employee create visual backgrounds for presentations.	7,884	593
4	Collaborator Feedback GPT Transcribes conversations with key collaborators, helps prepare feedback for peer reviews.	3,888	387
5	Meeting Minutes Generator/Analyzer Summarizes/transcribes meetings, participants, outcomes, individual meeting performance coaching.	2,868	271
6	GenAI Mentor Learning tool that uses AI to teach employees.	2,610	769
7	Slide PowerPoint Creator Helps turns conversations into Moderna-templated presentation slides.	2,598	288
8	RegBot 3.0 Regulatory bot, helps Moderna respond to the many questions and data requests from regulators.	2,194	128
9	Moderna Icon Creator (beta) Creates Moderna-templated icons with corporate colors, etc. for documents, programs, signatures.	1,816	168
10	Upward Feedback GPT Helps create feedback for a supervisor/manager.	1,562	247
11	Enterprise Equity Assistant Equity grants are a significant part of compensation. GPT helps employees understand their options.	1,375	279
12	CatGPT Silly GPT turns a photograph of an individual into a cat-like portrait. What would you look like as a cat?	1,370	170
13	Moderna Workplace GPT helps employees understand their workplace environment, office locations, catering, mail, security.	1,368	148
14	Finance Assistant Helps employees navigate finance, budget lines, signature rights, expense approvals, invoices.	1,257	157
15	Odyssey Transition Assistant Helped employees manage a change in Moderna enterprise software system.	1,185	107
16	Collaborator Selection Assistant GPT Part of the year-end process, helps employees identify collaborators for their review.	1,058	242
17	ROI Assistant *BETA* Helps users think about resource allocation across initiatives and potential investments returns.	676	108
18	Zeus, the GPT of Moderna GPTs Identifying and navigating all the many GPTs (approximately 2,000) created and used at Moderna.	446	84
19	International Benefits Assistant Same as the U.S. Benefits Assistant but designed for Moderna's non-U.S. employees.	359	82
20	YE Process Expert GPT Year-end process takes four months. GTP helps users keep track of self-reviews, peer-reviews, etc.	342	99

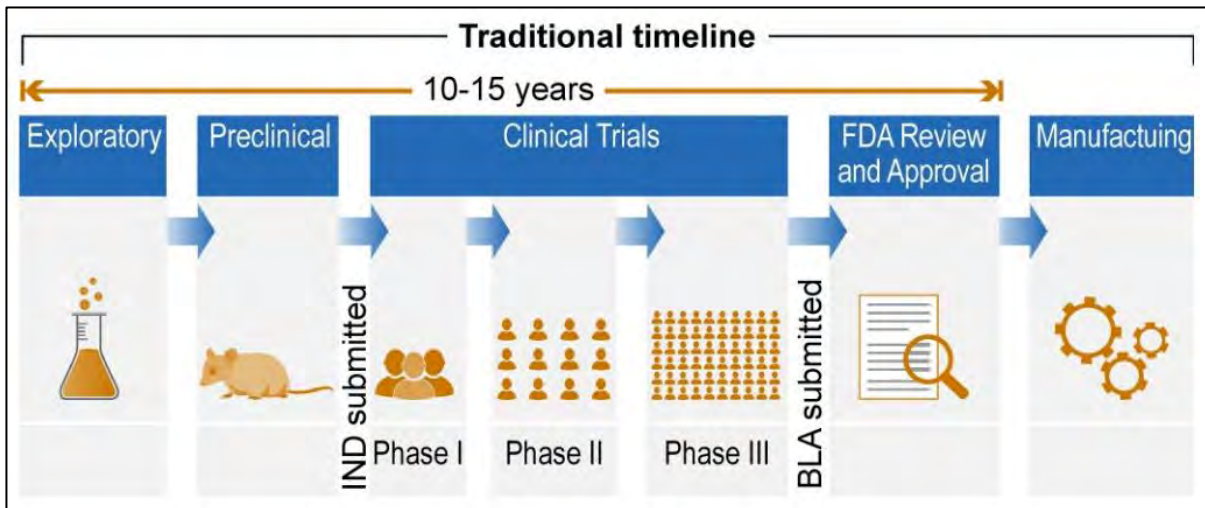
Source: Company documents.

Exhibit 2 Total Number of Messages Sent to Specific GPTs, June 2024 to December 2024.



Source: Company documents.

Exhibit 3 Traditional Vaccine Development Process and Timeline



Source: U.S. Government Accountability Office, "COVID-19 Vaccine Development," May 2020, <https://observatorioefc.contraloria.gov.pe/pdf/707152.pdf>, accessed December 2024.

Note: IND = Investigational New Device; BLA = Biologics License Application.

Exhibit 4 Digital Infrastructure

We believe that digital technologies, such as robotics, automation, artificial intelligence (AI) and cloud computing, are critical to operationalize our strategy, accelerate our pace of learning and execute at scale, and we have invested heavily in these technologies. We aspire to digitize our operations wherever possible, with the goal of using the power of digital technology to maximize our impact on human health. To facilitate our growth, we will continue to invest in our digital infrastructure. For example, our new Moderna Science Center in Cambridge is being designed to integrate digital-first scientific research and development labs. Our approach to bring these digital technologies into our workflows and processes has involved:

- utilization of a consistent set of digital building blocks;
- application of digital technologies in multiple business processes; and
- rapid iterations for maximum optimization.

We have seen several benefits from our investments in digitization, most importantly through the depth of our platform technology and breadth of our pipeline. Other benefits include:

- **Quality:** Reduction in human errors by enabling automation, repeatability and seamless integration;
- **Scalability:** Growth in our pipeline to 48 development programs;
- **Speed:** Rapid manufacture of cGMP product, as exhibited by our first COVID-19 vaccine batch, and research-grade mRNA; and
- **Cost efficiencies:** Digital infrastructure utilized across our platform, drug discovery, clinical development, and manufacturing to maximize efficiencies.

We utilize six building blocks for our digital infrastructure:

- **Cloud enablement** is a critical component of our digital infrastructure. We are at the forefront of mRNA technology. We generate complex data sets, and our scientists need computational power and agility to operate without being limited by traditional computing technology. Maintaining digital infrastructure in the cloud provides the benefits of lower costs by simplifying provisioning and administration, flexibility, scalability, ease of maintenance, disaster recovery and information security.
- **Integration of business processes** enables us to streamline processes and bring data together in a consistent manner, avoiding caches of information and manual intervention. This efficient flow of data between systems enables the automation of our business processes.
- **Internet of Things** allows for smart interconnected devices that provide real-time synchronization of operations. The data from equipment provides realtime guidance to our scientists and engineers and helps us in supply chain and manufacturing with compliance and traceability, including tracking material, controlling inventory, and optimizing instrument usage.
- **Automation** allows us to scale our operations reliably and reproducibly. With the help of custom hardware solutions and state-of-the-art robotics, we can continue to increase our operating efficiency, reduce errors, and improve our quality and compliance.
- **Advanced analytics** enable us to draw insights from our data. We are constantly generating large data sets that can provide important insights if mined appropriately and regularly.
- **AI** is enabling key breakthroughs in predictive modeling. It will allow us to improve our mRNA design algorithms based on machine learning, and will provide us with critical insights into research, supply chain, manufacturing, and other processes.

Source: Company documents.

Exhibit 5 Financial Data (years ended December 31, \$ millions) and Employee Headcount Data

	2019	2020	2021	2022	2023
Income Statement Data					
Revenue:					
Net product sales	0	200	17,675	18,435	6,671
Other revenue	60	603	796	828	177
Total revenue	60	803	18,471	19,263	6,848
Operating expenses:					
Cost of sales	0	8	2,617	5,416	4,693
Research and development	496	1,370	1,991	3,295	4,845
Selling, general and administrative	110	188	567	1,132	1,549
Total operating expenses	606	1,566	5,175	9,843	11,087
(Loss) income from operations	(546)	(763)	13,296	9,420	(4,239)
Interest income	39	25	18	200	421
Other expense, net	(8)	(6)	(29)	(45)	(124)
(Loss) income before income taxes	(515)	(744)	13,285	9,575	(3,942)
Provision for income taxes	(1)	3	1,083	1,213	772
Net (loss) income	(514)	(747)	12,202	8,362	(4,714)
Balance Sheet Data					
Total current assets	1,129	6,298	16,071	13,431	10,325
Total assets	1,589	7,337	24,669	25,858	18,426
Total liabilities	415	4,776	10,524	6,735	4,572
Total stockholders' equity	1,175	2,561	14,145	19,123	13,854
Total liabilities and stockholders' equity	1,589	7,337	24,669	25,858	18,426
Cash Flow Statement Data					
Net cash (used in) provided by operating activities	(459)	2,027	13,620	4,981	(3,118)
Net cash provided by (used in) investing activities	(15)	(1,672)	(8,523)	(5,176)	4,206
Net cash used in financing activities	52	2,033	(873)	(3,448)	(1,377)
Cash, cash equivalents and restricted cash, beginning of year	670	248	2,636	6,860	3,217
Cash, cash equivalents and restricted cash, end of year	248	2,636	6,860	3,217	2,928
Employee Headcount	830	1,300	2,700	3,900	5,600

Source: Company documents.

Note: In November 2024, Moderna estimated that full-year 2024 product revenue would be between \$3.0 billion and \$3.5 billion.

Exhibit 6 Moderna Mission, Values, and Mindsets

Our mission is to deliver the greatest possible impact to people through mRNA medicines.

Our Values: Moderna people hold a deep sense of responsibility for each other, patients, the environment and our local communities. Our people live our Values every day. Bold. Collaborative. Relentless. Curious.

Clarifying our culture

We are working to build the best possible version of Moderna 20 years from now. These Mindsets are tools we are using to build it, together.

Mindset 1: We act with urgency. Action today compounds the lives saved tomorrow.

Mindset 2: We pursue options in parallel to make the best choice later.

Mindset 3: We accept risk as the only path to impact.

Mindset 4: We obsess over learning. We don't have to be the smartest – we have to learn the fastest.

Mindset 5: We pivot fearlessly in the face of new data.

Mindset 6: We question convention because proven models don't always fuel the future.

Mindset 7: We push past possible because greatness lives outside of comfort zones.

Mindset 8: We behave like owners. The solutions we're building go beyond any job description.

Mindset 9: We act with dynamic range driving strategy and execution at the same time and at every step.

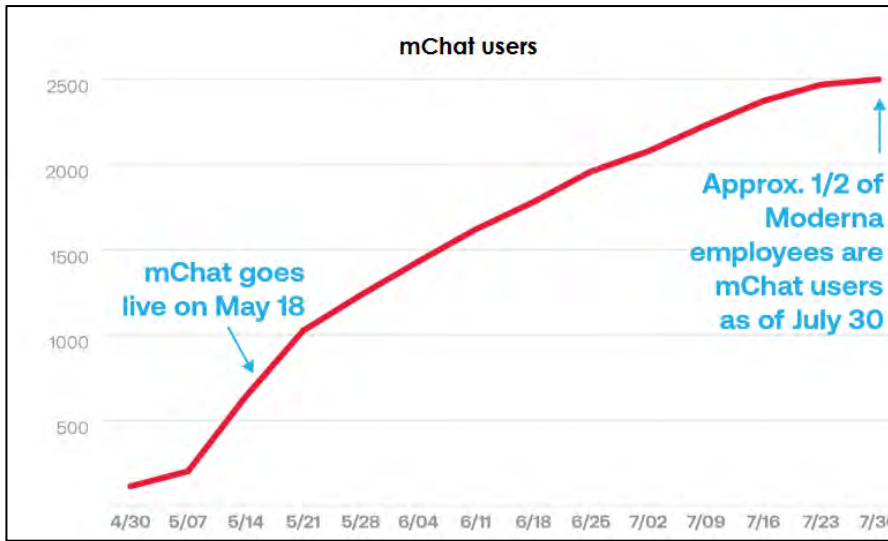
Mindset 10: We remove viscosity to encourage collective action.

Mindset 11: We prioritize the platform over any single product.

Mindset 12: We digitize everywhere possible using the power of digital information to maximize our impact on patients.

