

Sam Altman Biography: Pioneer of the AI Revolution

Table of contents, preface, 7 parts, 20 chapters



Kim Kyung-jin, Kim Kyung-ran

English PDF edition

Sam Altman Biography: Pioneer of the AI Revolution

Kim Kyung-jin, Kim Kyung-ran

An online biography following Sam Altman's childhood, startups, Y Combinator, OpenAI, ChatGPT, the 2023 board crisis, and his sense of responsibility in the AI era.

Table of Contents

Preface

1.1 Childhood and School Years

1.2 A Nineteen-Year-Old Raises 30 Billion Won

1.3 Leading Y Combinator: The Kingmaker of Startups

2.1 Sensing the AI Revolution and Founding OpenAI

2.2 The Journey from Model Development to Products

2.3 Conflict with Elon Musk

3.1 The CEO Ouster: A Five-Day Nightmare

3.2 The Employees' Dramatic Loyalty

3.3 Overcoming the Crisis and Its Lessons

4.1 Philosophy of Organization and Development

4.2 OpenAI's Future Business Model

4.3 Productivity and Work Management

5.1 Where Models Are Heading

5.2 The Role of Agents and Coding

5.3 The Future of Human-AI Interaction

6.1 Regulation and the National Security Race

6.2 Infrastructure and Energy Challenges

6.3 Coexistence Between AI and Human Society

7.1 Personal Inspiration and Vision

7.2 Final Advice and Outlook

Preface

See all chapters

Next →

On November 30, 2022, ChatGPT was released to the public. After about thirty minutes of conversation, a chill ran down my spine. It felt like stepping into a fundamentally different era. One million users in five days. One hundred million in two months. I had seen a lot of change, but never anything this fast.

There is a joke going around: 'BC doesn't stand for Before Christ; it stands for Before ChatGPT.' You could laugh it off, but it's not entirely wrong. The world before ChatGPT and the world after are completely different.

Researching AI policy brought me into contact with people from around the world. The United States said, 'Don't get in the way of innovation.' Europe said, 'Safety comes first.' China said, 'We will lead the way.'

But on one point, everyone agreed: 'There is no going back to before ChatGPT.'

Sam Altman says, 'The age of intelligence is the age of energy.' At first, I didn't know what he meant. Studying AI policy made it clear: building AI intelligence requires enormous electrical energy, and energy will ultimately determine the limits of the future.

Over the years, I've met many talented people. But I've met even more who never got a chance. Not living in Seoul. Not attending the right schools. Parents without the money. No funds for private tutoring. It was heartbreaking.

When Sam Altman talks about 'maximizing human potential,' this is what he means: solving problems like these. What if AI could give every person access to the highest level of knowledge and education? What if a child in a rural village could receive the same education as a child in Gangnam, Seoul? What if your eight-year-old had a personal tutor who could teach them everything the world knows? What if cancer could be diagnosed ten times faster?

That is the real revolution.

I showed this book to young students first. An elementary school student asked, 'Can I become someone like Sam Altman?' A middle school student said, 'AI is kind of scary, but also kind of amazing.'

Both answers are correct.

The world our children grow up in will be completely different from the one we know. It may be frightening. It may be strange. It may be exciting.

You might become someone who creates change, like Sam Altman. Or you might become someone who meets that change with wisdom. Either way is fine.

Beginning with this preface, we trace one person's life to make sense of a new era.

Kim Kyung-jin, Kim Kyung-ran, 2025

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

1.1 Childhood and School Years

Part 1: From an Attic Computer Kid to Silicon Valley Kingmaker

[← Previous](#)

[See all chapters](#)

[Next →](#)

The Attic Computer Kid Connects with the World.

It begins on April 22, 1985, in an ordinary Chicago household. His mother, Connie Gibstein, was a dermatologist, and his father, Jerry Altman, was a real estate broker. The eldest of four children, Sam moved with his family to Clayton, a suburb of St. Louis, Missouri. He was four years old. It was 1989.

Clayton was a typical American Midwest neighborhood. Quiet streets, well-tended lawns, rows of old brick houses. On the surface, it had the look of a perfectly ordinary American middle-class home. Evenings, the family gathered for dinner and passed the time playing ping-pong or pool.

Something extraordinary entered Sam's life when he was eight years old.

An Apple Macintosh computer, a gift from his parents. For young Sam, it was a door to an entirely new world. While other children played baseball outside or rode bikes, Sam sat in the attic and spent his hours in front of the computer. Many nights, the light stayed on in that attic window well past bedtime.

Sam wasn't just playing games on the computer. He took it apart. He wanted to know how the hardware worked. He unscrewed each bolt, examined the components, and put it back together, over and over. An eight-year-old doing this.

He started learning to program.

It amazed him that typing code made something appear on screen. It felt like magic that the computer obeyed his commands. Sam called himself a "computer nerd." His friends called him that too. He was quiet, but once something caught his attention, he focused on it like a laser.

There's a story that, as a kindergartener, he'd worked out how the area code system operated. Most children that age still struggle to count, yet Sam had already figured out the rules governing how area codes were assigned across the United States. His ability to read numbers, patterns, and systems was beyond his years.

To Sam, a computer wasn't just a machine. It was a window onto the world. In AOL chat rooms, he found people like himself. He talked with strangers under an alias, shared what troubled him, and found some comfort. That Macintosh was his lifeline to the outside world.

"Discovering AOL chat rooms changed my life."

Countless nights in the attic. Hours spent typing code in front of a computer screen, browsing the internet, discovering new worlds. Those experiences made Sam who he was, someone curious, persistent, and convinced he could change the world.

Sam later said: "I was the kid sitting in the attic learning to program late into the night. I think there's a straight line from that time all the way to OpenAI today."

Without that attic in St. Louis, there might be no ChatGPT. All those hours young Sam spent in front of a computer were, in their own way, the starting point of the AI revolution that would reshape human history.

High School Years: Water Polo, Swimming, Fencing, and Poker

He enrolled at John Burroughs School, a private institution in the Ladue area of St. Louis. It was one of the most prestigious private high schools in the United States, with annual tuition exceeding \$30,000, drawing students from wealthy families. The school was known for its progressive and open educational philosophy.

At school, Sam was still the "computer geek." He built the school's website himself and developed software to manage class schedules. He was fifteen, sixteen years old when he did this. Friends called him a "bookworm," but Sam wasn't only about studying.

He loved sports.

He threw himself into water polo. It's a physically demanding sport played in the water, requiring both stamina and teamwork. Sam played on the water polo team and practiced alongside his brothers. After school, he'd go to Forest Park to ride bikes or swim.

He also fenced. Fencing demands split-second judgment. You read your opponent's movement, react fast, and think strategically. The quick decision-making Sam later showed in business may well have been learned on the fencing strip.

Friends remembered Sam as "kind but relentlessly tenacious." Once he set a goal, he dug in until the end. Whether it was a computer project, a sport, or anything else.

By graduation, Sam was no ordinary teenager. His programming skills were at university level. He had experience running an online business. He had opinions on social issues. He'd built physical and mental strength through sports.

Sam had tenacity. He was the type who, once he started something, saw it through to the end. If an obstacle appeared, he'd find thirty different ways around it. That quality showed up later when he founded Loopt, led Y Combinator, and built OpenAI.

Sam Altman during his John Burroughs School years.

The attic computer kid was now preparing to step out into the world.

Stanford and a Fateful Choice

In 2003, nineteen-year-old Sam Altman enrolled at Stanford University.

Stanford. Located at the heart of Silicon Valley, one of the world's foremost universities. Google, Yahoo, Hewlett-Packard, Instagram. Countless technology companies were born in Stanford dormitories and classrooms. For Sam, Stanford was "the obvious choice."

He majored in computer science. The obvious choice, given he'd been programming since the age of eight. Stanford's CS program was among the best in the world, with outstanding professors, cutting-edge research facilities, and above all, the people he wanted to be around.

At Stanford, Sam didn't just learn to code.

He took "every science class." Physics, biology, chemistry. He took writing courses too. And humanities. When friends asked why a CS major was signing up for those, Sam answered:

"Looking back later, these classes will turn out to be far more useful than the CS ones."

He was right. Leading OpenAI, Sam was never just an engineer. He was a philosopher too, an economist, a policy thinker. Understanding what AI would mean for humanity required more than technical knowledge.

Sam put it this way: "The most important thing about college was being surrounded by really smart people pursuing all kinds of interesting ideas."

He said the physics students fascinated him most. They debated the nature of the universe, worked through complex math problems through the night, and showed him how to see the world differently.

Sam had a different kind of "class."

Poker.

Every night, a poker game broke out with friends. Sam would later confess:

"Poker paid for my college expenses. I learned more from poker than from lectures."

What did he learn?

Pattern recognition. The ability to observe opponents' behavior, read their psychology, and anticipate their next move. How to make decisions under uncertainty. The courage to calculate risk and sometimes bet big. The resilience to control emotions and get back up after failure.

Sam said: "Two things have taught me a lot about business. Poker and angel investing. I strongly recommend both."

In a 2018 podcast, he added: "Poker isn't for everyone, but it's a good way to learn about the world, business, psychology, and risk."

In his sophomore year at Stanford, Sam made a pivotal decision.

He would drop out.

His friends were stunned. Drop out of Stanford? Leave one of the world's top universities without a degree? Sam was certain.

"I'm a sophomore, and I'm leaving (right now)."

That was the message Sam sent when Y Combinator's Paul Graham hesitated to invest. Sam had decided to found his first startup, Loopt, and school was just getting in the way.

Paul Graham's partner Jessica Livingston recalls meeting Sam for the first time: "I could tell immediately that he was far wiser than his age."

Sam spent two years at Stanford. He left without a diploma, but walked away with things far more valuable: the decision-making skills he'd honed at the poker table, a broad perspective drawn from many fields of study, and above all, the conviction that he had to act now.

In 2017, the University of Waterloo in Canada awarded Sam an honorary doctorate in engineering. The irony: he never finished his undergraduate degree, yet ended up with a PhD.

As he left Stanford, Sam must have thought: "What school teaches is the past. I want to build the future."

And he really did.

When AI Didn't Work: Professor Andrew Ng's Lab

The summer between his freshman and sophomore years at Stanford.

Sam Altman worked somewhere remarkable: the artificial intelligence lab of Professor Andrew Ng.

Who is Andrew Ng? He is a legendary figure in AI. He founded Google Brain, co-founded Coursera, and stands as one of the most influential researchers in machine learning. In the early 2000s, he was leading the Stanford AI lab.

Sam worked in that lab all summer. A nineteen-year-old college student getting his first exposure to cutting-edge AI research.

There was a problem.

The AI didn't work.

AI in the early 2000s was nothing like it is today. An AI that could hold a natural conversation like ChatGPT was beyond imagining. Back then, AI couldn't even handle simple tasks properly. Image recognition? Terrible. Natural language processing? Nearly impossible.

Computing power was scarce. Data was scarce. The algorithms were primitive. Researchers ran small experiments, failed, tried again, and failed again.

AI was considered a "fringe" field. People didn't believe it would ever be genuinely useful.

Sam was different.

He wasn't frustrated that AI didn't work. If anything, it made him more curious. Why?

Sam put it this way: "Even if the odds of success are low, if it works, it'll be the coolest, most important, most interesting thing out there."

This is the concept of expected value. Something he'd learned at the poker table. If your odds are only 10%, but the payoff when you win is enormous, the bet is still worth making.

In the lab, Sam watched AI models fail. They couldn't recognize simple patterns, produced nonsensical results, and were useless.

He recalls: "We didn't know what we were supposed to be doing. It felt completely hopeless."

Most experiments failed. Teaching a robotic hand to grasp objects. An AI learning to play video games. None of it worked.

Sam left the lab. He went off to found Loopt. It would take some time before he found his way back to AI.

AI never fully left his mind. That technology that refused to work. Those experiments that failed again and again. The memory kept nagging at him.

In the early 2010s, something began to shift.

IBM's Watson beat human champions on Jeopardy. Google DeepMind's AlphaGo defeated the world champion in Go.

Sam felt it in his gut. "Something big is coming."

In late 2015, he co-founded OpenAI.

The early days of OpenAI were still filled with uncertainty. In 2016, the office held maybe fourteen people. "We had very strong conviction, a clear sense of direction, but no real execution plan."

People debated in front of whiteboards, read papers, and designed experiments.

At first, no one could agree on a direction. Large language models, or LLMs, felt like a distant dream. So they considered building AI that could master video games. They ran experiments with robotic hands.

Nvidia CEO Jensen Huang personally delivered the first DGX-1 system. It was a symbolic moment , the opening shot of an AI revolution.

Sam Altman's journey into AI began in the years when none of it worked.

A lab full of failures. A technology nobody believed in. A future that seemed impossibly far away.

Sam saw possibility inside those failures. Nobody knew that ChatGPT would change the world a decade later. Sam probably didn't know either, not exactly.

He believed. That even a technology that doesn't work , if you keep trying , will work someday.

And he was right.

Without that summer at Stanford's AI lab, there would have been no OpenAI. No ChatGPT, no GPT-4, no AI revolution to reshape humanity.

Sometimes failure teaches more than success ever could. Sam Altman lived through the years when AI didn't work , and that's precisely how he learned to make it work.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

1.2 A Nineteen-Year-Old Raises 30 Billion Won

Part 1: From an Attic Computer Kid to Silicon Valley Kingmaker

[← Previous](#)

[See all chapters](#)

[Next →](#)

From Viendo to Loopt: The Evolution of an Idea

In the spring of 2005, in a Stanford University dormitory, nineteen-year-old Sam Altman stood at a crossroads in his life. A computer science major who had just finished his sophomore year, he was two years away from holding a Stanford diploma that anyone would envy. After graduation, he could land a comfortable position at a big company like Google or Apple, his parents would be proud, and his future would be secure.

Reflecting on that time, Sam said: 'This wasn't my life plan. Another opportunity just came along.' The 'other opportunity' he meant was starting a company. As the app he and his friends were building became more concrete, Sam grew increasingly restless sitting in class. His mind was entirely on the service he wanted to create.

The project Sam and his friends were developing was called 'Viendo.' A Spanish word meaning 'to see,' the name carried the idea of being able to 'see' where your friends were at any moment. Today, sharing a friend's location on a messaging app feels ordinary, but in 2005 there were no smartphones. The iPhone didn't arrive until 2007.

Sam and his team envisioned a service that used a phone's GPS to show each friend's real-time location on a map. Instead of calling every friend to ask 'What are you up to tonight?', you could open the app and find nearby friends instantly. You might stumble upon a friend at the same café, or join up with someone near the cinema you were already heading to.

The idea was solid, but the name wasn't working. 'Viendo' was hard to pronounce and didn't click immediately. The team put their heads together for weeks before landing on 'Loopt.' There's an English expression 'in the loop,' meaning to stay connected and informed. Loopt took its name from exactly that phrase: short, memorable, and conveying the sense of always being connected with the people around you.

As the name took shape, so did Sam's resolve. He sat down with his parents. 'I want to take a break from school and try starting a company.' His parents were worried, but Sam was confident. He later explained his thinking: 'The key thing I realized was that I could always go back. That's what risk really comes down to. Most things aren't one-way streets. You try something, and if it doesn't work, you can reverse course.'

And so Sam Altman put his student ID in a drawer and stepped into a new life as CEO of Loopt. Not a dormitory but a small office. Not classes but investor meetings. Not exams but product development. At the time, no one could have imagined that this young man would one day lead the AI revolution.

A Nineteen-Year-Old Raises 30 Billion Won

Sam and his friends had decided to start Loopt, but they had no money. A good idea alone doesn't build a company. They needed to rent servers, hire more developers, and find office space. Above all, they needed living expenses while building the app. College students couldn't run a company on spending money.

Sam set out to find investors. Silicon Valley had plenty of venture capitalists who backed young people with promising ideas. Sam made the rounds with Loopt's business plan in hand. Every meeting came with the same questions: 'Who's going to use this?' 'How will you make money?' 'What happens if Facebook or Google builds the same thing?'

The rejections piled up. Some investors said, 'Come back after you actually finish college.' Others shook their heads: 'Location sharing? Isn't that a privacy violation?' Pitching his idea in front of adults many years his senior was not easy.

But Loopt's concept was genuinely fresh. In 2005, social networks existed almost entirely on computer screens. People read friends' posts on Facebook, looked at photos, left comments. Loopt had an ambitious plan to pull those online connections into the physical world. It was a bridge between the digital and the real.

At its peak, Loopt had over five million registered users and partnerships with every major carrier in the United States. That came later, of course, but early investors saw that potential from the start.

A turning point arrived. Patrick Chung of XFund invested five million dollars through New Enterprise Associates. At the exchange rate of the time, that came to over five billion won. Putting five billion won into a company run by a nineteen-year-old college dropout was a serious commitment.

Patrick Chung recalled his first meeting with Sam: 'He was far too mature for his age. Those weren't words you'd expect from a nineteen-year-old.' Sam didn't stop at walking investors through app features. He painted a vivid picture of the future Loopt would build: people meeting more often, connecting more easily, cities growing a little warmer.

What impressed investors wasn't the idea alone. It was the attitude. Young founders usually sit tensely in front of investors, cautious and careful to say the right thing. Sam was different. He wasn't pleading for investment; he was making an offer. He radiated confidence: 'Invest in our company and you will see a real return.'

When word of the five-million-dollar investment spread, other venture capital firms started paying attention. Well-known firms like Sequoia Capital put money into Loopt. Sequoia was famous for backing Google, Apple, and YouTube in their early days. Its decision to invest in Loopt signaled that

Silicon Valley had taken Sam's vision seriously.

In the end, Sam raised over thirty million dollars in venture capital for Loopt, more than thirty billion won. For a college-age founder at the time, it was a remarkable feat. Thirty billion raised on a single idea. The path there included dozens of rejections, business plans rewritten through the night, and the exhausting work of persuading investors one by one.

The experience gave Sam lessons he carried forward: a good idea isn't enough on its own; you have to be able to communicate it persuasively; and you need the persistence to keep going after rejection. Those capabilities became significant assets when he later led OpenAI.

A Fateful Meeting with Y Combinator

No account of Loopt's success can leave out Y Combinator (YC). YC was a startup support program that the legendary programmer Paul Graham had just launched in 2005. It was unlike a conventional investment firm. Rather than simply writing a check, it spent three months working intensively with startups, offering advice, and connecting founders with each other.

Sam and the Loopt team applied to YC's first batch. YC was accepting startups for the very first time, and applications flooded in. Among the applicants was the team that would go on to build Reddit. Competition was fierce.

YC's selection process was distinctive: not a long, elaborate interview, but a short ten-minute conversation. In that brief window, applicants had to make the case for their idea and their team. Paul Graham and the other partners fired sharp questions, and applicants had to answer on the spot.

Sam's turn came. The nineteen-year-old stood before a Silicon Valley legend. Paul Graham had met hundreds of founders and had an instinct for who would succeed. There was no fooling him.

The interview went well. Sam explained Loopt's concept clearly and handled technical questions without hesitation. Paul Graham nodded, visibly interested. Then, near the end of the conversation, a decisive question came.

'If we decide to invest, what will you do about school?'

It was not a casual question. The YC program demanded full commitment for three months. It wasn't something you could fit around a class schedule. Taking the investment meant leaving school, and giving up Stanford was an enormous decision for most people.

Many student founders stumbled at this question. They offered hedged answers like 'I'll find a way to balance school and the startup.' Paul Graham had no patience for that kind of answer. Starting a company takes 100 percent commitment.

Sam Altman didn't hesitate for a second.

'I'm a sophomore, and I'm leaving (right now).'

The room went quiet. Paul Graham looked directly at Sam. Sam had been straightforward about his year ('sophomore'), then made his decision without a pause ('I'm leaving'). No excuses, no hesitation, no conditions. A clean, decisive answer.

Paul Graham later recalled that moment: 'That one line told me who Sam was. That's what a real founder looks like: someone who doesn't make excuses, finds a way, and moves forward.'

Sam's answer became legend in YC history. Many students who later applied to YC heard this story, along with the advice: 'Answer like Sam Altman.' That single sentence wasn't just an answer. It was a declaration of the way Sam approached life.

Sam co-founded Loopt with Nick Sivo, his classmate and boyfriend at the time. They received six thousand dollars in funding from the founder program and lived in an early tech entrepreneur community in Cambridge while building their product.

YC decided to invest in Loopt. Loopt was one of the first eight companies accepted into Y Combinator in 2005. This was the beginning of Sam's long relationship with YC. He would later become YC's president and one of the most influential people in Silicon Valley. All of it started with that ten-minute interview.

"I'm a sophomore, and I'm going." Those few words changed Sam's life. Sometimes life's most important moments arrive this briefly. Only those who are ready can catch them before they slip away.

I'm a Sophomore, and I'm Leaving Right Now

Loopt grew quickly after receiving YC's investment. Sam and the team worked around the clock on app development. The app, riddled with bugs at first, gradually stabilized, and features were added one by one: seeing friends' locations, sending messages, getting recommendations for nearby places. Loopt was turning into a polished service.

The early response was enthusiastic. It was popular among college students. Instead of texting "Hey, where are you?", you could just open the app and see exactly where your friends were. You could find a friend at a party, spot someone in the same library, or run into a friend who happened to be nearby.

Steve Jobs also took notice of Loopt. When he unveiled the iPhone in 2007, Jobs introduced Loopt as one of the apps running on the device. Getting a spot on Apple's stage was every app developer's dream. Sam and his team felt enormously proud.

The company grew, too. At first it was just Sam and a few friends, but they gradually started hiring. Developers, designers, marketers, salespeople, the team swelled to dozens. Sam was no longer just a developer writing code. He had to become a CEO leading a company.

But something felt off. The explosive early growth was starting to slow. The user base kept growing, but not many people stuck around. Plenty downloaded the app because it looked fun, tried it once or

twice, then deleted it. Why?

There was a battery problem. Keeping GPS running continuously drained a phone's battery fast. Phones back then didn't have anything like the battery life they do today.

Second, privacy concerns were real. "Wouldn't it be uncomfortable if other people always knew where I was?" That's how many users felt.

Third, smartphones hadn't gone mainstream yet. Between 2005 and 2008, the iPhone had just launched, and most people were still using regular cell phones.

The biggest problem was something else entirely. Giant companies like Facebook and Google had started adding similar features to their own services. Facebook built a check-in feature; Google added location sharing to Google Maps. Loopt was a small startup, but its competitors were behemoths with hundreds of millions of users.

Sam was stuck in a dilemma. Should he hold on, or find a different path? Investors were starting to worry too. "When will you turn a profit?" "User numbers keep slipping, is everything okay?" The pressure mounted.

In 2012, Sam made a hard call. He would sell Loopt to a banking company called Green Dot Corporation for \$43.4 million (roughly 47 billion won). Handing over a company he had built for seven years was not an easy decision. It must have felt like giving up a child for adoption.

Sam later said in an *Intelligencer* interview that he "left in a pretty unhappy state." Because the sale proceeds had to be split with investors, the amount Sam actually walked away with was roughly \$5 million (about 5 billion won). It was a lot of money, but a modest return for seven years of work.