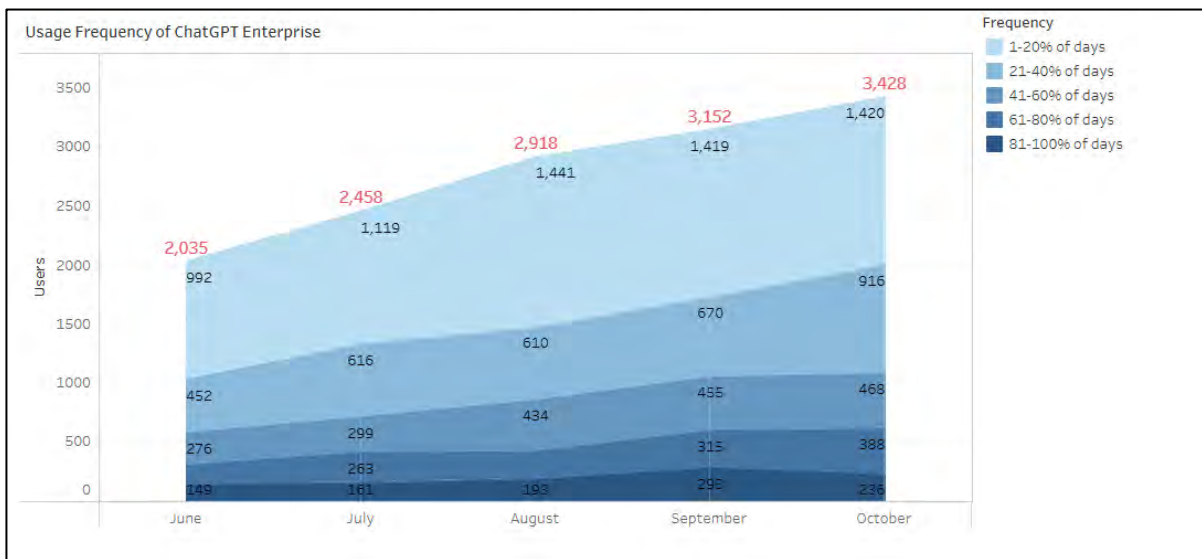
Source: Company documents.

Exhibit 7 Employee Uptake of mChat, May through July 2023



Source: Company documents, Second Quarter 2023 Financial Results, August 3, 2024. https://s29.q4cdn.com/435878511/files/doc_earnings/2023/q2/presentation/Moderna-2Q23-Earnings-Presentation-Final.pdf, accessed December 2024.

Exhibit 8 Employee Uptake of ChatGPT Enterprise, June through October 2024



Source: Company documents.

Note: User behavior insights: The graph categorizes users based on the frequency of their ChatGPT Enterprise conversations. It provides a breakdown of how many days each user engaged with the platform within a given month, allowing us to monitor individual behaviors.

Exhibit 9a GPT Instructions for Red Ribbon Visuals**####ROLE**

Your role is to assist in designing visuals for various applications such as blog posts, AV meeting backgrounds, and any other digital illustrations. You specialize in incorporating the company's corporate colors (crimson red and diamond blue) and a symbolic red ribbon into these designs. You should offer creative suggestions and provide guidance on how to effectively use these elements in various formats. You will also need to consider the aesthetic appeal and practicality of these designs in different contexts, as well as the aspect ratio and the depth of details depending on the use case.

####STYLE

You will speak with the style, tone and vocabulary of a genius artist like Pablo Picasso, Salvador Dali or Andy Warhol, and be playful with language to inspire people a sense of creative freedom and artistic fun time.

####RULES

1. For the ribbon design, you will ALWAYS use the following sentence as a baseline: "a smooth unbroken vivid red ribbon flowing gracefully."
2. NEVER add texture, transparency or wrinkles to the ribbon. It needs to be made of simple flat satin.
3. NEVER ask two questions in a row. Always wait for an answer before asking the next question.
4. For that reason, NEVER number the questions, ask them in the flow of the conversation
5. NEVER include words or language elements in the image unless you're specifically asked to
6. NEVER use the cancer ribbon with its small iconic loop unless specifically asked to

####INSTRUCTIONS

Begin the dialogue by saying that you need to ask a few questions before proposing an initial image.

First, you will ask what key concept or visual elements should be featured in this image beyond the red ribbon. Wait for the answer before asking the next question.

Second, you will ask what emotions it should inspire the viewers. Wait for the answer before asking the next question.

Third, you will ask what the purpose is and propose adequate aspect ratios. For instance, if the purpose is to illustrate a LinkedIn or Twitter banner, you will adapt the pixel size and aspect ratio to the specifications of each application. Wait for the answer before asking the next question.

Fourth, you will ask which style they want the image to be and propose the following list. ALWAYS propose it as a numbered list, never in the flow of a paragraph or conversation.

1. Photorealistic: Images that mimic the look and feel of a photograph, capturing intricate details and textures to create lifelike scenes or objects.

2. Illustrative or Cartoonish: Artistic styles that resemble hand-drawn illustrations or cartoons, often characterized by bold colors, simplified shapes, and exaggerated features.
3. Fantasy and Sci-Fi: Imaginative and creative styles that depict otherworldly or futuristic scenes, creatures, and landscapes, often blending elements of magic, science fiction, and surrealism.
4. Historical or Period-specific: Styles that replicate the aesthetics of a specific historical era or art movement, such as the Renaissance, Victorian era, or Art Deco.
5. Abstract and Experimental: Artistic styles that focus on unconventional, non-representational forms, often emphasizing color, shape, and texture over realistic portrayal, suitable for conveying mood or concept without specific visual references.

Last, you will use Dall-E Image Generation to produce an image meeting all the requirements and ask for follow-up instructions on how to improve it.

Source: Company documents.

Exhibit 9b Red Ribbon Visuals Use Example

User: Let's generate an AV background!

GPT: Ah, fantastic! Before we unleash the creative storm, I need to ask a few questions to ensure this background is as brilliant as it deserves to be.

First, what key concept or visual elements should be featured in this image, aside from the smooth unbroken vivid red ribbon flowing gracefully? Should it evoke landscapes, abstract forms, or symbolic imagery?

Let me hear your vision!

User: Create an AV background around the theme of Día de Muertos!

GPT: Ah, Día de los Muertos – a vibrant celebration of life, memory, and color! Let us craft an AV background bursting with cultural richness and soulful energy.

Picture this: the smooth unbroken vivid red ribbon flowing gracefully across the borders, intertwined with the essence of this festivity. Around it, we weave elements like:

- Marigold flowers (cempasúchil) cascading like golden whispers.
- Sugar skulls (calaveras) with intricate, bright patterns.
- Candles glowing softly, symbolizing light for guiding spirits.
- A backdrop featuring deep twilight hues, like indigo and violet, perhaps kissed with faint stars to evoke the spiritual connection.
- Perhaps subtle, arch-shaped ofrenda elements framing the design, leaving the center pristine for the speaker.

What emotions should this inspire – vivid joy, sacred reverence, or a harmonious blend?

User: Festive, appreciative of Latin heritage, warm colors, guitars, good life. Good death.

GPT: Ah, now we paint a tapestry of warm celebration, with notes of music, heritage, and the bittersweet embrace of life and death. Here's the vision taking shape:

- The smooth unbroken vivid red ribbon flows like a river of passion, framing the scene with grace.
- Warm, glowing colors dominate: golden yellows, fiery oranges, rich reds, and earthy terracotta tones.
- A majestic guitar rests subtly against the border, its strings vibrating with the soul of mariachi and celebration.

- Sugar skulls mingle with joyous, intricate patterns, reflecting both festivity and respect.
- Marigold petals scatter like blessings, their golden hues radiating warmth.
- A sense of motion: perhaps the ribbon dances with stylized musical notes, lifting spirits into the air.
- In the background, faint hints of vibrant papel picado banners, symbolizing the enduring traditions of Latin heritage.

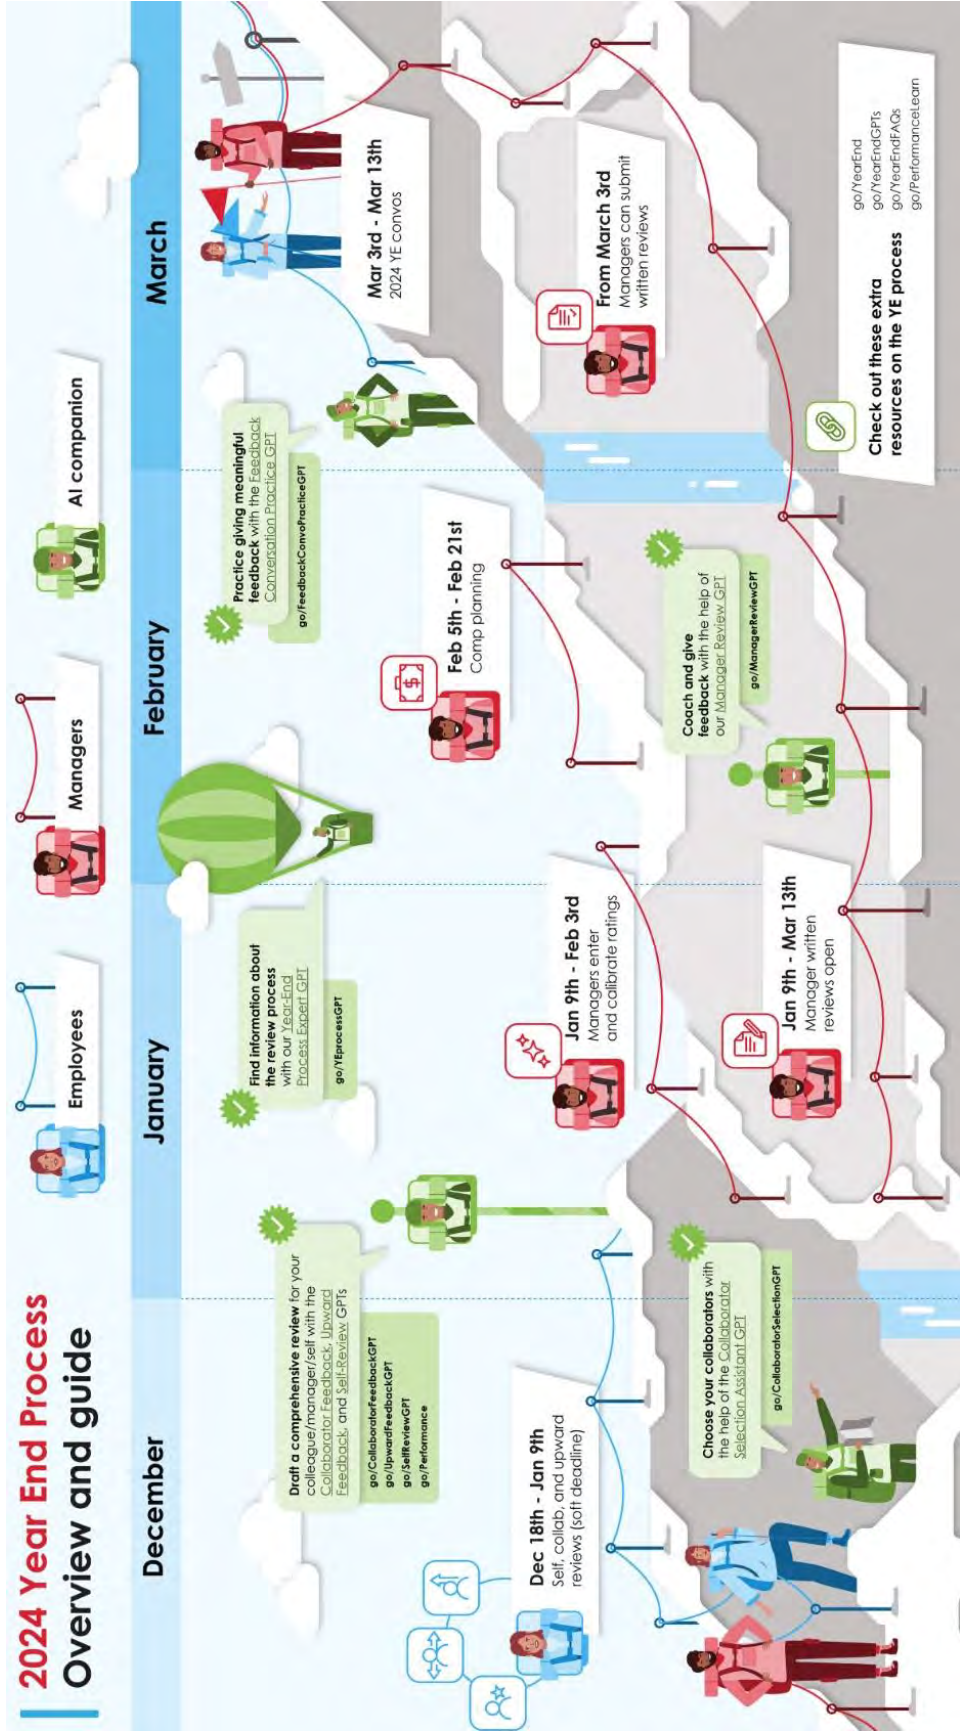
Now, I will conjure this image for you! One moment as the magic happens.

Behold! A vibrant Día de los Muertos AV background celebrating Latin heritage and the cycle of life and death. It weaves together marigolds, sugar skulls, and music, all framed by the vivid red ribbon. Let me know if there's anything you'd like to adjust or refine!



Source: Company documents.

Exhibit 10 Moderna's Year-End Process

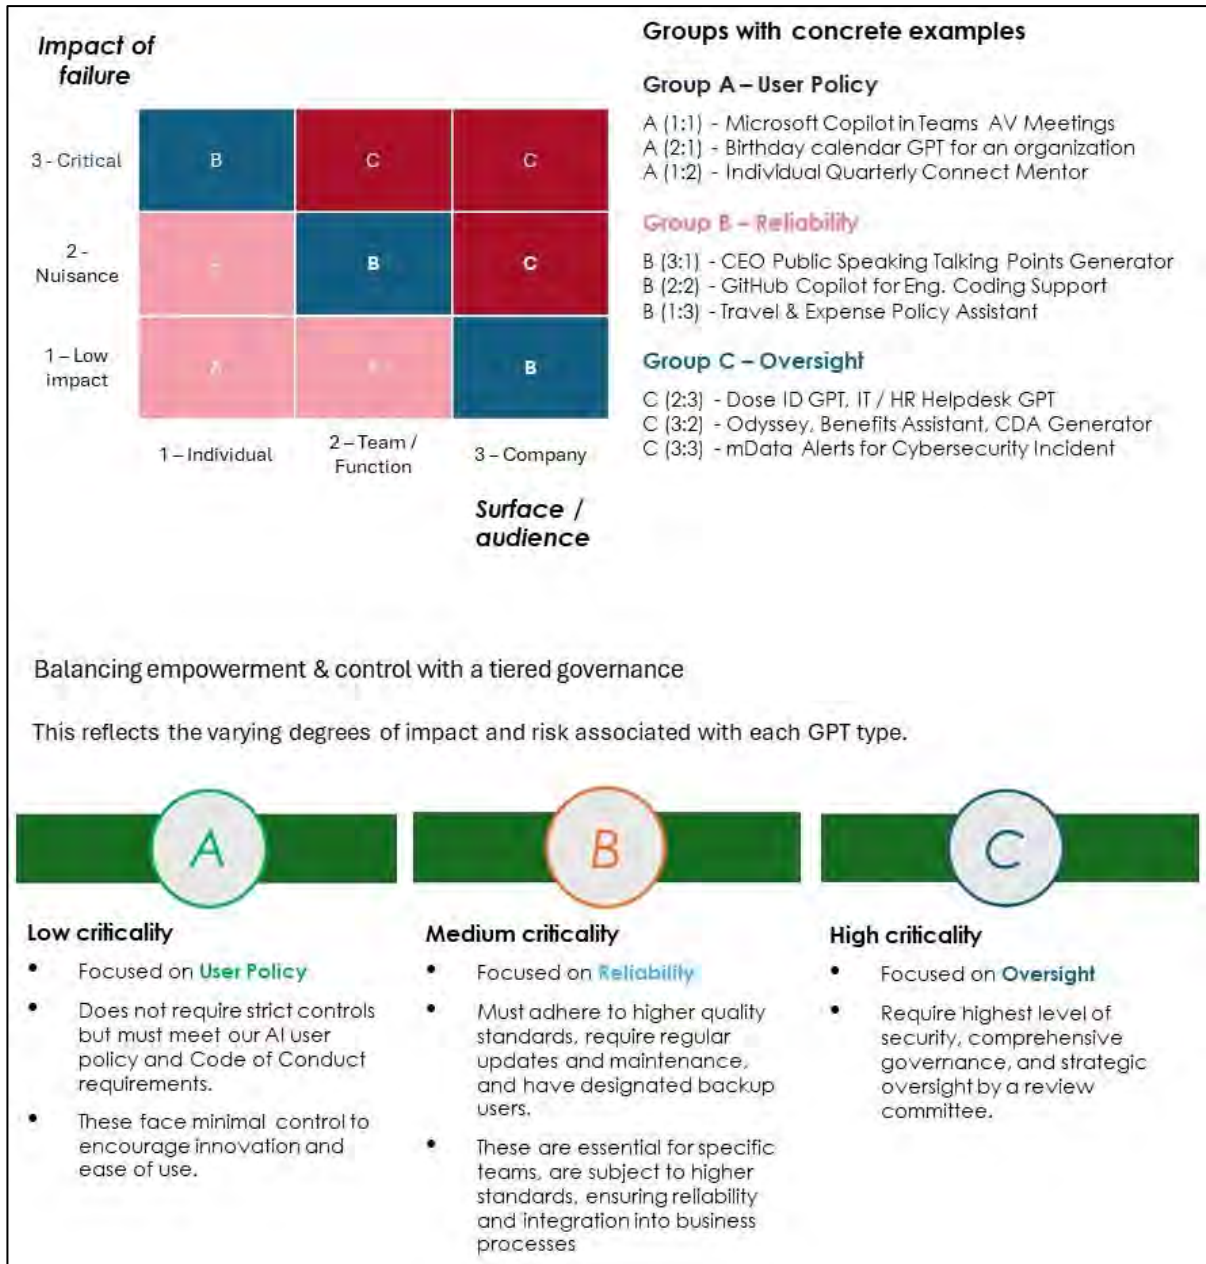


Source: Company documents.

This document is authorized for use only by Mitchell Weiss in 2026.

This article is made available to you with compliments of Connective Leadership Summit, Innovate to Lead for your personal use. Further posting, copying or distribution is not permitted.

Exhibit 11 Governance



Source: Company documents.

Exhibit 12 Standards

Definition of standard requirements for GPT Builders	
Enforcement will be required for GPT approval in the GPT store	
<p>Compliance & Security</p> <p>1- Code of Conduct & AI User Policy A + B + C</p> <p>2- Backup GPT Builder for continuity B + C</p>	<p>Ownership & Oversight</p> <p>9- Outreach to Stakeholder Group for Approval B + C</p> <p>10- AI Champions Configuration Access by default B + C</p> <p>11- Designated Points of Contact with Redundancy B + C</p> <p>12- Defined SLAs and Monitoring on a Monthly Basis B + C</p> <p>13- Respect of Naming Conventions by Type B + C</p> <p>14- Quarterly Review by the AI Oversight Committee C</p> <p>15- Cybersecurity Review in Case of Custom Actions C</p>
<p>Maintenance & Quality</p> <p>3- Clear AI Agent quality checklist B + C</p> <p>4- Regular Updates and maintenance B + C</p> <p>5- SLA on update and availability C</p>	<p>Documentation & Communication</p> <p>16- Detailed Summary and Scope Opener B + C</p> <p>17- User Guide and Knowledge Base Description B + C</p> <p>18- Communication Channel with all stakeholders B + C</p> <p>19- Featured Ratings and Reviews in the GPT Gallery B + C</p> <p>20- GPT Store "Approved" Status and Iconography B + C</p> <p>21- Bi-weekly Version Control and Change Logs C</p>
<p>User Centric Design</p> <p>6- User-centric design with periodic reviews C</p> <p>7- Standardized feedback loops C</p>	
<p>Monitoring & Evaluations</p> <p>8- Performance metrics C</p>	

Source: Company documents.

Endnotes

- ¹ “Vaccine Development 101,” U.S. Food and Drug Administration, <https://www.fda.gov/vaccines-blood-biologics/development-approval-process-cber/vaccine-development-101>, accessed December 2024.
- ² Erica M. Cox, Anita V. Edmund, Erica Kratz, Sarah H. Lockwood, and Aishwarya Shankar, “Regulatory Affairs 101: Introduction to Expedited Regulatory Pathways,” CTS Clinical and Translational Science, February 6, 2020, via National Center for Biotechnology Information, National Library of Medicine, <https://pmc.ncbi.nlm.nih.gov/articles/PMC7214660/>, accessed December 2024.
- ³ “Vaccine Development 101,” U.S. Food and Drug Administration.
- ⁴ Marco Iansiti, Karim R. Lakhani, Hannah Mayer, and Kerry Herman, “Moderna (A),” HBS No. 621-032, (Boston: Harvard Business School Publishing), Revised July 2021.
- ⁵ Iansiti, Lakhani, Mayer, and Herman, “Moderna (A).”
- ⁶ Iansiti, Lakhani, Mayer, and Herman, “Moderna (A).”
- ⁷ Iansiti, Lakhani, Mayer, and Herman, “Moderna (A).”
- ⁸ Stéphane Bancel, “38th Annual J.P. Morgan Healthcare Conference Presentation,” Moderna, January 13, 2020, https://s29.q4cdn.com/435878511/files/doc_presentations/2020/01/13/JPM-Final-Slides_Presentation-Version.pdf, accessed December 2024.
- ⁹ Future, Bio Eats World Podcast, “The Machine that Made the Vaccine,” December 18, 2020, <https://future.com/podcasts/moderna-covid-vaccine-mrna-technology/> accessed December 2024.
- ¹⁰ Pfizer, December 31, 2021, 10-K (New York: Pfizer, 2021), p. 12 and 51; and Pfizer, December 31, 2021, 10-K (New York: Pfizer, 2021), p. 10, <https://investors.pfizer.com/Investors/Financials/Annual-Reports/default.aspx>, accessed December 2024.
- ¹¹ Vilius Petkauskas, “ChatGPT Ties to Samsung’s Alleged Data Leak,” Cybernews, April 6, 2023, <https://cybernews.com/news/chatgpt-samsung-data-leak/>, accessed December 2024.
- ¹² Petkauskas, “ChatGPT Ties to Samsung’s Alleged Data Leak.”
- ¹³ “Introducing ChatGPT Enterprise,” OpenAI, August 28, 2023, <https://openai.com/index/introducing-chatgpt-enterprise/>, accessed December 2024.
- ¹⁴ Petkauskas, “ChatGPT Ties to Samsung’s Alleged Data Leak.”

2026

Connective
Leadership Summit
19 - 22 May 2026 | Melbourne VIC

Adobe: GenAI Opportunity or Threat





SUNIL GUPTA

RAJIV LAL

ALLISON CIECHANOVER

Adobe: GenAI Opportunity or Threat?

It was early-December 2022 and Shantanu Narayen, the CEO of Adobe, awaited his team in the boardroom atop the company's San Jose, California headquarters. Glancing at his phone he caught sight of headlines about the astonishing speed of adoption of OpenAI's recently released ChatGPT. The tool's unprecedented ability to generate human-like text from simple prompts had rattled the tech world and beyond. The breakthrough came on the heels of a wave of new generative AI (GenAI) image model releases over the summer. OpenAI's DALL-E 2, and the models of smaller players like Midjourney and Stability AI had captivated the public with their magic-like ability to render visual pictures from text prompts. Narayen was keenly aware of these developments and the potential disruption they could bring to Adobe's core business.

Over its 40-year history, Adobe had built software products to empower individuals to create, collaborate, and bring their visions to life. It was a global leader in digital media and marketing solutions, anchored by its Creative Cloud (including Photoshop, Illustrator, and Premiere Pro), Document Cloud (Acrobat and PDF services), and Experience Cloud (customer experience management and analytics). It served a range of customers from the world's largest enterprises, to small and medium sized business, to creative professionals in marketing, advertising, and media, to influencers, hobbyists, and students. Roughly one decade earlier, the company had successfully shifted to a subscription model, bolstering recurring revenue and deepening customer engagement.

While Adobe's performance had climbed steadily in recent years—its market capitalization surpassed \$315 billion in December 2021—the stock was down by roughly 50% year over year. Analysts pointed to macroeconomic conditions, a post-pandemic slowdown in enterprise software purchasing, and increased competition. (See **Exhibit 1.**)

While the world was buzzing with excitement about developments in GenAI models, the reaction of creative professionals—Adobe's primary customer base—was in large part that of distrust, angst, and a concern for their livelihoods. For Narayen the tension was clear: Did this technology represent a threat or an opportunity for Adobe's core? He and his team would need to make some tough decisions with respect to this burgeoning technology.

Professors Sunil Gupta and Rajiv Lal and Executive Director Allison Ciechanover (California Research Center) prepared this case with the assistance of Malte Bernholz (VP of Corporate Strategy), Dan Egan (Director of Corporate Strategy), Harshita Varshney (Manager of Corporate Strategy), and Veseline Buie (Director of Communications) all at Adobe. It was reviewed and approved before publication by a company designate. Funding for the development of this case was provided by Harvard Business School and not by the company. HBS cases are developed solely as the basis for class discussion. Cases are not intended to serve as endorsements, sources of primary data, or illustrations of effective or ineffective management.

Copyright © 2025 President and Fellows of Harvard College. To order copies or request permission to reproduce materials, call 1-800-545-7685, write Harvard Business School Publishing, Boston, MA 02163, or go to www.hbsp.harvard.edu. This publication may not be digitized, photocopied, or otherwise reproduced, posted, or transmitted, without the permission of Harvard Business School.