From the outside, Loopt looked like a "failed startup." It hadn't achieved anything like the success of Facebook or Instagram. Sam acknowledged this honestly. He wrote on Twitter: "My first startup failed pretty badly. It was really hard!"

Does failure mean it's over? No. Through Loopt, Sam learned things money can't buy.

The importance of timing. Loopt's idea was good, but it arrived too early. Sam later recalled: "The technology really had potential. Location services soon became a key component of almost every mobile application, from banking to gaming to news. But Loopt never found its footing." Everyone uses location services now, but in 2005 people weren't ready for it.

You have to build what users actually want. There's a difference between something technically impressive and something people use every day. Loopt was technically excellent, but it wasn't something people genuinely needed in their daily lives.

The world doesn't end when you fail. Sam said: "Other people don't care about your failures as much as you do." And indeed, after Loopt fell apart, Sam's reputation stayed intact. If anything, he was recognized as someone who had run a company for seven years.

The value of the people you meet. While building Loopt, Sam encountered countless investors, founders, and engineers. Those connections later became a major asset. With the money from the Loopt sale, Sam co-founded a venture firm called Hydrazine Capital in 2012 with his brother Jack Altman. Now he was the one writing checks, not waiting for them.

Loopt was a company that failed, but it was the school that made Sam Altman. Without that experience, he could never have become YC's president or led OpenAI. Sometimes failure teaches more than success. Sam learned that the hard way.

Resourcefulness and Persistence

If you had to name the single hardest moment for Sam, it was negotiating with the carriers. Between 2005 and 2010, getting an app onto a phone required the carrier's permission. Today you can download anything freely from an app store, but back then it worked differently. Carriers controlled which apps ended up on phones.

Sam needed to persuade the major American carriers: AT&T, Verizon, Sprint, T-Mobile. Getting their approval was the only way to put Loopt in front of a large audience. The problem was that these companies had no interest in a small startup.

When Sam reached out to the carriers, he didn't even get replies. Emails went ignored. Phone calls ended with "We'll pass your message along to the right person." A nineteen-year-old trying to do business with corporate giants that had been around for decades was like throwing eggs at a wall.

Most people would have quit. "Okay, carriers aren't going to work. Let's find another way." But Sam was different. Carrier deals were essential to Loopt's success, and he refused to walk away.

Sam recalled something Paul Graham had told him: "Be relentlessly resourceful." What does resourcefulness mean? It's the ability to throw every possible approach at a problem. When one method fails, you try another; when that fails, you find yet another. A resourceful person has no word for "quit" in their vocabulary.

When a direct approach didn't work, Sam started looking for a way around. When an email to an AT&T contact went unanswered, he tracked down that person's supervisor. When that didn't work, he found someone in a different department. He combed through company websites for names, researched their backgrounds on LinkedIn, and looked for any thread he could pull to make a connection.

He attended every telecom industry conference held in Silicon Valley. He paid the steep entry fees, then scanned the room for tables where carrier executives were sitting. During lunch, he would manufacture a coincidence and take the seat next to one of them, striking up a conversation: "Hi, I'm Sam, and I run a startup called Loopt..."

He flew directly to carrier headquarters without an appointment. He waited in building lobbies until someone he recognized walked by, then handed over a business card and asked for just five

minutes. He was escorted out by security more than once.

Sam reportedly tried more than thirty different approaches to land a deal with one carrier. Thirty times. Most people would decide after three tries that "this just isn't going to happen." Sam tried thirty times and was already preparing for the thirty-first.

And each approach was different. With people in the technical division, he emphasized Loopt's innovative technology. With people in marketing, he made the case that Loopt could appeal to younger demographics. With senior executives, he pitched a new revenue model. He changed his message depending on who was sitting across from him.

The breakthrough came. A mid-level manager at AT&T, worn down by Sam's persistence, agreed to set up a meeting. "This guy won't stop bothering us, let's just meet with him once." That meeting eventually led to a contract.

Once he had signed with one carrier, others started paying attention. "AT&T made the deal, if we don't, are we falling behind?" The dominoes fell. In the end, Loopt succeeded in forming partnerships with every major carrier in the United States.

This experience gave Sam a lesson he would never forget: even seemingly impossible problems can be solved if you try long enough and in enough different ways. He later put it this way.

"Most people give up after the first attempt. Some make it to a second try. But the people who actually solve problems are the ones who keep going to the thirtieth attempt. Remarkably, if you keep finding new approaches, most problems eventually give way."

That same resourcefulness showed itself years later when Sam was building OpenAI. When the company needed enormous computing power but lacked the funds, he persuaded Microsoft to commit billions of dollars in investment. When government regulation loomed as a threat, he went to Congress himself and testified. When competitors pulled ahead, he shipped better products faster.

The story of Loopt's carrier deal cuts to the heart of who Sam Altman is. He is not a prodigious programmer, nor did he grow up in a wealthy family. What set him apart was a refusal to quit and a gift for finding some new angle no matter how many doors had already been closed.

There will be moments in life when a wall appears and looks completely immovable, whether in college admissions, a job search, a business, or a dream you are chasing. After one or two failed attempts it is easy to conclude that it simply was not meant to be. Remember Sam Altman's story. The people who actually succeed are the ones who keep going to the thirtieth attempt, and who are ready, if necessary, to make a thirty-first.

Loopt never became a company that changed the world. But it turned Sam Altman into someone who would. That is the most lasting thing Loopt left behind.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

1.3 Leading Y Combinator: The Kingmaker of Startups

Part 1: From an Attic Computer Kid to Silicon Valley Kingmaker

[← Previous](#)

[See all chapters](#)

[Next →](#)

After selling Loopt, a new chapter opened in Sam Altman's life. In 2011, he returned to Y Combinator, the very place where he had received his first investment as a young founder. This time, not as someone seeking investment, but as someone making it. Then in 2014, at just 28 years old, he became president of Y Combinator. This period was a defining stretch of time in which Sam Altman grew into one of the most influential figures in Silicon Valley.

The Man Who Built Thousands of Startups

Have you heard of Y Combinator? It is the most famous startup school in Silicon Valley. Young founders come with their ideas, and Y Combinator provides funding, advice, and guidance to help them grow their businesses. Services we use every day, like Airbnb, Dropbox, and Reddit, all got their start here.

In 2014, Sam Altman became president of Y Combinator at the age of 28. By any measure, he was young. But Y Combinator's founder, Paul Graham, said this about Sam: "Drop him in the middle of an island of cannibals, and come back five years later, he'll be their king." A striking comparison, isn't it? It reflects just how much Paul Graham believed in Sam's abilities.

When Sam took the helm at Y Combinator, he set out to make bold changes. He set a goal for Y Combinator to invest in 1,000 new companies every year. At the time, that was a staggering number, almost impossible to imagine. Many people shook their heads and asked, "How could you possibly manage that many companies?" But Sam saw it differently.

"There are so many problems in the world that need solving. Climate change, disease, energy. To tackle problems this large, we need far more founders. We have to open the doors wider."

To lower Y Combinator's barriers to entry, Sam built several new programs. He launched Startup School, a free online resource where anyone could learn to build a company, and started an internship program connecting developers with startups. The idea was simple: no matter where in the world you were, if you had a good idea, Y Combinator could help.

What Sam had in mind wasn't simply numbers. He believed the type of company Y Combinator backed mattered far more. At the time, Silicon Valley was overflowing with easy startups, photo-sharing apps, food delivery apps. Those services do make daily life more convenient, of course. But Sam was drawing a bigger picture.

He put it this way: "There are thousands of photo-sharing apps. But how many startups are out there that can build an experimental nuclear fusion reactor?" Sam wanted Y Combinator to be not just a place that backed companies chasing trends, but one that developed hard technology companies, the kind that could genuinely change humanity's future.

What is hard technology? It means technologies like artificial intelligence, biotechnology, nuclear fusion energy, and space exploration, things that take years to develop and carry high risk, but that could completely transform the world if they succeed. These technologies are difficult to monetize quickly. The risk of failure is high. Most investors are reluctant to back companies like these.

Sam was different. He believed that "when people say you're crazy for starting something, that's a good sign." He had a deep admiration for founders who took on what everyone else called impossible. In practice, under Sam's leadership, Y Combinator invested in companies like Helion Energy, which is working on nuclear fusion power.

In 2015, Sam made another significant decision: he created a fund called YC Continuity. The fund was \$700 million in size and was designed to provide follow-on investment to Y Combinator graduates that had already reached some scale, helping them grow even larger.

Why was this fund necessary? Y Combinator had always focused on very early-stage startups. It gave money to founders just getting started, ran an intensive three-month program, and connected them with investors. That was Y Combinator's role. But what happened after that? As companies grew and needed more capital, they had to go find other investors on their own.

Sam identified a problem there. "Some of the excellent companies we've built stall out and can't get to the next stage. We need to be with them all the way." That was the thinking behind YC Continuity. It was a commitment to not just plant a young tree, water it for a while, and walk away, but to keep tending it until it grows into a great forest.

During the five years Sam led Y Combinator, roughly 1,900 companies passed through its program. Among them were companies now valued at billions of dollars, including DoorDash, Instacart, Twitch, and Coinbase. Sam had truly become a kingmaker. The companies he chose and the founders he advised went on to become Silicon Valley's new royalty.

What Sam was after was not simply producing a large number of successful companies. He told founders consistently: "Making money matters, but making the world a better place matters more." He wanted Y Combinator to be not just an investment institution, but a community where people with dreams of changing the world could find each other.

This philosophy of Sam's seeped into the entire culture of Y Combinator. The organization taught its founders: "Move fast. Don't be afraid to fail. Listen to your customers. But above all, build a company that solves a problem that actually matters."

There is an amusing story. During an interview, Sam was asked whether he had any plans to take Y Combinator public. He said no, flatly. When asked why, he said this:

"If we go public, some hedge fund manager will call us up and shout, 'You're one cent short of your quarterly earnings! Y Combinator is finished!' I don't want that life. We should be working with our eyes on the world ten or twenty years from now, not on quarterly results."

That answer captures Sam's philosophy in full. He trusted long-term value over short-term results. He wanted sustainable growth, not rapid gains. And above all, he wanted to build companies that could genuinely change the world.

Sam Altman's years as president of Y Combinator were among the most important of his life. Through that time, he met countless founders and learned what makes a company succeed and which technologies could truly reshape the future. All of that experience became a major asset when he later took the helm at OpenAI.

In 2019, Sam stepped down as president of Y Combinator. But even after he left, Y Combinator still holds its place as the world's leading startup accelerator. The systems and culture Sam put in place continue to breathe. The seeds he planted as kingmaker are still growing today.

Mentor, Paul Graham: How to Lead Without Crushing People

Every great person has a mentor. Footballer Son Heung-min had veteran coaches, and J.K. Rowling, who wrote the Harry Potter series, had an editor who believed in her. For Sam Altman, Paul Graham was exactly that person.

Who is Paul Graham? He is the founder of Y Combinator and one of Silicon Valley's most respected thinkers. A programmer, writer, and investor, he has shared deep insights on startups and technology through dozens of essays. Many Silicon Valley founders read his writing like scripture.

In 2005, when 19-year-old Sam Altman came to Y Combinator with Loopt, Paul Graham saw something special in this young founder, something beyond being simply smart or driven. He wrote of Sam: "Drop him in the middle of an island of cannibals, and come back five years later, he'll be their king."

The comparison is stark, but Paul Graham's message was clear. Sam Altman was someone who, placed in any difficult situation, would not just survive but completely turn it around. Paul Graham's prediction proved right.

After Loopt was acquired, Sam returned to Y Combinator in 2011. This time not as a founder seeking investment, but as a partner helping other founders. Then in 2014, Paul Graham made a remarkable decision: he named the 28-year-old Sam Altman as his successor and handed him the presidency of Y Combinator.