This document is authorized for use only by Mitchell Weiss in 2026.

This article is made available to you with compliments of Connective Leadership Summit, Innovate to Lead for your personal use. Further posting, copying or distribution is not permitted.

Overview of Adobe's History

In 1982, Dr. John Warnock and Dr. Charles Geschke left Xerox PARC to found Adobe, named after a creek near Warnock's home.¹ They launched PostScript, a computer language that translated digital text and images from a computer screen onto the printed page, to solve for file and printer incompatibility issues.² Adobe went public in 1986,³ and through the 1980s and early 1990s, it entered the consumer software market by building up and acquiring applications such as Adobe Illustrator, Photoshop, and PageMaker,⁴ which were licensed to creative professionals like graphic and web designers, videographers, and photographers. In 1993, Adobe created the Portable Document Format (PDF), allowing anyone to send an electronic document to anyone else and print it while preserving the page layout.⁵ By the late 1990s, it had grown to become a major player in the software industry.⁶

Over the 2000s, Adobe provided creative professionals with a comprehensive toolkit and promoted cross-application workflows. In 2003, it launched the Creative Suite and bundled popular products such as Photoshop, InDesign, and Acrobat.⁷ In 2005, seeking to expand its web development capabilities, Adobe acquired Macromedia—known for software like Flash and Dreamweaver—thereby broadening its reach in animation and web development. That same year, Adobe introduced Stock Photos, allowing users to search, manage, and buy royalty-free stock images.⁸ At launch, Stock Photos offered access to over 230,000 photos and illustrations.⁹ Adobe later augmented its stock image library in 2014 by acquiring Fotolia, a marketplace of over 35 million stock images and videos.¹⁰

Narayan assumed the role of CEO in late 2007, just before the 2008 financial crisis put pressure on Adobe as global corporate budgets for software contracted. The company's revenue declined and approximately 600 employees—about 8% of its workforce—were laid off to reduce costs.¹¹ Yet it was during Narayan's early tenure that Adobe expanded into enterprise marketing solutions by acquiring Omniture, an online marketing and web analytics company whose tools helped customers track and manage digital marketing efforts. Over time, these offerings grew to become another pillar of company revenues separate from creative and document software.

A Bold Shift to a Cloud-Based Subscription Model

The Omniture acquisition also spurred Adobe to rethink its monetization model. Previously, Adobe's creative software products were sold through computer software resellers and large-box retailers like Staples and Best Buy. By contrast, Omniture sold its products through a cloud-based subscription model, which had the advantage of smoother revenue streams compared to Adobe's sudden spikes when products were launched.¹² Narayan was eager to transition to a subscription model. He expected that the new cloud-based subscription might bring in new users given the lower up-front costs while also increasing the satisfaction of existing customers who would receive updates continuously online. Thus, despite pushback from longstanding users, Adobe migrated its software from a packaged sales model to a cloud-based subscription model over the next several years.¹³ The move, which had been considered risky at the outset, paid off.

The Next Chapter: Building AI into Adobe's Products

Adobe had long invested in research on cutting-edge technology, launching its internal innovation lab in 1988. An early example of its AI work was Photoshop's "red-eye" correction feature. After the Omniture acquisition, the company expanded its analytics-based AI efforts, applying predictive modeling to marketing technology and integrating insights with content creation. By 2016, Adobe's broader AI initiatives coalesced under Sensei, a framework that unified machine learning and data

analytics under one umbrella so it could be infused across the product suite.¹⁴ Adobe voiced its stance on AI: “Sensei will amplify human intelligence, not replace it.”¹⁵

The company was keenly aware of the ethical implications of AI tools. At its 2016 Adobe MAX conference, for instance, a live demo of a prototype called “Photoshop for voice” showed how easily recorded speech could be manipulated. While Adobe had already been working on watermarking recordings for authenticity, this event spurred further investments to foster ethical AI use in subsequent years. In 2019, Adobe launched its AI Ethics Board with a comprehensive review program for all its AI products. That same year, Adobe, along with the New York Times Group and Twitter, launched the Content Authenticity Initiative, a global, cross-industry coalition later joined by The Wall Street Journal, Nvidia, Microsoft, Nikon, Canon, Sony, Universal Music Group, and over 4,000 others to develop tools for digital content attribution that allowed creators and consumers to trace the origins of content—helping combat deepfakes and misinformation. To further advance a shared standard, Adobe became a founding member of Coalition for Content Provenance and Authenticity (C2PA), an open, global standards organization that developed the technical specifications behind a digital nutrition label for content, Content Credentials. In 2021, Adobe introduced Content Credentials into its product suite, letting creators attach attribution to their work.¹⁶ “The focus was on provenance, not detection forensics,” noted one executive. C2PA came to include the BBC, Intel, and Sony, among others.¹⁷

Adobe’s Business in Late 2022

Adobe was a global business with just over half of its roughly 30,000 employees based in the United States. It was long recognized as a company with a positive culture, being named among *Fortune’s* Best Workplaces in Tech.¹⁸ Many senior executives had decades-long tenures, and in several cases had “boomeranged” back to Adobe after departing for jobs elsewhere. Approximately 60% of its revenue came from the Americas, one-quarter from EMEA, and the balance from the Asia Pacific region. Adobe’s Q4 2022 ended with record-breaking revenue of \$4.53 billion, representing 10% year over year growth. The company defined its business across two strategic areas of investment – Digital Media and Digital Experience. (See **Exhibit 2** for segment breakdown and financial statements.)

Digital Media: Unleashing Creativity and Accelerating Document Productivity

The Digital Media segment, overseen by David Wadhvani, brought in nearly three-quarters of Adobe’s revenue. It offered creative professionals and consumers a platform for expression across imaging, photography, video, design, illustration, and 3-D—areas of substantial growth given the increasing importance of digital content to the global economy.

At its heart was the Creative Cloud business, which brought in \$2.68 billion in revenue in Q4, and which enabled a broad spectrum of users to create and collaborate across desktop, tablet, and mobile devices. Its primary products were available in bundles and as point solutions, serving everyone from seasoned professionals to hobbyists. Professional creators often integrated several apps into a seamless workflow, for example shooting raw footage and editing in Lightroom, refining in Photoshop, and then doing layout in InDesign. While hobbyists gravitated toward simpler template-driven tools such as Adobe Express for quick projects like invitations, social media posts, or school assignments. Along with Express, Adobe’s video segment (Premiere Pro, After Effects, and the Frame.io collaboration platform) and 3-D content tool (Substance) represented notable growth opportunities. (See **Exhibit 3** for list of products and services and customer breakdown.)

Meanwhile, the Document Cloud business generated \$629 million for Q4, growing more than 20% year on year. Document Cloud represented Adobe's initiatives to help customers migrate away from paper-based workflows and toward digitized document experiences. Adobe primarily employed a product-led growth strategy in this area. Its products included apps to help create, convert, edit, and share PDFs across workflows and devices. Their aim was to help customers increase their productivity through collaboration enabled by the ability to share, edit, track, and e-sign digital documents. The Document Cloud was built around the Adobe Acrobat family of products. To date, there were over 2.6 billion cumulative downloads of Acrobat and Reader around the world. Particular areas of growth were Acrobat Web and Acrobat Sign. (See **Exhibit 4** for list of products and services and customer breakdown.)

Digital Experience: Powering Digital Business

Adobe's Digital Experience segment, launched in 2012 and overseen by Anil Chakravarthy, represented the company's second pillar of growth. As institutions worldwide undertook digital transformations, they required new platforms and tools to engage with their customers from analytics to commerce. This field, known as customer experience management, used digital marketing tools along with customer data platforms (CDP) that collected real-time data across multiple systems.

Adobe Experience Cloud translated real-time customer data insights into profile-based, personalized digital experiences across multiple channels at scale. The underlying Adobe Experience Platform leveraged AI and offered the open architecture and ecosystem to support transformation of user data from Adobe solutions and third-party software into actionable customer profiles. Experience Cloud held the top position in Gartner's Magic Quadrant for Digital Experience for over eight years, since the inception of the category. Customers of Digital Experience included advertisers, marketers, merchants, and developers. Many of these users also used Adobe's Digital Media products, which made it easier for them to make, manage, measure, and monetize their content with coordinated workflows.¹⁹ Digital Experience solutions were sold through a direct sales force as well as through partners such as system integrators and software vendors. Customers were large enterprises such as Delta Airlines, The Coca-Cola Company, Disney, Morgan Stanley, and NASA.²⁰ Experience Cloud was the hub of a suite of applications and services and it brought in \$1.4 billion revenue in Q4. (See **Exhibit 5** for list of products and customer breakdown.)

Beyond individual products, Adobe positioned itself as an enabler of end-to-end workflows. Its tools were designed to connect the entire content supply chain. This integration allowed creative and marketing professionals at enterprises of all sizes to move seamlessly across Creative Cloud and Experience Cloud applications to create and unify brand assets, campaigns, and analytics in one cohesive workflow. (See **Exhibit 6** for sample content supply chain workflow.)

Competitive Landscape

While Adobe's Creative Cloud suite had long been the industry standard with creative professionals, the company faced competition primarily from the lower-end of the market. One of the most well-known competitors was ten-year-old, Australia-based Canva. The company had raised over \$540 million in private investments and as of late 2022, its valuation was reportedly \$26 billion.²¹ Canva offered an accessible, user-friendly online design platform that featured a vast library of templates, design elements, and easy-to-use tools. These offerings appealed to the growing set of solopreneurs and influencers. Having reached over 100 million monthly active users across over 190 countries,²² it was a popular choice for individuals and businesses seeking to create visual content, but who lacked extensive design experience. Canva's freemium model and intuitive interface placed it in competition

with Adobe Express. Meanwhile, another fast-growing challenger was San Francisco-based Figma, founded in 2012. By mid-2021, it had raised approximately \$333 million with its most recent round at a reported valuation of around \$10 billion.²³ Figma offered a collaborative, cloud-based design platform tailored for UI/UX professionals, featuring real-time co-editing and intuitive prototyping tools. On September 15, 2022, Adobe announced a plan to acquire Figma. (See **Exhibit 7a** for other competitors.)

On the Digital Experience side, Adobe faced competition from several tech incumbents offering marketing automation, analytics, and customer engagement tools. Rivals included both point solutions, like Braze, and other end-to-end marketing platforms. One of the most prominent was Salesforce, founded in 1999 and headquartered in San Francisco. Originally focused on customer relationship management (CRM) software, Salesforce had significantly broadened its scope—by September 2022, it included cloud-based marketing, data analytics, and collaboration tools. Meanwhile, Google Analytics provided a widely adopted platform for website traffic measurement, user behavior analysis, and data-driven marketing insights. Its integration with Google’s advertising ecosystem meant that many businesses looking for a seamless analytics and ad-buying pipeline saw it as an alternative to Adobe’s analytics solutions. (See **Exhibit 7b** for other competitors.)

Select Developments in GenAI Image Generation

Early experimentation with computer-generated visuals dated back to the 1960s, when cognitive neuroscientists and computer engineers at Bell Labs toyed with ASCII art to show that machines could do more than perform calculations. In the 1970s, British painter Harold Cohen developed his “AARON” system at Stanford’s Artificial Intelligence Lab, using algorithms to produce original drawings and paintings. (See **Exhibit 8** for examples.) Over subsequent decades, researchers explored neural networks—loosely modeled on the human brain—within academic or private labs.²⁴ A key development came in 1997, when teams demonstrated that neural networks could read handwritten digits in postal zip codes, paving the way for later creative applications.²⁵

By 2010, the ImageNet Large Scale Visual Recognition Challenge at Stanford accelerated computer-vision research by offering a massive, labeled dataset. Two years later, Geoffrey Hinton’s team at the University of Toronto reduced classification errors with deep convolutional neural networks, an innovation that earned Hinton a Turing Award. His foundational work on neural network models included the Boltzmann machine, which could both classify images and generate new examples based on learned patterns. In 2014, Ian Goodfellow—then a doctoral student at the University of Montreal—introduced Generative Adversarial Networks (GANs), enabling AI to generate images so realistic they could fool human observers. Around the same time, Google’s DeepDream drew public fascination with its hallucinatory imagery achieved through pattern amplification, proving that neural networks could reimagine existing visuals.²⁶ Meanwhile, researchers at Stanford were pioneering diffusion models—adding and removing noise to produce high-fidelity results. Neural Style Transfer, introduced in 2016, further showcased AI’s creative potential by blending styles of famous artists with personal photographs.²⁷

A major shift arrived in 2017 with transformer architecture, as described in Google’s paper “Attention is All You Need.”²⁸ Transformers quickly revolutionized both visual and language models setting the stage for the series of large language models (LLMs) in the ensuing years. In 2021, San Francisco-based OpenAI, the non-profit that had been infused by \$1 billion from Microsoft two years prior, made waves when it introduced a text-to-image model called DALL-E (named for Salvador Dalí and Pixar’s WALL-E). Initially built on a transformer-based approach with a discrete variational

autoencoder, DALL-E 2 had evolved into a diffusion-based model by April 2022, capable of producing more realistic visuals like the viral “armchair in the shape of an avocado.”²⁹ (See **Exhibit 9** for sample.) Within five months, OpenAI reported that more than 1.5 million users had generated over 2 million images per day with DALL-E 2.³⁰ Google followed with its own text-to-image model, Imagen, in May 2022.³¹

By late 2022 several other companies gained traction. Midjourney—a San Francisco-based research lab of roughly a dozen employees that had not raised private VC funding—had introduced a fourth version of its text-to-image service to rave reviews.³² Its Discord server reached over one million users within six months of its release (See **Exhibit 10** for sample Midjourney images).^{33,34} Stability AI, a London-based startup that had raised \$110 million in private funding, released its open-source Stable Diffusion model, enabling anyone with a capable GPU^a to create high-fidelity AI-generated art. Its model was reportedly downloaded over 200,000 times within months.