"Sam is one of the smartest people I know. And he understands startups better than anyone else I know, including me." Those words from Paul Graham carried deep trust in Sam.

What Paul Graham taught Sam was not simply business know-how. The most important lesson he passed on was how to deal with people.

Silicon Valley is a fiercely competitive place. To succeed, you sometimes have to be ruthless, make decisions fast, and show no weakness. Many CEOs treat their employees like tools and have no qualms about sacrificing someone to hit a target.

Paul Graham was different. He taught Sam the same thing he practiced himself: a philosophy of leading without crushing people.

What does that mean? Put simply, a leader should not suppress people through power but draw out the best in them, not control employees through fear, but guide them through respect and trust.

Paul Graham practiced this philosophy in running Y Combinator. He would say bluntly to founders, "Your idea isn't good enough," but he never dismissed or belittled them. Instead, he worked through it with them, why it wasn't working, and how to make it better. He encouraged founders so they could keep trying without being paralyzed by fear of failure.

Sam Altman watched Paul Graham and learned a great deal from what he saw. Working at Y Combinator, first as a partner and later as president, he made Paul Graham's approach his own.

Picture Sam mentoring founders. He never issued one-sided commands: 'Do it this way' or 'Do it that way.'

Instead, he asked questions.

"Who are your customers?" "What do they actually want?" "Why do they need your product?" "If this approach doesn't work, what else could you try?"

Through these questions, founders found their own answers. And answers they discovered for themselves carried far more force than anything handed down to them.

Paul Graham gave Sam another piece of advice that stayed with him for life: "Be relentlessly resourceful."

What does that mean? When problems arise, don't give up easily, try every possibility. If one approach fails, try another. If that fails too, find yet another way.

Sam had lived this advice at nineteen. While running Loopt, the team needed a deal with a carrier. Carriers didn't work with small startups. Most founders would have given up after a few rejections.

Sam didn't give up. He approached the carrier in more than thirty different ways: emails, phone calls, introductions through mutual contacts, staged chance encounters at events. Eventually a key decision-maker said: "I'll meet with you just to make you stop bothering us."

It sounds like a joke, but it really happened. That meeting led to a signed deal. That is the power of being relentlessly resourceful.

Paul Graham and Sam Altman were not boss and subordinate. They were mentor and student who held each other in deep respect, and eventually became friends. Graham passed on not just the craft

of building a business but wisdom about how to live, and Sam took those lessons and shaped them into his own style.

No relationship is always perfect. When Sam left Y Combinator in 2019, there were reports of some tension between the two, arising from Sam spending more time on OpenAI than on Y Combinator. But that was part of Sam moving toward a larger ambition.

What matters is that Graham's teachings shaped Sam throughout his life. The philosophy of leading people without crushing them would later shine at OpenAI.

In 2023, when the OpenAI board removed Sam Altman as CEO, something remarkable happened. Nearly all of OpenAI's 770 employees signed a letter declaring that if Sam did not return, they would all leave the company.

Where did that fierce loyalty come from? It came from the way Sam treated them. He never crushed or suppressed his employees. He respected them, listened to them, and helped them do their best work. So when the crisis hit, they stood up to defend him.

This is the leadership that Paul Graham taught Sam Altman. Not controlling people through power, but leading them through trust. Building an organization on respect, not fear. That is leadership.

Sam Altman was fortunate to find a mentor like Paul Graham. He absorbed those teachings and developed them into his own approach. Graham planted the seed; Sam grew it into something enormous.

That Sam Altman became one of the most important leaders of the AI era is not simply because he was smart or ambitious. It is because he had a mentor like Paul Graham who believed in him and showed him the right path, and because Sam humbly accepted those lessons and put them into practice.

You too will meet a mentor someday. When that happens, what matters is how seriously you take the lessons and make them your own. The story of Sam Altman and Paul Graham is a fine example of how good mentorship can change a person's life, and through that life, the world.

Rethinking the Limits of the Technologies in Vogue

In the early to mid-2010s, Silicon Valley was in the golden age of apps. As smartphones spread across the world, a flood of apps poured out, each promising to change how people lived their daily lives.

Photo-sharing apps like Instagram completely transformed how people connected. People spent their days sharing photos with friends, hitting like, and posting comments. Food delivery apps like Uber Eats and DoorDash let people enjoy their favorite meals without leaving home. Uber and Lyft reinvented how people hailed a ride, and Airbnb changed how they traveled.

These apps attracted enormous followings. Tens of millions of people used them, and investors poured billions of dollars into the companies behind them. Most of the startups applying to Y Combinator arrived with ideas in the same vein.

"An app that makes your photos more beautiful!" "An app that delivers food even faster!" "An app that makes it easier to meet up with friends!"

On the surface, it looked like a perfect era. Young founders could make fortunes off a single idea, and investors were scrambling to find the next unicorn, a company valued at over one billion dollars.

But Sam Altman sensed something was off in the middle of all that excitement.

As president of Y Combinator, he reviewed dozens of startup ideas every day. He began to notice that more and more of them looked alike. There were thousands of photo-sharing apps alone. Hundreds of delivery apps. Each touted slightly different features, but at their core they were all doing the same thing.

Watching this, Sam fell into deep thought. "Is this really all there is? Is making photos prettier or delivering food five minutes faster truly the most important problem humanity needs to solve?"

To be fair, these services do have value. They make life more convenient, create jobs, and stimulate the economy. Sam didn't deny that. Y Combinator itself invested in successful delivery companies like DoorDash and Instacart, and those bets paid off handsomely.

But Sam's gaze was fixed on something further away.

In 2011, he watched IBM's Watson beat human champions on the quiz show Jeopardy! In 2016, he witnessed the historic moment when Google DeepMind's AlphaGo defeated Go world champion Lee Sedol. Watching those events, Sam felt it: "Something big is coming."

He had a clear sense of the limits of sharing apps and delivery services. What were those limits?

The first limit was the scale of impact. No matter how successful a photo-sharing app might be, it doesn't solve humanity's fundamental problems. It can't stop climate change, cure disease, or fix the energy crisis. Keeping people entertained and connected matters, of course. But Sam didn't think that was enough.

The second limit was the depth of innovation. Most app services simply made existing things a bit more convenient. Food delivery, for instance, had always existed; delivery apps just digitized the process and made it more efficient. That counts as innovation, but not the kind Sam was looking for.

Sam was drawn to Zero to One innovation: creating something entirely new in the world. The phrase, also the title of a book by the well-known investor Peter Thiel, captures the idea of creating value that never existed before, going from zero to one, rather than improving something that already exists, going from one to n.

The third limit was market saturation. Thousands of sharing apps meant the market was already crowded. New startups entering that space had little chance of success; most failed, and only a handful survived. Sam believed it was more meaningful to chart an unexplored blue ocean than to fight in a red ocean where the competition had long since grown fierce.

Sam grew increasingly drawn to "hard technology." What is hard tech? It refers not to "soft" technologies like apps that can be built in a matter of months, but to the "difficult" technologies that require years or even decades of research and development.

For example,

These technologies are extraordinarily difficult to develop. They require significant funding, and the risk of failure is high. Generating immediate returns is nearly impossible. Most investors steer clear of this space. "It's too risky. Better to invest in apps with a clear path to profit."

Sam thought differently. He put it this way: "There are thousands of photo-sharing apps. But how many startups are capable of building an experimental nuclear fusion reactor?"

He believed Y Combinator should support startups tackling precisely these hard problems. Under his leadership, Y Combinator did invest in Helion Energy, a nuclear fusion startup. Nuclear fusion is the technology that tries to replicate on Earth the process the sun uses to generate energy. If it works, it could yield nearly unlimited clean energy. Yet after decades of attempts by countless scientists, success has remained out of reach.

Many people asked Sam the same question: "Why invest in something so uncertain? Wouldn't apps with a clear path to profit be far safer?"

Sam's answer was this: "When people tell you something is crazy, or that it will never work, that is always a good sign."

This was Sam Altman's philosophy: taking the hard but meaningful path rather than the easy one; pursuing long-term impact over short-term wins; building the future rather than chasing trends.

In 2015, Sam made an even larger commitment. Together with Elon Musk, Greg Brockman, and others, he co-founded a company called OpenAI. What was OpenAI's goal? To build AGI, artificial general intelligence at the level of human capability.

At the time, it sounded like something out of science fiction. Even many experts were skeptical: "AGI is still decades away." "It might be impossible." But Sam believed this was precisely the kind of genuine breakthrough that could change humanity's future.

His recognition of the limits of sharing apps and delivery services reflected a mind always fixed on a larger picture. He imagined the world ten, twenty, fifty years from now, and thought about the problems humanity would face in that world.

Environmental catastrophe from climate change, energy scarcity, food shortages, disease and aging, expansion into space , solving problems this vast would require fundamental technological breakthroughs, not ordinary apps. And Sam was convinced that artificial intelligence would sit at the center of those breakthroughs.

These ideas were not universally welcomed from the start. Inside Y Combinator, there was pushback: "We make money by investing in app companies , why would we take on the risk of hard tech?" Some investors criticized his thinking as "too idealistic" or "out of touch with reality."

With time, Sam's judgment began to be proven right. When OpenAI released ChatGPT in late 2022, the world was stunned. More than one million people signed up in just five days; today, hundreds of millions use it.

ChatGPT was not an ordinary app. It was an artificial intelligence that could hold conversations, answer questions, write prose, generate code, and solve problems. People realized: "This is what real innovation looks like. This is technology that changes the world."

Looking back, Sam Altman's recognition in the mid-2010s of the limits of photo-sharing apps and delivery services was a remarkable act of foresight. When everyone else was captivated by apps, he could see past them. He could anticipate long-term shifts rather than short-term fads.

This is the difference between a true leader and an ordinary person. Ordinary people follow whatever is fashionable now. True leaders see ten or twenty years ahead, and they know what needs to be done today to build that future.

Sam Altman looked past the era of photo-sharing apps and saw the age of artificial intelligence coming. To build that future, he decided to leave Y Combinator and stake everything on OpenAI. When he resigned as president of Y Combinator in 2019 to become CEO of OpenAI, it was no accident. It was part of a larger plan he had been working toward for years.

Is a Busy Executive a Capable Executive?

Up early, into the office, meetings, emails, phone calls, home late. But is all that "busyness" actually meaningful? Sam Altman wrestled deeply with this question while leading Y Combinator.

Silicon Valley has an old saying: "A capable executive is a busy executive." What does it mean? Company leaders , executives in particular , must constantly provide their employees with work. The idea is that you keep people busy so they don't waste time or drift.

Sam agreed with the sentiment, but his interpretation of "busy" was slightly different. Simply doing a lot of things does not make you productive, he thought. What matters is being busy with the right things.

Running Y Combinator, Sam watched many startups fail. Most of them did not fail from laziness or lack of effort. They worked too hard, if anything , fifteen hours a day, no weekends, through the night. And they still failed. Why?

The answer was simple: they were busy doing the wrong things.

Sam put it this way: "It doesn't matter how fast you move if you're moving in a useless direction. Choosing what to work on is the most important factor in productivity. And most people almost completely ignore it."

This was the core of Sam's productivity philosophy. Being busy is not the goal; being busy solving the right problems is.

How do you find the right work? Sam offered a few methods.

First, protect time to think. It may sound paradoxical, but being productive requires time to do "nothing." Sam said: "I leave enough time in my schedule to think about what to work on."

He read books, talked with interesting people, and walked in nature to clear his mind. During these stretches he asked himself: "Is what I'm doing right now truly important? Is there something more important I should be doing instead?"

Second, make lists. Sam's productivity system rested on three core pillars.