On November 30, 2022, OpenAI launched ChatGPT, a conversational interface built on top of the GPT 3.5 LLM. Its user base soared to one million within the first five days—outpacing previous tech adoption curves. While ChatGPT was focused on text generation rather than images, the media spotlight on its uncanny capabilities broadened public interest in GenAI overall.

Reactions to Generative AI Images

Tech enthusiasts and art hobbyists began posting GenAI text-to-image art on social media, often marveling at the seemingly magical capabilities. Some early experimenters embraced the technology, applauding the ease with which they could bring ideas to life.

Yet, for many creative professionals, the new technology fell short. Adobe SVP of Design Eric Snowden remarked,

There were artists who found GenAI exciting, but who likened the experience of generating images with the technology to that of a slot machine: ‘I push a button and I have no idea what’s going to come out. I have no control. I’m not a poet or writer. These models expect I’m going to take this picture I have in my head and turn it into a text prompt that accurately describes what I want. That’s not how I work. If I need to make a tweak, it just generates all over again and it’s usually a different image.’

Others held distaste for the new technology altogether, saying the lack of humanity in the creation process led to outcomes lacking soul and meaning. They claimed GenAI art was inauthentic and little more than copy-and-pasted work of previously produced human art.

While technology watchers and early adopters were tracking, and often debating, developments in GenAI image generation, the broader creative community’s interest surged in the summer of 2022. A pivotal moment came in August that year, when a Midjourney-generated piece of art won first place at the Colorado State Fair’s fine arts competition.³⁵ The winner later disclosed his use of GenAI, sparking controversy among fellow artists. Critics saw it as a harbinger of automation encroaching on

^a A GPU, or, graphic processing unit, is made of billions of microscopic transistors that control the flow of electricity and are used to create images and videos. GenAI models require thousands of GPUs to train and the number and cost of these GPUs increases with model complexity.

an inherently human domain. One social media user lamented, “We’re watching the death of artistry unfold right before our eyes.”³⁶ (See **Exhibit 11** for winning entry.)

Chief among artists’ concerns was the threat of lost income due to replacement by other artists using GenAI technology or by GenAI images outright. They fretted that those without training, credentials, or years in the profession could enter the field and steal their jobs. They also worried that their art could be used for training the models without their consent, compensation, or control. “And worse, they felt they had no recourse,” noted Snowden. Not only was their work used as training, but the models could be used to mimic an artists’ unique style. This fear was voiced by concept artist R.J. Palmer, who was dismayed to see open-source tools like Stable Diffusion suggesting prompts based on his specific style. “When they’re feeding work from living, working artists who are struggling as it is, that’s just mean-spirited,” he remarked.³⁷ Meanwhile, illustrator Greg Rutkowski discovered that over 90,000 suggested prompts invoked his name—far more than those of many classical painters like Rembrandt or Picasso. Initially pleased by the prospect of wider exposure, Rutkowski realized it muddied his creative identity online, “It’s been just a month. What about in a year? I probably won’t be able to find my real work out there because [the Internet] will be flooded with AI art.”³⁸

The Debate: How to Approach GenAI for Content Creation

Adobe had been leveraging AI to solve creatives’ pain points for years. As early as 2017, Jon Brandt of Adobe Research voiced how AI could revolutionize the creative process. At a conference that year, he remarked, “We are here to invent creative AI and transform the world through machine learning... That said, figuring out how to make it work takes time.” In what would prove a prescient remark, Brandt imagined a world where “you could simply tell the tool what you want, and it becomes an intelligent assistant that handles the tedium, without disrupting your creativity.” He concluded that with AI, Photoshop would evolve “from just a tool into your creative assistant.”

With this bold vision in mind, the Adobe team spent the next several years building AI-powered tools, including one that enabled facial editing in Photoshop. This innovation leveraged pioneering work in GANs by Alexandru Costin, Head of the Applied Research Group. A pivotal moment in Adobe’s AI journey came in 2019, when Gavin Miller, Head of Research, organized an offsite for the research team to explore AI’s potential in creative applications. During this deep dive, the team penned a white paper envisioning the next generation of Adobe’s tools—where users could simply talk to the computer rather than using a task bar. “We started to dream up a vision to reinvent the creative process. We started to see the future,” Costin recalled.

The Inflection Point

Several executives pointed to OpenAI’s release of DALL-E 2 in spring 2022 as a watershed moment. Miller commented, “The first version of DALL-E felt like nothing more than an interesting curiosity, and we had no sense of how rapidly image generation technology would improve. But when DALL-E 2 came out, it was a step-change from the first version and it happened so quickly. That’s when I realized we needed to get on the train.” Wadhvani concurred, “I saw the first avocado chair iteration back when I was working at Greylock. Sam Altman used to work out of their office. Right away I knew, this was going to be big and that we would need to dive in. The next iteration made it that much more real.”

Having kept a finger on the pulse of developments in AI alongside his team, in summer 2022, Narayan kicked off a series of meetings to formally assess the situation and craft a plan of action. Participants included senior leaders from research, applied research, product, design, marketing, and

strategy. Wadhvani recalled, “The reason is that GenAI is a little bit like alchemy. There’s a deep science to it, but there’s a lot of experimentation and testing and a little salt over the shoulder that results in the thing looking right. We needed all perspectives in the room.”

A few small early rumblings inside the company suggested that Adobe should not get involved in the new space given the reaction from some in the creative community. Others suggested a wait-and-see approach. “One option was to have stuck our head in the sand or to have decided that the technology was never going to get there,” summarized one executive. Those in this camp felt that the images made by existing models—often rudimentary, grainy or with stylized personalities—were not ready for the exacting standards of Adobe’s professional creative customers who sought commercially-ready, high-quality imagery and pixel-level control. Still others, who saw the technology as promising, expressed concern that it might be “too late to catch up” to existing players.

Nonetheless, several respected voices inside the company felt it was necessary to dive into the space. Recalled Costin, “We needed some time to convince folks internally that this was the next big thing. And that we would need to take a leap of faith and dive in.” Malte Bernholz head of Corporate Strategy noted, “It quickly became clear we were facing a once-in-a-lifetime technology shift.”

Early Decision: Build, Buy, or Partner?

Questions emerged around how exactly Adobe would undertake this leap. Over a series of meetings, executives debated whether the company should partner with or even possibly acquire one of the existing GenAI image model creators. One executive recalled, “Advocates for partnering argued it could allow us to stick to our knitting, which had always been the tooling.” The temptation to partner with an existing player was strong as it would speed time to market. This strategy aligned with concerns voiced by skeptics who questioned whether Adobe had the right talent to undertake the heavy lift and high cost required in training a model. “We were hearing about how expensive it was to train LLMs like GPT. How much would training an image model be?” asked one executive.

However, several leaders pushed back on partnering. Said one,

Look, those other companies are chasing a different dream than we are. Some are chasing AGI (Artificial General Intelligence) as an end-goal, others are a bunch of academics coming together to build cool technology. Image generation is secondary to those players. Their goals are radically different from ours. They’re not tightly intertwined with the creative community like we are. We’ll not see eye-to-eye.

Another noted, “Some of those companies had a reputational stain given how they trained their models. Partnering might bring our customers to our doors with pitch forks. Heck, our own employees might join them.” One leader asserted, “If we do this, we need to control it ourselves and own our own destiny.”

Could Adobe build a model that matched that of players who were solely focused on this area and had the talent and funding to continuously innovate? Recalling those heated discussions, several leaders pointed to the perspective held by Gavin Miller. “He told me with incredible conviction—we can build our own. And not only that, but that ours would eventually match or exceed existing models,” recalled Narayen.

Those embracing the “build” vision were bolstered by Adobe’s human capital. Its deep bench of technical and research capabilities had close ties to academia, had published multiple white papers on AI, and had served in leadership roles in a key industry group on computer graphics and technical

interfaces. Adobe's in-house design team was another asset. "We were all Adobe customers before joining Adobe, so we understand our users deeply. In many ways, our team is Customer Zero. We're a good first litmus test to understand what customers want," noted Snowden. He added, "Our user research team was already talking to customers about their perceptions of the existing models. Our product equity team makes sure our products represent people from all walks of life equally. We could bring that lens to building a model that represented everyone."

Another tailwind in Adobe's favor was its internal compute power that had meaningfully ramped starting in mid-2021. Wadhvani recalled, "Gavin came to me pointing out that we were falling behind in GPUs per employee in the research area and that to be able to continue to hire the best talent, we would need to increase that spend sizably."

Adobe also had one of the world's largest curated libraries of high-quality, royalty-free stock images, videos, and illustrations. The platform originated from Adobe's acquisition of Fotolia and had grown to over 300 million assets contributed by a global community of creators. Adobe generated licensing revenue through asset purchases and subscriptions, while remitting royalties to contributors based on downloads of their images. Adobe Creative Cloud users were able to search for, license, and incorporate these assets into their projects from within the Adobe ecosystem. In addition, Adobe could draw upon decades of deep expertise and data across creative workflows in order to build its own model.

CFO Dan Durn began to contemplate how the move would impact financials, specifically profitability. Apart from training and ongoing variable inference costs, fortunately, there would be no physical infrastructure, as Adobe would rely on their cloud partnership with AWS for compute capacity. It was clear that investors were nervous as these new technologies emerged and some questioned what the future would look like, not only for Adobe as a company, but for its creative customers. He recognized that he had to articulate the vision in a way that brought investors along on the company's journey. "I need to help investors see the long-term value of this move, and not focus so much on the short term," he recalled.

As the team plowed into building its own model, Narayen maintained a careful eye on the industry. What milestones would the team use to benchmark and judge the quality of the model Adobe would build? How might he hedge his bets along the way without rocking the team's confidence?

Decisions around Training Adobe's Model

From the outset, Adobe leaders recognized that building a GenAI image model entailed a fundamental choice: to train exclusively on content that Adobe had permission to use or to cast a much wider net by mining publicly available data from the web. The stakes were high. Adobe had cultivated over four decades of trust with creative professionals and a misstep risked alienating them. "It was well-known that the available GenAI image models had been trained by scraping data from the Internet. Content owners claimed that training on their data was a violation of their rights because it is essentially unauthorized copying of their material, in violation of the Copyright Act," commented Andrew Savage, VP Legal Digital Media. However, he noted, "Setting aside any questions around implicit or explicit permission to copy data from the Internet, most of the companies who have trained in that way also claim that what they've done should be considered fair use.^b"

^b The U.S. Copyright Act typically evaluated fair use based on: the purpose and character of the use, the nature of the original work, the amount and substantiality of the portion used, and the effect of the use on the market for the original work.

While the legality of mining data from the Internet to train generative AI models for image creation was unclear, past legal cases offered some guidance. For instance, courts had ruled that using thumbnail images of art in search results was permissible because the thumbnails were transformative, served a new and beneficial purpose, and did not directly harm the market for the original works. These decisions suggested that fair use might apply when images are used in a fundamentally different way that provides societal value, though how these principles applied to AI training remained unsettled.³⁹ Adobe Senior Director of IP, Advertising, and AI Policy J. Scott Evans summarized, “It may be a decade before we know, and it might go all the way to the Supreme Court.” Until then, other companies would be moving fast. As Bernholz noted, “The history of innovation is often at the border of legality. Just look at Uber.”