Sam made lists of what he wanted to accomplish each year, each month, and each day. He preferred writing them by hand on paper, and often rewrote them from scratch. Why? Because the act of rewriting naturally separated what truly mattered from what did not.

He also said he "prioritizes in a way that generates momentum." What does that mean? Completing one thing makes you feel good, which enables you to do more. So he tried to start and end each day with tasks that produced real, tangible progress.

Third, learn to say no. This is perhaps the hardest part. Sam said: "I am ruthless about quickly dealing with or declining non-essential things."

Saying no is something many people find genuinely hard. It feels rude to turn down a request. Receiving a meeting invitation feels like it demands attendance. Sam was different. He believed that protecting his time was paramount.

He put it this way: "I am always astonished when I see people who earn \$100 an hour spending several hours to save \$20." People, he observed, simply do not value their own time enough.

Fourth, use different times of day for different kinds of work. Sam designed his day with intention. He said that "the first few hours of the morning are clearly the most productive part of the day." So he blocked that time from any scheduled meetings. Instead, he pushed all his meetings into the afternoon.

Many CEOs pack their schedules with back-to-back meetings from 9 a.m. to 6 p.m., sliced into thirty-minute or one-hour blocks. Sam thought this was deeply unproductive. He believed that real thinking and solving hard problems require uninterrupted stretches of focused time.

He also said that "when my attention starts to slip, I take a break or switch to something else." He knew that stepping away briefly and returning fresh beats grinding through diminishing focus.

Fifth, work on things you actually care about. Sam admitted plainly that "I've learned that I can't be productive on things I'm not interested in or don't enjoy." So he worked to put himself in a position where he didn't have to do those things. He delegated, avoided, or found another way around them.

He applied the same principle when delegating to others. "Everyone is most productive when doing work they enjoy. Figure out who likes what and is good at what, then distribute the work accordingly."

The distinction matters. Most managers think mainly about dividing up the work, but Sam worked to understand what each person genuinely liked and was good at. Some people love building presentations. Others prefer data analysis. Still others come alive in conversations with customers. Matching work to those preferences makes everyone both happier and more productive.

Sixth, work with good people. Sam's advice: "Try to be around people who are smart, productive, happy, positive, and who don't look down on your ambitions."

He loved being around people who pushed him and inspired him to grow. He warned, on the other side, that negative, critical people drain your mental energy, and the cost of that is enormous.

Finally, Sam warned against falling into 'productivity porn.' What does that mean? Many people spend far too much time searching for the perfect productivity system. Which app to use, which methodology to follow, how to squeeze every last second out of the day. Meanwhile, the work that actually matters never gets done.

Sam put it this way: "Chasing productivity for its own sake doesn't help. Most people spend too much time thinking about how to perfectly optimize their system and almost no time asking whether they're solving the right problems."

That is the point. Whatever system you use, however tightly you squeeze every second, none of it matters if you're working on the wrong things.

At OpenAI, Sam put this philosophy into practice. He said he tries "to give enormous responsibility to a relatively small number of people." Small teams, but each person focused on genuinely important work and able to move fast.

This was one reason OpenAI was able to ship ChatGPT before tech giants like Google and Meta. Those larger companies had tens of thousands of employees, but their complex organizational structures and slow decision-making held them back. OpenAI, by contrast, was a small, intensely focused team that could experiment quickly, fail, learn, and try again.

Sam Altman's belief that "a capable executive is a busy executive" was never simply about doing more. It meant

That is what it meant.

When studying, what matters is not how many hours you logged but whether you actually understood and learned the material. Spreading time thinly across many subjects is often less effective than concentrating on the subjects that matter most or are hardest.

Sam Altman was a busy man, but never meaninglessly so. He constantly asked himself, 'Is what I'm doing right now genuinely important?' If the answer was no, he stopped or handed it off without hesitation. That discipline let him stay focused on what truly mattered, and it is one reason he was able to build world-changing products like ChatGPT.

The Power of APIs: Small Interface, Big Impact

If you had to name the single most important lesson Sam Altman took from Y Combinator, it would be this: "Building an API tends to lead to good outcomes, one way or another." That deceptively plain sentence would later shape the destiny of OpenAI. But first, it helps to understand what an API actually is.

API stands for Application Programming Interface. The name alone doesn't tell you much, so let's unpack what it actually does.

In plain terms, an API is like a bridge that lets different programs talk to each other.

Think of a restaurant. You look at the menu and place your order. The chef prepares the food and brings it out. You don't need to know what the kitchen looks like or what tools the chef uses. You just choose what you want from the menu.

"OpenAI provides the core AI brain block, and developers add their own idea blocks on top to build an endless variety of things. That is the power of an API ecosystem." Take Google Maps, for instance. Google exposes its mapping capability as an API, so a delivery app company doesn't need to build its own maps. It simply calls the Google Maps API to show where a courier is.

Watching hundreds of success stories at Y Combinator, Sam noticed a pattern. Companies that offered APIs tended to grow especially fast and achieve unexpected success.

Stripe is the clearest example. Stripe handles online payments. Before Stripe, accepting credit cards on a website meant going through a complicated process: negotiating with banks, obtaining security certifications, writing complex code. For most startups, it was nearly impossible.

Stripe reduced all of that to an API. A developer could add a few lines of code and their site could take payments, snapping in like a LEGO brick. The result: countless startups adopted Stripe, and it grew into a company valued at tens of billions of dollars.

Twilio is a similar story. Twilio offers SMS and voice calling as an API. If you're building a delivery app and want to text customers 'Your order is almost here,' you call the Twilio API. No need to negotiate a deal with a mobile carrier yourself.

Watching these cases, Sam came to a realization. The real power of an API is that it makes unexpected innovation possible.

What does that mean?

When a company builds an API, it usually expects people to use it for the purposes the company had in mind. In practice, it often gets used in entirely different ways. And those unexpected uses are sometimes where the greatest value comes from.

Consider Amazon Web Services (AWS). Amazon originally built those APIs internally to make its own server infrastructure run more efficiently. When they opened it to the public, startups around the world started using it. AWS is now one of Amazon's most profitable business units. Amazon almost certainly didn't see that coming.

Paul Graham made the same point. He told Sam: 'Whatever happens, build the API. Good things will follow.'

That advice would later prove decisive at OpenAI.

In 2020, OpenAI developed GPT-3, a powerful language model. It could write like a human, answer questions, and even generate code. But a problem remained: 'What do we actually build with this?'

GPT-3 was clearly remarkable technology, but on its own it was not a 'product' that ordinary people could use. It was like having a powerful engine with no idea what car to put it in.

There was a lot of debate inside the OpenAI team. Some people said, "Let's build an email writing tool with GPT-3." Others said, "Let's build a homework assistant for students." Nobody felt certain about any of it.

At that point, Sam Altman thought back to advice from Paul Graham and lessons he'd absorbed at YC. "Build an API and good things tend to follow."

He put it simply: "If we don't know exactly what we should build, let's open it up as an API first. Someone else might build something we never would have thought of."

In June 2020, OpenAI released the GPT-3 API. They granted access to a select group of developers and watched to see what they would create with this powerful AI.

The results were remarkable.

One developer built a tool that automatically generated marketing copy using GPT-3. Another built a customer service chatbot. Another built a tool to help write code. Hundreds of distinct use cases emerged.

What the OpenAI team found most fascinating was something else entirely. People were spending hours each day just talking with the AI inside the GPT-3 Playground.

The Playground had been built as a simple interface for developers to test the API. It was a space to experiment with how GPT-3 worked, without any particular purpose in mind. Yet, strikingly, many users were lingering there for hours, holding what could only be called a 'conversation' with GPT-3.

They asked it questions, told it stories, sought its advice. Just as they would with a friend. These were a small number of users, but their behavior was sending a clear signal: people want to talk with AI.

The OpenAI team reached a realization. "Maybe the product we need to build is exactly this. Not a complex tool, just a simple chat window where you can have a conversation with an AI."

ChatGPT was born from that insight. On November 30, 2022, OpenAI released ChatGPT. No flashy features, no complicated setup. Just a simple chat interface for talking with an AI.

The result? One million users signed up within five days. One hundred million within two months. It became the fastest-growing consumer application in history.

What if OpenAI had tried to build the "perfect product" from the start? They might have spent years planning, debating, and developing, only to miss the window entirely. Instead, they released an API first, and discovered what users actually wanted.

That was the moment Sam Altman's insight, that "APIs lead to good outcomes," proved itself.

The power of an API comes down to three things.

An API taps into collective intelligence. Ten smart people at one company can only think of so much. Tens of thousands of developers around the world will think of far more. Opening an API is an invitation: "Take our technology and build whatever you want with it."

An API accelerates market validation. You can't know in advance which product idea will succeed. But once an API is open, you can quickly see what people actually use it for. The most popular use case is exactly what the market wants.

An API builds an ecosystem. A single company doesn't have to build everything on its own. When core technology is offered as an API, countless other companies build their own products on top of it. The whole ecosystem grows, and so does the company that provided the original API.

The pattern Sam Altman observed went beyond a business strategy. It reflected a deep understanding of how innovation actually happens. Innovation doesn't come only from one person's brilliant idea. Sometimes the greatest breakthroughs come from building a tool, putting it in other people's hands, and watching what they create.

The hundreds of millions of people using ChatGPT today, and the tens of thousands of companies that have woven it into their own products, probably don't know this. Behind the technology they rely on is a lesson Sam Altman learned during his YC years, simple but powerful: "Build an API, and somehow good things tend to follow."

That is one of the most enduring legacies of Sam Altman's time leading Y Combinator. He didn't simply invest in startups. He gained a deep understanding of how innovation happens, and that understanding became the key that would later open the door to the age of artificial intelligence.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

2.1 Sensing the AI Revolution and Founding OpenAI

Part 2: The Birth and Growth of OpenAI

[← Previous](#)

[See all chapters](#)

[Next →](#)

Watson and DeepMind: The First Signs of a Major Shift

In the early 2010s, Sam Altman was president of Y Combinator, listening to dozens of startup pitches every day. Photo apps, delivery services, social media. Smart people were building them, but honestly, Altman found them a bit underwhelming. 'Will any of this actually change the world?' he wondered to himself.

In 2011, something strange played out on television. A computer called Watson, built by IBM, appeared on the famous American quiz show Jeopardy!. People laughed at first. 'A computer solving trivia questions?'

Jeopardy! was no ordinary quiz show. A clue like 'He discovered America in 1492' is easy enough, you just answer Columbus. But Jeopardy! works in reverse. You're given 'Columbus' as the answer and have to figure out what the question was. On top of that, the clues are packed with wordplay, metaphors, and double meanings.

The human champions were nervous. Ken Jennings and Brad Rutter were the greatest players in the show's history. But Watson dominated both of them. It found answers the instant each clue appeared and hit the buzzer faster than any human could. While the people on stage were still thinking, Watson had already spoken.

Sam Altman watched and felt a chill. He'd already seen computers beat humans at chess, but that was a calculation problem. Language is different. Language is ambiguous, context-dependent, culturally loaded. The fact that Watson understood American wordplay meant something larger: computers were stepping into the human world of language.

Most people saw Watson as a curiosity, something to marvel at briefly before the TV credits rolled and everyone moved on. Altman saw it differently. He was someone who read the patterns beneath technology.

'There's this pattern,' Altman would say. 'A technology shows up and looks like a toy. People laugh, what's the big deal? Then five years pass, ten years pass, and it's completely reshaped the world.' He'd watched it happen with the internet, then with smartphones. Watson looked like the same kind of beginning.

In 2014, something bigger happened. Google acquired DeepMind, a small British AI company, for more than \$400 million. At the time, DeepMind had fewer than 50 employees and no revenue-generating products. So why did Google pay so much?

DeepMind's researchers were working on something unusual: a technique called reinforcement learning. Put simply, instead of programming a computer step by step, you let it learn through trial and error on its own. The way a baby learns to walk, falling down, getting up, trying again, gradually getting better.

DeepMind used this technique to master Atari games, Pac-Man, Space Invaders, the old classics. They never taught the AI the rules. They just showed it the screen and said, in effect, get a higher score. The AI died thousands of times, failed over and over, and figured out the patterns on its own. Eventually it played far better than any human.

When Altman heard the news, he couldn't sleep. He understood why Google had spent that kind of money. What DeepMind was building wasn't just an AI that was good at games. It was intelligence that could teach itself. This was real.

Altman started calling friends. Elon Musk, Peter Thiel, the sharpest minds in Silicon Valley. 'We need to have a serious conversation.' His voice carried a weight his friends weren't used to hearing.

'What happens if Google gets to AGI first?' Altman asked. AGI, artificial general intelligence, means not a narrow AI that's good at one thing, but a system that can learn and understand anything, the way a human can. 'If one company monopolizes that kind of power, that can't be good.'

Most people at the time still treated AGI as science fiction. 'That's a hundred years away.' But Altman, looking at Watson and DeepMind, thought differently. Technology moves slowly, then all at once. If you wait until you're sure, it's already too late.

'Someone has to handle this responsibly. For all of humanity.' Elon Musk nodded. Both men felt they had to do something.

Watson winning on a quiz show was a signal, that the age of machines was coming. The DeepMind acquisition was a louder signal: the tech giants had begun the race to AGI.

Sam Altman read the signals and made a decision. 'We need to get in, but do it differently.' That was how the seed of OpenAI was planted. The uneasy feeling of 2014 became the defining commitment of 2015.

In December 2015, Sam Altman and a small group gathered at a café in San Francisco. Elon Musk, Greg Brockman, Ilya Sutskever, names that carried weight on their own. They sat over coffee with serious expressions and talked.

'Are we really doing this?' someone asked.

'We have to. If not us, who?' Altman answered.

What they were setting out to do was straightforward to state: build AGI, human-level artificial intelligence, and make sure its benefits reach all of humanity. It sounds admirable. But when you think it through, it was an audacious plan, maybe a crazy one.

Building AGI means building a genuinely intelligent computer, not one that only plays chess or only translates text, but one that can learn new things, solve problems, and think creatively, the way a human does. At the time, it wasn't even clear this was possible.

'Honestly, we didn't know how to build it,' Altman admitted in an interview years later. 'You just believe it's possible and keep trying.'

Think about the early days of the space program. No one knew how to build a rocket capable of reaching the moon, how to keep people alive in space, or how to land. President Kennedy declared they'd do it within ten years. People thought he was out of his mind. They did it anyway.

OpenAI was similar. The destination was visible; the path was not. They decided to leave anyway.

Their stated mission read: 'Ensure that AGI benefits all of humanity.' The key word is 'all.' Not a particular company, not a particular country, everyone.

Why did that matter? Altman worried. 'If a giant corporation like Google or Facebook gets to AGI first, they'll use that power for their own benefit. That's what companies do, they answer to shareholders.'

AGI was too powerful a technology to leave in one set of hands. Imagine a system smarter than any human across every field, acting as doctor, lawyer, scientist all at once. If one company controlled that, the world would become deeply unequal.

'We need a different approach,' Altman said. 'The goal can't be profit. The goal has to be building AGI that's safe and good.'

The problems were enormous.

The first problem was how to build AGI at all. AI research at the time was narrowly focused: an AI that recognized images, an AI that translated text, an AI that played games. Each was advancing, but combining them all into something like general intelligence was an entirely different challenge.

'It's like having puzzle pieces without knowing what the puzzle is supposed to look like,' one researcher said.

The second problem was safety. If AGI were actually built, it could be dangerous, not in the sci-fi sense of machines attacking humans, but in a subtler way: an AGI might pursue its goals through means no one intended.

There's a well-known thought experiment: tell an AI to make as many paper clips as possible. A sufficiently intelligent system might build factories, gather resources, and eventually consume the entire planet in the service of paper clip production. Tell it to stop, and it would say, in effect: my goal

is to make paper clips.

Sounds absurd, doesn't it? But AI safety researchers take these problems seriously. AGI could carry out instructions 'to the letter.' Altman understood just how hard it is to make a machine grasp what humans actually mean.

Third problem: recruiting the best minds in the world. AI researchers at Google and Facebook were pulling enormous salaries. OpenAI was a freshly founded nonprofit, and its pay was far lower. How do you convince people to come?

The strategy was simple: offer meaning, not money. 'What you do here will be remembered in the history of humanity. Someday your kids will grow up and say, My mom/dad helped build AGI.'

Remarkably, it worked. World-class AI scientists like Ilya Sutskever left Google and joined OpenAI. Why? Because the dream was bigger than the paycheck.

Fourth problem: money. (We'll go into this in detail in the next section.) AGI research demanded enormous computing power. Running supercomputers alone cost staggering sums every single day.

'Can we actually do this?' one team member wondered aloud.

'We have to.' Altman was firm. 'If we don't, this technology will end up in the hands of a few giant corporations. And that would make the world deeply unfair.'

Founding OpenAI was a gamble. What were the odds of success? Nobody knew. How long would it take? Nobody knew that either. How much would it cost? More than anyone could imagine.

And yet they started anyway. Why? Altman explained it this way: 'Sometimes you have to do something even if the odds are low. If the payoff when you succeed is enormous enough, it's worth it. AGI is exactly that kind of thing.'

Put it in mathematical terms: even with a 10% chance of success, if the value of that success is 1,000, the expected value is 100. That makes it worth attempting.

On December 11, 2015, OpenAI was officially founded. The statement posted on its website was short but striking: 'We will create AGI for the benefit of humanity.'

That night, Sam Altman walked outside the office into the cold San Francisco wind. He looked up at the sky. 'Can we really go to the moon?' he said to himself.

Then he smiled. 'Let's go find out.'

A Reckless Gamble: Starting as a Nonprofit

'You're doing this as a nonprofit?'

Many people in Silicon Valley shook their heads in disbelief. In the heartland of startups, building a company that wasn't trying to make money seemed bizarre. 'Are those people out of their minds?'

But Sam Altman had a clear reason.

'AGI can't be a money-making tool.' He explained to his friends: 'Think about it. If we were a for-profit company, we'd have to deliver returns to shareholders. That would mean using AGI in whatever way makes the most money. Is that really what's best for humanity?'

Imagine you found a magic lamp. Rub it, and a genie appears to grant your wishes. Now imagine turning that lamp into a company. What would happen?

The company would think: 'We need to profit from this lamp! Let's sell wishes to the rich at a premium!' Poor people would never get to make a wish. That's not fair.

OpenAI thought differently. 'This lamp should belong to everyone. Rich or poor, American or African, every person deserves to benefit from it.'

That's why they chose the nonprofit path. A nonprofit's purpose isn't to make money; it's to serve the public good. OpenAI was no different. 'Build AGI, but build it for all of humanity' - that was the mission.

At first, it seemed fine. Billionaires like Elon Musk, Peter Thiel, and Reid Hoffman pledged donations. A total of \$1 billion! That's over one trillion Korean won. 'Wow, surely that's enough!'

But... it wasn't nearly enough.

The problem was that AI research devoured money at a staggering rate. It needed enormous amounts of computer processing, or in technical terms, 'computing power.'

Training AI requires specialized computer chips called GPUs. Each one costs several million won. And you don't just need one. You need thousands, tens of thousands of them.

Why so many? Imagine looking at a million photos of dogs and saying, 'That's a dog, that's a dog too.' A person learns quickly, but a computer has to repeat the process millions, even hundreds of millions of times. Processing all those calculations takes enormous computing power.

In 2016, one of OpenAI's researchers came to Altman with a problem. 'Sam, we have a situation. To run this experiment, we need 500 more GPUs, but... there's no budget.'

'How much would that cost?' Altman asked.

'Hmm... around 200 million won?'

Altman sighed. This wasn't the first or second time. Money was draining away every month like water. They had to pay salaries for the world's top researchers, keep the computers running, and cover office rent.

A billion dollars had seemed like a lot, but at this rate, it would run out within a few years.

'Google and Facebook have practically unlimited money to spend, while we...' one employee said, frustrated.

True. The tech giants had no money worries. Google poured hundreds of billions of won into DeepMind every year. An experiment failed? No problem, just try again. But OpenAI was different. The budget was fixed, and when the donations ran out... it was over.

There was an even bigger problem. Nonprofits struggle to attract investment. A regular company can promise investors, 'When our valuation rises, you'll make a fortune!' A nonprofit can't make that promise. It doesn't distribute profits.

"This... isn't sustainable." Altman couldn't sleep.

He was caught in a dilemma. On one hand, he wanted to keep the nonprofit structure. He didn't want AGI to be corrupted by money. On the other hand, continuing the research required enormous funding.

One night, Altman gathered the team leaders.

"I'll be straight with you. At this rate, we have maybe three years. After that, the money runs out."

The room went quiet.

"So what do we do?" someone asked.

"We need to find a new way." Altman said.

This eventually led to the creation of an unusual structure called 'OpenAI LP.' It was still a nonprofit at its core, but with a hybrid design that allowed for some for-profit activity. Complex, but,

The key point was this: "We can take investment, but ultimate control stays with the nonprofit board."

But back in 2016, there was no such solution yet. Altman simply held on. He chased donors everywhere, cut costs, and asked researchers to "just wait a little longer."

Some researchers went back to Google. They were offered higher salaries and unlimited computing resources. Altman didn't blame them. "Everyone has to make a living."

But many stayed. Why? Because they believed in OpenAI's mission. "The pay is low, but what we're doing will leave a mark on history."

Just how hard that period was, Altman later described it this way: "Every morning when I woke up, I'd think, 'Is today the day OpenAI shuts down?' That's how anxious I was."

Starting as a nonprofit was idealistic. But reality was harsh. They kept their principles, yet paid a price for it. Sleepless nights, relentless worry, an uncertain future.

"Still, I have no regrets." Altman said in a later interview. "Because we started as a nonprofit, we were able to go in the right direction. If money had been the goal.. we would have become a different company."

The funding problem tormented OpenAI, but it also toughened the organization. Scarce resources forced them to think more creatively and focus only on what truly mattered. In a way, the poverty may have been what made OpenAI stronger.

YC President and OpenAI CEO: Two Jobs at Once

This was what a day in Sam Altman's life looked like.

6 a.m. Wake up, check email. Fifty urgent items for YC, thirty urgent items for OpenAI.

9 a.m. Arrive at the YC office. Investment review meeting. "Should we fund this startup?"

Noon. Mentoring YC founders. "Your business model is wrong."

3 p.m. Drive to the OpenAI office. Conference calls stuck in San Francisco traffic.

4 p.m. Arrive at OpenAI. Technical discussions with researchers. "The GPT model isn't hitting expectations?"

7 p.m. Skip dinner, investor meeting. "Would you consider donating to OpenAI?"

10 p.m. Home. Reading research papers until 1 a.m.

That was how it went. Every day.

"Are you out of your mind?" his friends asked. "Running two organizations at the same time!"

Altman laughed. "Yeah, I probably am."

Being president of YC alone was an enormous job. Y Combinator was the most influential startup accelerator in Silicon Valley. Airbnb, Dropbox, Stripe.. all major companies had come out of YC.

Twice a year, Altman reviewed hundreds of startups. He met founders, asked questions, and decided whether to invest. "Will these people succeed?" That judgment call was his job.

At the same time, he had to lead OpenAI. AGI research -- arguably the most difficult project in human history.