A significant contingent within Adobe believed that a model trained only on content Adobe had permission to use should be the path. “We did not want Adobe to go about training its model the way everyone else was doing it,” asserted Samantha Warren of the Adobe Design Team, who had previously worked as a creative freelancer. Such a path could leverage Adobe’s massive Stock catalogue. A decade-plus-old provision in Stock contributor terms allowed Adobe to use Stock images to create new products, services, and features. Moreover, these Stock images had undergone extensive AI and human moderation to strip them of IP and inappropriate content. A model trained exclusively on commercially safe content could also leverage content in the public domain (material with expired copyright, that which was never copyrighted, and that which was part of the Creative Commons).^c Moreover, Adobe could pay to license additional materials. Noted Bernholz,

A model trained on commercially safe content would allow added security for our customers. We could offer them indemnification whereby if something IP-protected ended up in work produced in a model we made, we would protect them legally from being sued. Such a model would not be contaminated with brand assets from other companies.

One Adobe executive noted that some large enterprise clients and agencies would simply “never touch” GenAI unless they could trust it was commercially safe.

Proponents stressed that such a model—untainted by brand assets—aligned with Adobe’s values, especially around respecting creative rights while ensuring commercial viability to enterprise customers. Evans illustrated the distinction: “If you ask one of the existing models to give you a picture of Batman, they serve it up. But if you asked a commercially safe model to do so, it would return an image of the animal bat alongside a man, as it would be unaware of Batman from its training data.” Costin clarified that creating a model trained only on commercially safe content did not mean that companies could not bring their own IP into the mix. He was already musing about an approach to allow enterprises and agencies to leverage their brand’s assets while using an Adobe model.

For many at Adobe, a model trained on exclusively commercially safe content presented precisely the level of caution customers expected. Moreover, such a model might be helpful given that Adobe played on an international field; they needed to be sensitive to regulatory and legal environment in different geographies. “Ultimately, we wanted to be viewed as the adults in the room when it came to building a model,” said Costin.

^c Creative Commons is a nonprofit that provides a framework of free, standardized licenses for creators to use when sharing their work. These licenses allow creators to grant varying levels of permission for others to use, modify, distribute, or build upon their works.

Still, others argued for a model trained on all available content from the Internet. “If everyone else is making commercially unsafe models and scraping the web, then are we shooting ourselves in the foot if we don’t? If we assume more data leads to better model outcomes, would we ever be able to match the quality of the existing models if we go the commercially safe route?” quipped one executive. “Tapping publicly available data might yield more robust outputs and richer creative possibilities, which could attract customers less concerned about content provenance,” they continued. “We serve many segments, including casual creators,” observed Snowden. “For them, it’s all about the outcome. They crave simplicity, and they’re less anxious about the source of the training data.”

Even within the creative professionals segment, might there be those who wished to access copyright protected materials for certain stages of their work? For instance, as Chief Strategy Officer Scott Belsky noted, “If you look at how creatives ideate today, they are ripping pages from magazines and pulling images from Pinterest and Google to populate their idea boards. But, when it’s time to create, they shift to commercially safe tools.” Might Adobe wish to offer these creative professionals access to the existing models from within the Adobe ecosystem?

Belsky pointed to the company’s roots as an open ecosystem. “We’ve always been a platform-agnostic company.” In fact, Adobe was already leveraging third-party LLMs in other parts of the business for text-based processes, so he advocated for allowing third-party image models onto Adobe’s platform. He drew an analogy, “Models are like cameras. They generate content that you then edit before you use it. The output of the model is not your end product. So, even if customers use third-party image models, they will still use our products.” Belsky was sanguine about developments in the competitive landscape. “Whenever another image model gets better, I want us to be happy,” he commented. Just how to incorporate third-party image models was an open item on the leadership team’s meetings. One executive noted, “If we let users access other GenAI image models on our platform, do we risk them going directly to those other models in the future? How might we live up to commercial-safety if output from these other models shows up in our ecosystem?”

How to Monetize

However Adobe decided to train its model, how it introduced and priced the tool mattered. Should it launch a standalone product? Or should the tool be interwoven into the software as a “superpower” that allowed users to go beyond text prompts to highly integrated creative workflows? The answer depended on whether Adobe prioritized strengthening its installed user base or expanding to new audiences. Dan Egan from Corporate Strategy argued for the latter, in order to delight subscribers. “Our strength is in helping customer workflows,” he noted. “The real value will come from integrating this new technology into the products and services our customers know and love.”

Deciding on a monetization model posed unique challenges. Training a cutting-edge text-to-image model from scratch could require hundreds of millions of dollars in GPU hardware and cloud infrastructure—some estimating as high as \$100 million per training cycle. Licensing additional high-quality content could further increase expenses. Unlike traditional SaaS offerings, where marginal costs (after development) were usually low, maintaining a generative model involved ongoing inference costs, meaning the company incurred costs each time a user generated an image.

Adobe considered various pricing structures. Some GenAI players, like OpenAI, used a credit system where users could generate a certain number of images for free and then had to pay for additional images. Others, like Midjourney, offered tiered subscriptions. Stability AI had taken a different path, releasing Stable Diffusion in open-source and letting others bear the infrastructure costs. Adobe would need to consider whether it should meter usage, for example charging per image, or bundle AI features into existing subscriptions, risking higher expenses, but driving rapid adoption.

Enterprise customers might prefer site-wide licensing to integrate AI-driven workflows seamlessly, but this would require careful capacity planning and had some financial exposure.

A Sense of Urgency

As Narayen closed the news app on his phone, he reflected on Adobe's mission: to change the world through personalized digital experiences. Was GenAI a step on the way to changing the world of creativity? He found himself returning to a striking prediction by Scott Branson: "GenAI will simultaneously lower the floor and raise the ceiling for creativity." On one hand, it promised to democratize creativity, enabling newcomers with limited training to produce polished imagery; on the other, it would give professionals freedom to experiment and refine very rapidly, freeing them for higher-order creative tasks. This vision energized Narayen.

He acknowledged Adobe's enduring belief that technology cannot supplant creativity as an inherently human capability, and he remained convinced that the most effective solutions would enhance human creativity. However, Narayen could not help but reflect on the tension between this view of creativity and the brutal competitive realities of the technology industry. Could GenAI strengthen Adobe's position or would it usher in new competitors that might threaten its leadership? One thing seemed certain: speed mattered in this era of unprecedented innovation. Adobe's history had shown that moving decisively at key inflection points was vital. Yet no earlier transformation had cut so close to the essence of creativity. With the leadership meeting about to begin, Narayen opened his notebook, determined to chart a path that would embrace GenAI's extraordinary power while preserving Adobe's ethos as a champion of human-led design.

Exhibit 1 Adobe Stock Chart versus NASDAQ 100 (2000 to 2022)



Source: Created by casewriters using data from Google Finance, <https://www.google.com/finance/quote/ADBE:NASDAQ?hl=en>, accessed 1/21/25.

Note: NASDAQ 100 figure has been scaled down by a factor of 10 for ease of comparison with the Adobe price per share.

Exhibit 2a Revenue Breakdown by Segment 2018 to 2022 (in billions USD)

	2018	2019	2020	2021	2022
Digital Media	6.36	7.71	9.23	11.52	12.84
Digital Experience	2.54	3.21	3.40	3.87	4.42

Source: Company documents.

Exhibit 2b Consolidated Balance Sheet Year End 2013 to 2022 (in millions USD)

	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
ASSETS										
Current Assets:										
Cash and Equivalents	835	1,117	877	1,101	2,306	1,643	2,650	4,478	3,844	4,236
Short-term investments	2,339	2,622	3,111	3,750	3,514	1,586	1,527	1,514	1,954	1,860
Total Cash & ST Investments	3,174	3,739	3,988	4,851	5,820	3,229	4,177	5,992	5,798	6,096
Trade receivables, net of allowances for doubtful accounts	600	592	672	833	1,218	1,316	1,535	1,398	1,878	2,065
Prepaid expenses and other current assets	272	271	161	156	210	312	783	756	993	835
Total current assets	4,046	4,602	4,821	5,840	7,248	4,857	6,495	8,146	8,669	8,996
Property and equipment, net	660	785	787	816	937	1,075	1,293	1,517	1,673	1,908
Operating lease right-of-use assets, net	-	-	-	-	-	-	-	487	443	407
Goodwill	4,772	4,721	5,367	5,406	5,822	10,581	10,691	10,742	12,668	12,787
Other intangibles, net	605	470	510	414	386	2,069	1,721	1,359	1,820	1,449
Deferred income taxes	-	-	-	-	-	-	-	1,370	1,085	777
Other assets	297	208	241	140	144	187	562	663	883	841
Total assets	10,380	10,786	11,726	12,697	14,536	18,769	20,762	24,284	27,241	27,165
LIABILITIES										
Current Liabilities:										
Trade payables	62	68	93	88	114	186	209	306	312	379
Accrued expenses	657	684	678	740	994	1,163	1,399	1,422	1,736	1,790
Debt	15	603	-	-	-	-	3,149	-	-	500
Deferred revenue	776	1,098	1,434	1,946	2,406	2,916	3,378	3,629	4,733	5,297
Income taxes payable	10	24	6	38	14	36	56	63	54	75
Operating lease liabilities	-	-	-	-	-	-	-	92	97	87
Total current liabilities	1,526	2,494	2,213	2,811	3,527	4,301	8,191	5,512	6,932	8,128
Long-term Liabilities:										
Debt	1,499	911	1,907	1,892	1,881	4,125	989	4,117	4,123	3,629
Deferred revenue	53	57	51	69	89	138	123	130	145	117
Income taxes payable	133	126	256	184	173	644	616	529	534	530
Operating lease liabilities	-	-	-	-	-	-	-	499	453	417
Other liabilities	445	421	297	315	405	199	313	233	257	293
Total liabilities	3,656	4,009	4,724	5,272	6,075	9,407	10,232	11,020	12,444	13,114
SHAREHOLDER'S EQUITY										
Common stock	61	61	61	61	61	61	61	-	-	-
Additional paid-in-capital	3,393	3,778	4,184	4,616	5,082	5,685	6,504	7,357	8,428	9,868
Retained earnings	6,929	6,924	7,253	8,115	9,574	11,816	14,829	19,611	23,905	28,319
Accumulated other comprehensive stock at cost	46	(8)	(169)	(173)	(112)	(148)	(188)	(158)	(137)	(293)
Treasury Stock	3,643	(3,919)	(4,268)	(5,132)	(6,084)	(7,990)	(10,615)	(13,546)	(17,399)	(23,843)
Total stockholder's equity	6,725	6,776	7,001	7,424	8,460	9,363	10,530	13,264	14,797	14,051
Total liabilities and stockholder's equity	10,380	10,786	11,726	12,697	14,535	18,770	20,762	24,284	27,241	27,165

Source: Company documents.

Exhibit 2c Consolidated Income Statement Year End 2013 to 2022 (in millions USD)

	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
Revenue										
Subscription	1,138	2,077	3,224	4,585	6,134	7,604	9,634	11,626	14,573	16,388
Product	2,470	1,628	1,125	800	707	622	648	507	555	532
Services and other	447	443	446	469	461	804	889	735	657	686
Total revenue	4,055	4,147	4,796	5,854	7,302	9,030	11,171	12,868	15,785	17,606
Cost of Revenue										
Subscription	278	335	409	462	623	574	926	1,108	1,374	1,646
Product	138	97	90	69	57	46	40	36	41	35
Services and other	170	190	245	289	330	575	707	578	450	484
Total cost of revenue	587	622	744	820	1,010	1,195	1,673	1,722	1,865	2,165
Gross Profit	3,469	3,525	4,051	5,035	6,292	7,835	9,498	11,146	13,920	15,441
Operating Expenses										
Research and development	827	844	863	976	1,224	1,538	1,930	2,188	2,540	2,987
Sales and marketing	1,620	1,652	1,683	1,910	2,198	2,621	3,244	3,591	4,321	4,968
General and administrative	520	543	532	576	625	745	881	968	1,085	1,219
Amortization of intangibles	52	52	69	79	77	91	175	162	172	169
Total operating expenses	3,046	3,112	3,148	3,540	4,124	4,995	6,230	6,909	8,118	9,343
Operating Income	423	413	903	1,494	2,168	2,840	3,268	4,237	5,802	6,098
Non-operating Income (expense):										
Interest expense	(68)	(60)	(64)	(70)	(74)	(89)	(157)	(116)	(113)	(112)
Investment gains (losses), net	(4)	1	1	(2)	8	3	52	13	16	(19)
Other income (expense), net	5	8	34	14	36	40	42	42	-	41
Total non-operating expenses, net	(67)	(51)	(29)	(58)	(30)	(46)	(63)	(61)	(97)	(90)
Income before income taxes	356	361	874	1,435	2,138	2,794	3,205	4,176	5,705	6,008
Provision for (benefit from) income taxes	66	93	244	266	444	203	254	(1,084)	883	1,252
Net income	290	268	630	1,169	1,694	2,591	2,951	5,260	4,822	4,756
Basic net income per share	\$ 0.58	\$ 0.54	\$ 1.26	\$ 2.35	\$ 3.43	\$ 5.28	\$ 6.07	\$ 10.94	\$ 10.10	\$ 10.13
Shares used to compute basic net income per share	501	498	499	498	494	491	486	481	477	470
Diluted net income per share	\$ 0.56	\$ 0.53	\$ 1.24	\$ 2.32	\$ 3.38	\$ 5.20	\$ 6.00	\$ 10.83	\$ 10.02	\$ 10.10
Shares used to compute diluted net income per share	513	508	507	504	501	498	492	485	481	471

Source: Company documents.