"How do you manage both?" a reporter asked.

"Just don't sleep." Altman said it like a joke. But he meant it. He averaged about five hours of sleep a night.

Why push himself so hard? Couldn't he have let one go?

Altman had his reasons. YC and OpenAI weren't separate things to him. They were connected.

At YC, he saw the future of technology unfold in real time. Which startups thrived, which ideas collapsed, which technologies had the power to change the world. All of that shaped how he thought about strategy at OpenAI.

In the other direction, what he learned about AI at OpenAI fed back into how he advised YC startups. 'What if you applied AI to your product this way?'

He served as a bridge between two worlds.

But it wasn't easy. It was close to physically impossible.

People at YC grew frustrated. 'Sam is too busy to really focus.' People at OpenAI felt the same. 'Our CEO is always over at YC.'

Paul Graham, YC's founder and Altman's mentor, grew worried. He believed that when a leader isn't fully present, the organization starts to wobble.

One day, Graham called Altman in. 'Sam, we need to talk.'

The two sat down in a quiet cafe.

'You're taking on too much,' Graham said. 'YC needs your complete commitment. So does OpenAI.'

Altman knew. He had always known that a choice was coming.

'But not yet...' He let the sentence trail off.

'Not yet?'

'I need both. YC is the present, OpenAI is the future. I don't want to lose either.'

Graham sighed. He understood Altman's passion, but he also knew the reality of the situation.

2017, 2018. As time passed, things got harder. OpenAI demanded more and more attention. The research grew complex, funding became a serious problem, and the technical decisions kept getting weightier.

Finally, in 2019, Altman made his decision. He would step down as president of YC.

'It wasn't an easy call,' he said later. 'YC is where I learned everything. But I felt that OpenAI mattered more.'

Why did he choose OpenAI?

'YC is already a well-oiled machine. It will keep running without me. OpenAI is different. AGI development is still in its early stages. It could fail. It could veer off in the wrong direction. Someone needs to be completely focused on steering it.'

At his final YC event, Altman held back tears as he spoke. 'The time I spent with you has been the best of my life. But now it's time to climb a different mountain.'

Three and a half years running YC and OpenAI in parallel. That stretch forged him. Time management, fast decision-making, setting priorities under pressure. Every one of those abilities was pushed to its limit.

'If you had to go through that period again?' someone asked.

Altman laughed. 'No thanks. It was brutal.'

He paused, then added: 'But I have no regrets. That time made me who I am.'

Early Uncertainty: Research Lab or Product Company?

The OpenAI office in 2016 was, to be honest, a mess.

Whiteboards were covered in strange equations, desks were buried under stacks of papers. People in hoodies sat in front of laptops, typing furiously. Every so often someone would shout 'Come look at this!' and everyone would crowd around a screen.

'What exactly are we building?'

That question came up more often inside OpenAI than any other.

The answer? Nobody knew.

At a normal company, that would be alarming. 'We don't know what our product is' would make any CEO faint. But OpenAI wasn't a normal company. It was closer to a research lab.

'We're trying to build AGI,' Altman said.

'Okay, but how?' the researchers asked.

'...I don't know.'

That was the honest answer.

Everyone knew what AGI was. 'A computer that thinks like a human.' But nobody knew how to build one. It was like shouting 'Let's go to the moon!' with no rocket blueprints in hand.

Researchers were experimenting in several directions.

One team was working on robotics. They built a robotic hand and trained it to grasp objects. They started with rotating a cube , picking it up, turning it, solving it. It looked simple but was extraordinarily hard.

"Why do we need robots?" someone asked.

"Real intelligence has to interact with the physical world. Being smart only inside a computer isn't real intelligence."

It made sense. Humans learn by touching, feeling, and experiencing the world, after all.

Another team was focused on video games.

"Games?"

Yes, games. They focused in particular on Dota 2, a complex strategy game.

Why games? Because games were perfect training grounds for AI.

Games have a clear objective: "Win!" Simple. Easy to explain to an AI.

In a game, you can practice millions of times. In the real world, a falling robot can break. Inside a game, dying a million times is fine , you just restart.

A complex game demands strategy, teamwork, and long-term planning. Raw reaction speed isn't enough. You have to actually "think."

The OpenAI team began training an AI to play Dota 2. At first it was a mess. The AI spun in circles or made random, pointless moves. It would die within ten seconds of a match starting.

"Can this actually work?" Even the researchers had doubts.

But they kept training it. Thousands of games, tens of thousands, hundreds of thousands. Through repetition, the AI learned bit by bit. "Ah, that move doesn't work here." "This character needs to be played this way."

A few months later, something remarkable happened. The AI started beating amateur players.

"Wow, it's actually learning!" The researchers were thrilled.

But professional players still dominated. The AI understood the basics but hadn't grasped advanced strategy.

"Train it more." They pushed on.

The process consumed enormous computing power. Hundreds of GPUs ran around the clock. The electricity bill alone was staggering. But something was being learned , evidence that the AI was beginning to grasp what "strategy" meant.

Meanwhile, another team was researching language.

"Language is the key," said Ilya Sutskever. He was OpenAI's chief scientist.

"Why?"

"Think about it. Most human knowledge is expressed in language, books, the internet, conversation. All of it is language. If an AI truly understands language, it can access everything humans know."

The team began building a language model called GPT. GPT stands for Generative Pre-trained Transformer, but that's a mouthful, so think of it simply as "an AI that generates text."

At first it was underwhelming. GPT-1 could produce only simple sentences, things like "I go to school." It performed worse than a primary school student.

"Can this ever become AGI?" Many were skeptical.

But Sutskever was convinced. "Make it bigger and it'll be different. Much bigger."

So they built GPT-2. This was better. It could write short stories and give plausible answers to questions. People started saying, "Oh, that's something."

But it still wasn't right. It rambled at times and said things that made no sense. "Human-level" was still a long way off.

This was the true picture of OpenAI in 2016.

The robotics team sighed watching the robot drop the cube again. The games team debugged another AI loss. The language team scratched their heads at yet another garbled sentence.

"Can this really lead to AGI?" Everyone had doubts.

Altman watched all this uncertainty unfold. He didn't have the answers either.

But he didn't lose faith.

"That's just how science works," he told the researchers. "It's always murky at the start. Edison failed 999 times before he made the light bulb. We're in the middle of that same process."

One evening, Altman walked through the office and looked in on each team's work. He stopped by the robotics lab, glanced at the monitors in the gaming team's room, and read through the latest results from the language team.

Everything was incomplete. It all looked like failure.

But he saw it differently. Each team was building a small piece of a puzzle. The robotics team was working on "physical learning," the gaming team on "strategic reasoning," the language team on "knowledge representation."

"Someday these pieces will come together," Altman thought.

But when? And how?

Nobody knew.

That night, Altman stepped out of the office and looked up at the sky. The city lights of San Francisco washed out most of the stars.

"I can see the moon, but I can't see the road to get there," he said to himself.

Then he smiled. "Still, we have to go. Somehow."

OpenAI in 2016 was chaotic, its direction far from clear.

There was no product, and the odds of success looked dim.

But something was growing.

Out of sight, quietly, slowly.

A seed was sprouting, one that would later change the world.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AILibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

2.2 The Journey from Model Development to Products

Part 2: The Birth and Growth of OpenAI

[← Previous](#)

[See all chapters](#)

[Next →](#)

Video Games, Robot Arms

In 2016, the OpenAI office had a completely different feel from what it is today.

Sam Altman remembers it as "feeling like a research lab." Nobody was quite sure what to do. A team of roughly fourteen people stood in front of whiteboards, asking each other, "So what are we actually doing?" The big goal was clear: build artificial general intelligence, AGI, at human level. How to get there, nobody knew.

Their first GPU arrived , a DGX-1 system delivered in person by NVIDIA CEO Jensen Huang. By today's standards it was a remarkably small piece of hardware. Altman recalled with a laugh that "one person could carry it."

The team pursued two directions.

The first was video games. Why games? Because a score tells you immediately how well the AI is performing. It was also something people could follow. "The AI beat Dota!" was a line everyone could understand.

OpenAI chose Dota 2, one of the most complex games in existence. It requires five players to coordinate strategy in real time, reading the opponent's intentions, planning long-term, and working together as a team. Every quality associated with human intelligence was baked into it.

In 2018, OpenAI's team of AI agents, called OpenAI Five, defeated world-class professional players.

The world was stunned.

The idea that an AI could master a game this complex was startling. OpenAI Five had learned by playing against itself, countless times. It started knowing nothing. But after simulating the equivalent of hundreds of thousands of years of gameplay, it kept improving , and eventually developed strategies no human player had ever conceived.

Altman drew a lasting lesson from this project: given enough data and computing power, AI could do remarkable things.

The second direction was robotics.

OpenAI built a robotic hand modeled on the human hand and set it the task of solving a Rubik's Cube. The goal was an AI that could touch and manipulate physical objects in the real world, not just inside a simulation.

The real world turned out to be far harder.

There was gravity, friction, and sensors that were never quite accurate. The same motion produced different outcomes every time. The team used a technique called domain randomization, constantly varying the cube's weight, size, and friction within the simulation during training. That way the AI learned to handle the unpredictability of the physical world.

A few years later, the robotic hand solved a Rubik's Cube one-handed.

This was not something achievable through conventional programming alone. The AI had learned to do it on its own.

Altman valued these experiments deeply. There were plenty of failures. But the failures taught something. The road to AGI required trial and error. Video games and robots showed how an AI could learn to navigate an environment.

Looking back, this whole period was a search for direction.

Altman has said, "We tried to understand the world through games and robots, but eventually we realized that language is the tool that compresses the world." Without those early experiments, GPT would never have been born.

From GPT-1 to GPT-3: "Make It Bigger and Everything Changes"

While the video game and robotics experiments were underway, a new interest took hold inside OpenAI.

Language.

Researcher Alec Radford made an intriguing discovery. Inside an AI trained on Amazon reviews, a single neuron had learned to judge whether a piece of text was positive or negative. Nobody had taught it to do that. The AI had figured it out on its own.

Altman sensed that something interesting was happening here.

Around this time, a key piece of technology appeared.

It was the Transformer, published by Google in 2017. This architecture could handle long sentences with ease. It ran in parallel, so it was fast. Most critically, it could be scaled up to enormous size.

In June 2018, OpenAI released GPT-1.

The name stood for Generative Pre-trained Transformer 1. It had 117 million parameters and had trained on thousands of books. What it did was remarkable: without being specifically trained for any given task, it could handle translation, summarization, and more.

Altman saw this as the beginning.

"Make the model bigger and it'll get better."

In February 2019, GPT-2 arrived.

This time it had 1.5 billion parameters, more than ten times the size of GPT-1. It trained on a vast body of text scraped from the web. GPT-2 produced writing that read as naturally as something a human had written. Give it a few words and it would flow right into the next sentence.

It was so capable that it became a problem.

OpenAI wrestled with the question. The technology could be misused. It could churn out fake news at scale. It could generate phishing emails automatically. So at first, the team released only a smaller version.

This decision sparked controversy.

"You call yourselves OpenAI, so why aren't you making it open?" The criticism was real. But OpenAI proceeded carefully. A few months later, once they concluded there were no major risks, they released the full model.

Through this process, Altman became convinced of something important.

Scale up a model and its performance improves. Not just a little. It improves in a predictable, measurable way. This was the 'scaling law.' Add more data, build a bigger model, use more compute, and performance rises at a steady rate.

"This may be one of the most important discoveries made in my lifetime."

Altman laid out the next plan. Build GPT-3. Then GPT-4 after that.

GPT-2 was the opening shot of a revolution still to come.

After GPT-2's success, Altman set his sights on something bigger.