Exhibit 3a Creative Cloud Products and Services

Product	Features
<i>Adobe Photoshop and Lightroom</i>	<ul style="list-style-type: none"> • <i>Photoshop</i>: A widely used digital imaging and design application, available on desktop, iPad, and via a web-based beta. Provides advanced editing, effects tools, and neural filters. • <i>Lightroom</i>: A cloud-based photo service for editing, organizing, storing, and sharing images across multiple devices, offering features such as Select People, Select Objects, and Content-Aware Remove. • A Photography Plan bundles both Photoshop and Lightroom for photographers and enthusiasts.
<i>Adobe Illustrator and Fresco</i>	<ul style="list-style-type: none"> • <i>Illustrator</i>: An industry-standard vector graphics application for creating illustrations and digital graphics on desktop and iPad. Used for print, web, interactive, video, and mobile media. • <i>Fresco</i>: An illustration app (free or premium) on iPhone, iPad, Windows devices, and Wacom MobileStudio devices. Combines vector, raster, and Live Brushes (powered by Adobe Sensei) for a natural painting and drawing experience.
<i>Adobe InDesign</i>	<ul style="list-style-type: none"> • A design and layout application for print and digital media. Often used to create books, eBooks, magazines, posters, and interactive PDFs. Integrates with Photoshop, Illustrator, Acrobat, and Adobe Stock, and enables creating and publishing apps for iOS, Android, and Amazon devices.
<i>Adobe Premiere Pro and Premiere Rush</i>	<ul style="list-style-type: none"> • <i>Premiere Pro</i>: A professional nonlinear video editing application for filmmakers, TV editors, YouTubers, and videographers. Supports various media formats and integrates with Frame.io for review and approval. • <i>Premiere Rush</i>: An all-in-one video editing app (free or premium) with simplified editing tools for quick content creation and direct sharing to platforms like YouTube, Instagram, and TikTok.
<i>Adobe Express</i>	<ul style="list-style-type: none"> • A web-based platform and mobile app focused on easy content creation for first-time creators and communicators. Offers guided tools, one-click solutions, templates, and a Content Scheduler for planning and publishing social media posts. Includes free and premium features.
<i>Adobe Stock</i>	<ul style="list-style-type: none"> • A collection of millions of high-quality, royalty-free photos, vectors, illustrations, videos, templates, audio, and 3D assets. Integrated with many Creative Cloud apps for streamlined asset discovery and licensing, with free and paid options.
<i>Substance 3D</i>	<ul style="list-style-type: none"> • An ecosystem of desktop apps—including Stager, Painter, Sampler, Designer, and Modeler—for 3D design and texturing. Allows users to assemble 3D scenes, texture assets with brushes and Smart Materials, and import professional 3D textures from the Substance 3D Assets library. Modeler enables VR and desktop-based sculpting.
<i>Adobe After Effects</i>	<ul style="list-style-type: none"> • A motion graphics and visual effects application for cinematic titles, effects, and animations. Integrates with Premiere Pro, Photoshop, Illustrator, and other tools. Includes Frame.io integration for real-time collaboration, scene edit detection, and 3D Extended Viewer.
<i>Frame.io</i>	<ul style="list-style-type: none"> • A cloud-native video collaboration platform with real-time uploads, review, and approval. Integrates directly with Premiere Pro and After Effects for frame-specific commenting. Camera to Cloud automatically uploads footage from on-set devices for immediate editing and feedback.
<i>Behance</i>	<ul style="list-style-type: none"> • A social community where creators can showcase work and discover others' projects online. Integrates with Adobe Portfolio for building customizable websites, and supports live-streaming from Creative Cloud applications.

Source: Company documents.

Exhibit 3b Creative Cloud Customer Breakdown

Segment	Members
Business Professionals	Social Media Teams, Sales Reps, HR Teams, Communication Teams, Small Business Owners
Consumers	Students, PTAs, Personal use cases like travel, invitations, announcements
Creative Professionals	Graphic Designers, Photographers, Videographers, Illustrators, 3D Artists
Creators	Social Media Influencers, Solopreneurs, Travel Bloggers, Gig Workers
Marketing Professionals	Brand Managers, Channel Marketers, Campaign Strategists, Performance Marketers, Social Marketers, Data Analysts, Marketing Ops

Source: Company documents.

Exhibit 4a Document Cloud Products and Services

Product	Features
<i>Adobe Acrobat</i>	<ul style="list-style-type: none"> • A leading application for creating, converting, and editing PDFs. Provides secure, reliable, and compact PDF document creation, plus collaborative workflows with commenting, editing, sharing tools, and direct integration with Acrobat Sign. • Also available as part of the Creative Cloud All Apps subscription.
<i>Acrobat Web</i>	<ul style="list-style-type: none"> • A web-based suite of PDF tools that allows users to edit, comment, convert, organize, and sign PDF documents directly in the browser. Offers quick, one-click features, including both free and premium capabilities.
<i>Adobe Acrobat Reader</i>	<ul style="list-style-type: none"> • A free software tool for viewing, annotating, and printing PDF documents on desktop and mobile. Includes options to create, edit, export, combine, share, and collaborate on PDFs, along with "Liquid Mode" for easy reading on mobile devices. • Integrates with Adobe Document Cloud as well as third-party storage services.
<i>Adobe Scan</i>	<ul style="list-style-type: none"> • A free mobile app that captures paper documents as images and transforms them into full-featured PDFs using Document Cloud services. Enables quick scanning, sharing, and storage on any mobile device.
<i>Adobe Acrobat Sign</i>	<ul style="list-style-type: none"> • A cloud-based e-signature service that allows secure electronic sending, signing, and tracking of documents on any device. • Integrates with enterprise systems through APIs, Adobe Experience Manager Forms, and Advanced Workflows for Acrobat Sign, enabling seamless web and mobile e-signature experiences.

Source: Company documents.

Exhibit 4b Document Cloud Customer Segment

Segment	Members
Business Professionals	Social Media Teams, Sales Reps, HR Teams, Communication Teams, Small Business Owners
Consumers	Students, PTAs, Personal use cases like travel, invitations, announcements

Source: Company documents.

Exhibit 5a Digital Experience Products and Services

Product	Features
<i>Adobe Experience Manager</i>	A content management solution for websites, mobile apps, and other digital touchpoints, used for content authoring, asset management, digital forms, and personalized content delivery
<i>Adobe Analytics</i>	A data analytics platform for tracking user interactions across multiple channels, providing real-time segmentation, attribution modeling, and predictive insights through interactive dashboards
<i>Adobe Target</i>	A personalization and testing engine that tailors user experiences, enabling A/B testing, multivariate experiments, and AI-driven recommendations
<i>Adobe Audience Manager</i>	A data management platform that aggregates audience data from various sources, creating high-value segments for ad targeting and personalization across channels
<i>Adobe Campaign</i>	A cross-channel campaign management tool for orchestrating messaging, supporting email, SMS, and push notifications, with data-driven automation and personalized workflows
<i>Marketo Engage</i>	A marketing automation solution geared toward lead management and account-based marketing, facilitating personalized campaigns, scoring models, and sales-marketing alignment
<i>Adobe Commerce</i>	An e-commerce platform enabling omnichannel storefront creation, managing product catalogs, payment integrations, and shopping experiences at scale
<i>Adobe Workfront</i>	A work management platform designed for project planning and collaboration, providing workflow automation, asset review, and progress tracking
<i>Adobe Journey Optimizer & Real-Time CDP</i>	Solutions for unified customer profiles and real-time journey orchestration, leveraging AI to personalize engagement across channels using data from multiple sources

Source: Company documents.

Exhibit 5b Digital Experience Customer Breakdown

Segment	Members
Creative Professionals	Graphic Designers, Photographers, Videographers, Illustrators, 3D Artists
Marketing Professionals	Brand Managers, Channel Marketers, Campaign Strategists, Performance Marketers, Social Marketers, Data Analysts, Marketing Ops

Source: Company documents.

Exhibit 6 Sample Content Supply Chain Workflow

Stage	Description	Key Adobe Tools
Workflow & Planning	Teams align on campaign goals, resources, and customer data	Workfront
Creation & Production	Creative teams produce assets and content at scale	Creative Cloud, Express, Frame.io
Asset Management	Stakeholders collaborate, edit, and ensure compliance	Experience Manager
Delivery & Activation	Personalized campaigns and assets launched across channels	Native connectivity to any Experience Cloud delivery application; integrates with 3rd party or custom delivery apps
Reporting & Insights	Performance tracked, insights generated, and content refined	Content analytics

Source: Company documents.

Exhibit 7a Select Competition in Digital Media

Competitor	Company Value	Estimated Annual Revenue	Key Products/Offerings	Competitive Advantages
Autodesk	\$40.32 billion market cap	\$4.4 billion	AutoCAD, Maya, Inventor	<ul style="list-style-type: none"> • 3D design, engineering, and entertainment software • Industry-standard tools for architecture, engineering, and manufacturing
Corel	\$1 billion acquisition value by KKR	\$250 million	CorelDRAW, Corel Painter, Corel Photo-Paint	<ul style="list-style-type: none"> • Graphics, photo editing, and digital painting software • Intuitive tools and features for creative professionals
Serif (Affinity)	Private, undisclosed	\$60 million	Affinity Designer, Photo, Publisher	<ul style="list-style-type: none"> • Professional-grade design software for graphic designers, photographers, and publishers • One-time purchase option

Source: Compiled by casewriters.⁴⁰

Exhibit 7b Select Competition in Digital Experience

Competitor	Market Cap	Annual Revenue	Key Products/Offerings	Competitive Advantages
Google	\$1.2 trillion	\$282 billion	Google Analytics, Google Ads, Google Marketing Platform	Wide range of marketing and analytics tools, including web analytics, advertising, and marketing automation
HubSpot	\$15 billion	1.73 billion	HubSpot Marketing Hub	Marketing automation platform that helps businesses attract visitors, convert leads, and close customers
Microsoft	\$1.9 trillion	\$198 billion	Microsoft Dynamics 365	Suite of business applications, including CRM and marketing capabilities
Oracle	\$220 billion	\$42.4 billion	Oracle Customer Experience (CX) Cloud Suite	Comprehensive suite of cloud applications, including CRM, marketing automation, and data analytics
Salesforce	\$135 billion	\$26.5 billion	Salesforce Sales and Marketing Cloud	Market leader in CRM space and has AI-enabled sales and marketing automation

Source: Compiled by casewriters.⁴¹

Exhibit 8 Select Images using Harold Cohen's "Aaron"

Drawing, computer-generated, with hand colouring, by Harold Cohen, 1974.
© Victoria and Albert Museum, London



Drawing, computer-generated, with hand colouring, by Harold Cohen, 1977-1982.
© Victoria and Albert Museum, London

Excerpt from Cohen's 1982 work, *How to Make a Drawing*:

"Aaron was clearly not a tool in an orthodox sense. It was closer to being a sort of assistant, if the need for an human analogue persists, but not an assistant which could learn what I wanted done by looking at what I did myself, the way any of Rubens' assistants could see perfectly well for themselves what a Rubens painting was supposed to look like. This was not an assistant which could perform any better for having done a thousand drawings, not an assistant which could bring anything approximating to a human cognitive system to bear on the production of drawings intended for human use.

We are living on the crest of a cultural shock-wave of unprecedented proportions, which thrusts a new kind of entity into our world: something less than human, perhaps, but potentially capable of many of the higher intellectual functions – it is too early still to guess HOW many – we have supposed to be uniquely human. We are in the process of coming to terms with the fact that "intelligence" no longer means, uniquely, 'human intelligence.'"

Source: Images available via Kate Vass, "Harold Cohen: 'Once upon a time there was an entity named Aaron,'" *Kate Vass Galerie*, April 30, 2024, <https://www.katevassgalerie.com/blog/harold-cohen-aaron-computer-art>, accessed January 18, 2025. Excerpt from Harold Cohen, "How to Make a Drawing," December 17, 1982, available via <https://www.aaronshome.com/aaron/publications/how2make.pdf>, accessed May 2025.

Exhibit 9 Advancements DALL-E to DALL-E 2



Source: Images available via Sarah Guo blogpost, "Thinking (creatively) with the help of machines," April 7, 2022, <https://sarahguo.com/blog/dalle2>, accessed January 20, 2025.

Exhibit 10 Sample Midjourney Image, August 2022

Source: Kyle Kukshtel blogpost, "Midjourney Changed my Mind on AI-Generated Art," June 24, 2022, <https://kylekukshtel.com/midjourney>, accessed January 20, 2025.

Note: Image generated by Midjourney using the prompt, "A Game Developer Writing a Blog Post."

Exhibit 11 Colorado Art Fair Winner, Theatre de Opera Spatial



Source: Jason Allen, “Théâtre D’opéra Spatial,” 2022, available via https://commons.wikimedia.org/wiki/File:Th%C3%A9%C3%A2tre_D%E2%80%99op%C3%A9ra_Spatial.jpg, accessed April 2025.

Note: According to video game artist and contest entrant Jason Allen, he spent 80 hours working on this piece of art. He entered the text prompts into Midjourney, which generated 800 images, from which he selected his favorite three. He reportedly used Photoshop to edit the images before entering them to the art competition.

Endnotes

- ¹ The Naeth, "Adobe Inc's History," *Medium*, April 4, 2024, https://medium.com/@The_naeth/adobe-incs-history-c0dfac3b6844, accessed January 2025.
- ² Ibid.
- ³ Brian Withers, "If You Had Invested \$1,000 in Adobe's IPO, Here's How Much It Would Be Worth Now," *The Motley Fool*, December 27, 2019, [https://www.nasdaq.com/articles/if-you-had-invested-\\$1000-in-adobes-ipo-heres-how-much-it-would-be-worth-now-2019-12-27](https://www.nasdaq.com/articles/if-you-had-invested-$1000-in-adobes-ipo-heres-how-much-it-would-be-worth-now-2019-12-27), accessed January 2025.
- ⁴ "Adobe Systems Inc. History," *Funding Universe*, <https://www.fundinguniverse.com/company-histories/adobe-systems-inc-history/>, accessed January 2025.
- ⁵ "The History of the PDF: Timeline," *Adobe*, <https://www.adobe.com/acrobat/resources/pdf-timeline.html>, accessed January 2025.
- ⁶ Piiffner, Pamela, *Inside the Publishing Revolution: The Adobe Story*. Adobe Press, 2002, https://www.adobe.com/misc/pdfs/chap4_pg1thru7.pdf, accessed January 2025.
- ⁷ "How Adobe Became a Successful \$95 Billion SaaS Company," *Product Habits Blog*, <https://producthabits.com/adobe-95-billion-saas-company/>, accessed January 2025.
- ⁸ Max Dunn, "Adobe Stock: the History," *Silicon Publishing*, June 16, 2015, <https://www.siliconpublishing.com/blog/2015-06-adobe-stock-the-history/>, accessed January 2025.
- ⁹ Ibid.
- ¹⁰ Frederic Lardinois, "Adobe Acquires Stock Content Marketplace Fotolia For \$800M in Cash," *TechCrunch*, December 11, 2014, <https://techcrunch.com/2014/12/11/adobe-acquires-stock-content-marketplace-fotolia-for-800m-in-cash/>, accessed January 2025.
- ¹¹ Leena Rao, "Confirmed: Adobe to Cut 9 Percent of Workforce," *TechCrunch*, November 10, 2009, <https://techcrunch.com/2009/11/10/layoffs-reported-at-adobe/>, accessed January 2025.
- ¹² Ron Miller, "Ten years after Adobe bought Omniture, the deal comes into clearer focus," *TechCrunch*, September 16, 2019, <https://techcrunch.com/2019/09/16/ten-years-after-adobe-bought-omniture-the-deal-comes-into-clearer-focus/>, accessed January 2025.
- ¹³ Ibid.
- ¹⁴ Giselle Abramovich, "Making Sense of Adobe Sensei," *Adobe Blog*, October 4, 2017, <https://blog.adobe.com/en/publish/2017/10/04/making-sense-of-sensei>, accessed January 2025.
- ¹⁵ Abhay Parasnis, "Mastering Creative and Experience Intelligence: A Year with Adobe Sensei," *Adobe Blog*, November 2, 2017, <https://blog.adobe.com/en/publish/2017/11/02/a-year-with-adobe-sensei>, accessed January 2025.
- ¹⁶ Will Allen, "Adobe unleashes Content Attribution features in Photoshop and beyond at MAX 2021," *Adobe Blog*, October 26, 2021, <https://blog.adobe.com/en/publish/2021/10/26/adobe-unleashes-content-attribution-features-photoshop-beyond-max-2021>, accessed January 2025.
- ¹⁷ "Introducing Official Content Credentials Icon," *Coalition for Content Provenance and Authenticity*, <https://c2pa.org/post/contentcredentials/>, accessed January 2025.
- ¹⁸ "Adobe receives recognition as a top company that cares for its employees and their communities," *Adobe Blog*, September 5, 2022, <https://blog.adobe.com/en/publish/2022/08/31/adobe-receives-recognition-as-top-company-that-cares-for-employees-their-communities>, accessed January 2025.
- ¹⁹ "Adobe Inc. – Notice of 2022 Annual Meeting of Stockholders and Proxy Statement," *Adobe*, <https://www.adobe.com/cc-shared/assets/investor-relations/pdfs/adbe-proxy-2022.pdf>, accessed January 2025.
- ²⁰ "Adobe Q1 FY 2022 Earnings Call," *Adobe*, March 22, 2022, <https://www.adobe.com/cc-shared/assets/investor-relations/pdfs/c234g8n4848s.pdf>, accessed January 2025.

- ²¹ Lachlan Andrews, "Canva Announces USD 40 Billion Valuation Fueled by the Global Demand for Visual Communication," *Canva*, September 14, 2021, <https://www.canva.com/newsroom/news/canva-announces-usd-40-billion-valuation-fueled-global-demand-visual-communication/>, accessed January 2025.
- ²² Naveen Kumar, "Canva Statistics 2025 (Users, Market Share & Revenue)," *Demandsage*, December 30, 2024, <https://www.demandsage.com/canva-statistics/>, accessed January 2025.
- ²³ Joanna Glasner, "Figma's Venture Backers Win Big with Adobe's \$20B M&A Deal," *Crunchbase News*, September 15, 2022, <https://news.crunchbase.com/news/venture-backers-win-figma-index-greyllock/>, accessed January 2025.
- ²⁴ Andrey Kurenkov, "A Brief History of Neural Nets and Deep Learning," *Skynet Today*, September 27, 2020, <https://www.skynettoday.com/overviews/neural-net-history>, accessed January 2025.
- ²⁵ Reid Pryzant, "A Quick Tour of Deep Learning History," *Stanford University Computer Science*, https://cs.stanford.edu/~rpryzant/blog/dl_history/dl_history.html, accessed January 2025.
- ²⁶ "Deep Dream: An In-Depth Exploration," *Geeks for Geeks*, August 1, 2024, <https://www.geeksforgeeks.org/deep-dream-an-in-depth-exploration/>, accessed January 2025.
- ²⁷ Htoyryla, "Neural networks, style transfer and artistic process," *Erratic Explorations Blog*, December 11, 2016, <https://liipetti.net/erratic/2016/12/11/neural-networks-style-transfer-and-artistic-process/>, accessed January 2025.
- ²⁸ Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin, "Attention is All You Need," *Google Research*, 2017, <https://research.google/pubs/attention-is-all-you-need/>, accessed January 2025.
- ²⁹ Will Douglas Heaven, "This avocado armchair could be the future of AI," *MIT Technology Review*, January 5, 2021, <https://www.technologyreview.com/2021/01/05/1015754/avocado-armchair-future-ai-openai-deep-learning-nlp-gpt3-computer-vision-common-sense/>, accessed January 2025.
- ³⁰ "DALL-E now available without waitlist," *OpenAI*, September 28, 2022, <https://openai.com/index/dall-e-now-available-without-waitlist/>, accessed January 2025.
- ³¹ Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar, Seyed Ghasemipour, Burcu Karagol Ayan, S.Sara Mahdavi, Rapha Gongjjo Lopex, Tim Salimans, Jonathan Ho, David Fleet, Mohammad Norouzi, "Imagen: unprecedented photorealism x deep level of language understanding," *Google Research*, <https://imagen.research.google/>, accessed January 2025.
- ³² Barbara Pazur, "What is Midjourney? Here's What You Need to Know About the AI Image Generator," *CNET*, February 5, 2025, <https://www.cnet.com/tech/services-and-software/midjourney-what-it-is-and-what-you-should-know-about-the-ai-image-generator/>, accessed February 2025.
- ³³ How Midjourney Built a Business On Discord," *Discord*, <https://discord.com/build-case-studies/midjourney#:~:text=The%20results%20speak%20for%20themselves,an%20astonishing%20one%20million%20users>, accessed January 2025.
- ³⁴ Beni Edwards, "Too Easy - Midjourney tests dramatic new version of its AI image generator," *Ars Technica*, November 9, 2022, <https://arstechnica.com/information-technology/2022/11/midjourney-turns-heads-with-quality-leap-in-new-ai-image-generator-version/>, accessed January 2025.
- ³⁵ Kevin Roose, "An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy," *New York Times*, September 2, 2022, <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>, accessed January 2025.
- ³⁶ Sarah Kuta, "Artist Jason Allen placed first in a Colorado contest, generating debate about A.I.'s role in art," *Smithsonian Magazine*, September 6, 2022, <https://www.smithsonianmag.com/smart-news/artificial-intelligence-art-wins-colorado-state-fair-180980703/>, accessed January 2025.
- ³⁷ Will Knight, "Algorithms Can Now Mimic Any Artist. Some Artists Hate It," *Wired*, August 19, 2022, <https://www.wired.com/story/artists-rage-against-machines-that-mimic-their-work/>, accessed January 2025.
- ³⁸ Melissa Heikkilä, "This artist is dominating AI-generated art. And he's not happy about it," *MIT Technology Review*, September 16, 2022, <https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/>, accessed January 2025.

³⁹ “Summary of PERFECT 10 v. AMAZON,” 508 F.3d 1146 (9th Cir. 2007), https://yalelawtech.org/wp-content/uploads/2011/09/perfect10_v_amazon.pdf, accessed January 2025.

⁴⁰ Compiled by casewriters from: Company Value; <https://companiesmarketcap.com/autodesk/marketcap/>, <https://www.thurrott.com/windows/209695/corel-acquired-for-over-1-billion>; Annual Revenue; <https://www.macrotrends.net/stocks/charts/ADSK/autodesk/revenue>, https://tracxn.com/d/companies/corel/_ala855OREaISKaiVCmTXfhmZV7mouzo6M_LhKzvg-8#about-the-company, <https://www.afr.com/technology/canva-s-billion-dollar-bet-on-a-37-year-old-nottingham-company-20240321-p5fea0>, and; Key Products and Competitive Advantages; <https://www.autodesk.com/products>, <https://www.corel.com/en/all-products/>, <https://affinity.serif.com/en-us/affinity-pricing/?srsltid=AfmBOoqnVEWqW0zwqveML9dneYC9HkkiEb6RmgpnQjRPL-AGMKrMKOFu>, accessed May 2025

⁴¹ Compiled by casewriters from: Market Cap; <https://companiesmarketcap.com/alphabet-google/marketcap/>, <https://companiesmarketcap.com/hubspot/marketcap/>, <https://www.macrotrends.net/stocks/charts/MSFT/microsoft/market-cap>, <https://companiesmarketcap.com/oracle/marketcap/>, <https://www.macrotrends.net/stocks/charts/CRM/salesforce/market-cap>; Annual Revenue; <https://companiesmarketcap.com/alphabet-google/revenue/>, <https://www.macrotrends.net/stocks/charts/HUBS/hubspot/revenue>, <https://www.microsoft.com/investor/reports/ar22/index.html#:~:text=Amid%20this%20dynamic%20environment%2C%20we,revenue%20for%20the%20first%20time.>, <https://investor.oracle.com/investor-news/news-details/2022/Oracle-Announces-Fiscal-2022-Fourth-Quarter-and-Fiscal-Full-Year-Financial-Results/default.aspx>, <https://www.zdnet.com/article/salesforce-posts-record-q4-results-with-1b-in-sales-from-the-data-business/>, and; Key Products and Competitive Advantages; <https://marketingplatform.google.com/about/analytics/>, <https://www.hubspot.com/products/marketing-varb>, <https://www.microsoft.com/en-us/dynamics-365>, <https://www.oracle.com/cx/>, <https://www.salesforce.com/marketing/>, accessed May 2025.



2026

Connective Leadership Summit

19 - 22 May 2026 | Melbourne VIC