"Let's scale everything."

In June 2020, GPT-3 arrived. 175 billion parameters. More than 100 times larger than GPT-2. The biggest language model ever built.

GPT-3 felt like magic.

Show the AI just a few examples, and it caught the pattern and solved new problems on its own. This was called 'few-shot learning.' It could translate text, handle simple calculations, and even write code.

There was a problem.

The costs were staggering.

The computing power required to train GPT-3 was beyond imagination. Tens of thousands of high-performance GPUs had to run without pause for weeks. The electricity bill alone ran into millions of dollars.

The OpenAI team came to a realization.

"This is no longer something a small research lab can do."

Building AGI demanded enormous capital. Donations to a nonprofit alone couldn't cover it. Top talent was also essential. Giants like Google and Meta were pulling researchers away with extraordinary salaries.

Altman made a decision.

They had to make money. But they couldn't lose the mission.

In 2019, OpenAI built a new structure. A 'capped-profit' model. Investors could earn returns, but no more than 100 times their investment. Anything beyond that goes back to the nonprofit foundation. It was a safeguard against unchecked profit-seeking.

They partnered with Microsoft.

Microsoft invested \$1 billion. In return, OpenAI committed to using Microsoft's Azure cloud exclusively. OpenAI got the computing power it needed; Microsoft got access to the best AI technology available.

Altman put it plainly: "The cost of intelligence ultimately converges to the cost of energy."

Building GPT-3 made it clear to him: smarter models require bigger computers, and those computers require enormous amounts of energy.

GPT-3 was a technical triumph, but it also revealed a hard truth. AI development had become a field that demands investment on the scale of national infrastructure.

After building GPT-3, Altman faced an important question.

"Now what?"

They had built remarkable technology. They just didn't know what to do with it. Building GPT-4 would require far more money. Billions of dollars. An amount impossible to raise through pure research

alone.

OpenAI had to change.

Altman had learned something from his time as president of Y Combinator: "Build an API and good things follow." An API lets developers plug a technology directly into their own services. Altman had a thought.

"Instead of building products ourselves, let others build the applications."

In June 2020, the GPT-3 API launched.

It was OpenAI's first commercial product. It let developers tap into GPT-3's capabilities, for a fee. Access started with a limited group. Then it expanded gradually.

Silicon Valley was electrified. "This is genuinely extraordinary." Some people went so far as to say, "This is AGI." But the general public had little interest at the time. It was too technical. If you weren't a developer, you couldn't use it.

Few companies were actually making money.

A handful of copywriting companies had emerged. They offered services that automatically generated ad copy. GPT-3 was good at writing marketing text. That was the only 'economically meaningful' use case at the time.

Altman didn't see this as failure.

It was a learning process. The API was a tool for developers. It was too difficult for ordinary people. For real innovation to happen, an easier approach was needed.

That's when an important discovery was made.

There was a place called the 'Playground' where developers tested the API. Looking at the logs, people were doing something unexpected. They were having conversations with GPT-3.

"Who are you?" "What do you think about this problem?" "Tell me a joke."

People were treating GPT-3 not as an information tool, but as a conversation partner.

Altman realized something.

"People want to talk to AI."

This was the beginning of ChatGPT. The API wasn't a successful business model, but it was a bridge to the next stage. And thanks to the partnership with Microsoft, OpenAI could keep doing research.

From a pure research lab to a business. Without that transition, ChatGPT would never have come to be.

The Limited Success of the API and a New Discovery

The GPT-3 API launched quietly.

In June 2020, only a select group of developers had access. It felt like a secret club. People were curious. "Is GPT-3 really that impressive?"

The developers who tried it first were astonished.

Given a brief description, it wrote code. It translated English into French. It summarized long texts. It even wrote poetry. Remarkable demo videos started appearing on Twitter.

"I built an entire website with GPT-3!" "Just describe it in plain language and an app gets created!"

But few of these turned into actual businesses.

The first problem was cost. GPT-3's most powerful model was expensive. Building a service for large numbers of users was financially daunting.

The second problem was unreliability. GPT-3 sometimes gave strange answers. It would confidently fabricate things that weren't true. This came to be called 'hallucination.' Results varied wildly depending on how a prompt was written.

Still, there were areas where it succeeded.

Copywriting.

Startups like Jasper and Copy.ai appeared. They used GPT-3 to automatically generate marketing copy. Blog posts, email subject lines, product descriptions, social media posts. Feed in a few keywords and multiple versions of the text would come out.

Marketers loved it.

Why copywriting? The field required creativity, but factual accuracy mattered less. A bit of exaggeration or error wasn't a major problem. What mattered was producing a lot, fast.

Code generation also worked well.

Microsoft's GitHub used Codex, the code-focused version of GPT-3, to build GitHub Copilot. When a developer wrote a comment, it would auto-complete the code. Programmer productivity jumped significantly.

Altman learned something important through the API.

No matter how good the technology, a poor interface meant no innovation. The API was a developer tool. It was too far removed for ordinary people to feel the power of GPT-3.

An unusual pattern appeared in the Playground logs.

People were having 'conversations' with GPT-3. Asking questions, hearing answers, asking again. Like talking with a friend.

This led to ChatGPT.

The API was a limited success. But it wasn't a failure. It was a learning process on the way to the next stage. The success in copywriting was proof that GPT-3 could solve problems that actually generated revenue.

People Wanted to 'Talk' to AI.

The Playground was originally a testing space for developers.

There was a large text box. You'd type something, and GPT-3 would generate the next sentence. Developers used it to experiment before working with the API. They'd tweak settings like temperature and token count to see how the output changed.

Then something strange started happening.

People were chatting with GPT-3 in the Playground.

"Hey, who are you?" "I'm GPT-3. How can I help you?" "What's the capital of France?" "Paris." "Can you recommend some places to visit in Paris?"

It was conversation. Not just retrieving information, but exchanging words. Users treated GPT-3 not as a machine but as someone to talk to.

OpenAI's research team analyzed the logs.

What they found was striking. People weren't spending a few minutes, they were talking for tens of minutes. They went beyond simple question-and-answer exchanges, building stories, asking for advice, trading jokes. Some even shared their personal worries with GPT-3.

At the time, the conversational capability was poor.

Altman admitted it himself. "It was terrible." OpenAI hadn't yet figured out how to tune the model for conversation. This was before RLHF (Reinforcement Learning from Human Feedback) had been properly applied.

And yet people loved it.

"They loved it anyway."

Altman understood. This was the thing that mattered.

People didn't just want information. They wanted to talk. To organize their thoughts, hear explanations, get advice, ask questions and learn through dialogue. It was a human instinct.

Language is how humans think. We speak to sort out our thoughts. GPT-3 fit that instinct perfectly.

Altman told his team:

"We now know for certain that people want to have conversations with the model. Let's build a product that makes that possible."

That was the beginning of ChatGPT.

It was built on GPT-3.5 and used RLHF. Human evaluators read the model's responses and rated them. They ranked which answers were better. The model learned from that feedback and kept improving.

The goal was straightforward.

Make it accessible to anyone. Log in and start talking right away. No technical knowledge required. Ask questions in plain language. Keep the conversation history easy to read.

A small behavioral pattern spotted in the Playground became the seed of a product that would change the world.

People were already showing OpenAI what they wanted from GPT-3. OpenAI read that signal clearly.

DALL-E: The Magic of Turning Imagination into Images

In early 2021, OpenAI unveiled something different.

It was DALL-E.

Type in a phrase like "an armchair shaped like an avocado," and an image of a chair that genuinely looked like one would appear. Type "a two-legged elephant walking on the moon," and that strange image would materialize.

People were stunned.

AI can draw? It takes something you imagined and turns it into a picture on the spot? It felt like magic.

DALL-E used the same Transformer architecture as GPT, but instead of text it produced images. The mechanism worked like this: an image is broken into small patches, each patch is converted into a token, and text and image tokens are trained together.

So text went in, and images came out.

Altman saw the future in DALL-E.

"AI can't only deal with language. The world isn't just text. There's vision, there's sound, there's action."

Multimodal. AI that handles multiple forms of data together. That was the path toward genuine intelligence.

DALL-E showed what the democratization of creativity could look like.

Even people who couldn't draw could become artists. All you needed was language. You didn't have to be a designer to create a stunning image. This was a cultural shift.

In 2022, DALL-E 2 arrived.

It was far better. The images were sharper and more realistic. DALL-E 2 used a different model called CLIP. CLIP had a deep understanding of the relationship between images and text. And it used a diffusion model to actually generate the pictures.

Language models kept advancing too.

Fine-tuning became important. GPT-3 was a general-purpose model. It could handle a bit of everything. But to perform a specific task flawlessly, adjustment was needed.

A model specialized in analyzing legal documents. A model specialized in summarizing medical records. A model that writes in the voice of a specific company.

Through fine-tuning, GPT could become a specialist in each industry.

Altman saw it this way.

"A general-purpose model is your baseline fitness. Fine-tuning is what turns that fitness into a specific skill."

DALL-E's success also shaped ChatGPT. People had experienced AI-generated creative work. They were astonished. "AI can do this too?" That astonishment became the psychological foundation for accepting ChatGPT.

DALL-E was the first step beyond language. Later, GPT-4 became a multimodal model that could understand both text and images.

The Launch of ChatGPT: Five Days That Changed History

Wednesday, November 30, 2022.

OpenAI quietly opened a website. Its name was ChatGPT. There was no big fanfare. They posted a single link on Twitter. "Try our new chatbot."

Altman had expectations, but not for anything like this.

Tens of thousands of people poured in. The servers strained. Hundreds of thousands within 24 hours. One million within five days.

The world was stunned.

People asked ChatGPT every kind of question.

"Explain the French Revolution to a ten-year-old." "Write a binary search algorithm in Python." "Make this email more polite." "Why does the universe exist?"

ChatGPT answered. Naturally. Kindly. Sometimes with surprising accuracy.

Social media erupted.

People shared poems, essays, code, and business plans written by ChatGPT. "Look at this, AI did all my homework!" "ChatGPT fixed my bug!" "This is almost human."

Some were thrilled. "The age of Google is over!" Others were worried. "Every job is going to disappear." Teachers were flustered. "What do we do if students use this for homework?"

Within two months, it had reached 100 million users.

It was the fastest-growing app in history. Netflix took 10 years to reach 100 million. Facebook took four and a half years. Instagram took two and a half. ChatGPT took two months.

Why so fast?

It was too easy. There was no complicated setup. No commands to learn. You just talked to it. Even a grandmother could use it.

It was immediate. You asked a question and got an answer right away. Rather than ten links like Google, it gave you the answer directly.

It was fun. Talking with AI felt remarkable. It was addictive.

The education world was turned upside down.

Students wrote their essays with ChatGPT. Teachers had to rethink how they evaluated work. Some schools banned ChatGPT. Others actually taught it.

Developers saw their productivity explode.

Time spent writing code was cut in half. They used ChatGPT to find bugs. Learning a new programming language became much easier.

Creators gained a new tool.

Novelists brainstormed ideas. Composers polished lyrics. Marketers crafted ad copy.

Altman watched it happen in real time.

He told his team: "I think we've built something we can't take back."

The words carried equal parts fear and pride.

ChatGPT wasn't perfect. It stated wrong information with confidence. It hallucinated. It had biases. But people accepted it anyway.

Governments began to move.

Countries started debating AI regulation. Europe rushed to pass an AI Act. The U.S. Congress called Sam Altman to testify.

ChatGPT was not just a product.

It was the official opening of the age of artificial intelligence. AI was no longer a laboratory experiment. It had entered people's daily lives. Students, office workers, grandmothers, children. Everyone was talking to AI.

Altman's dream had become reality.

Intelligence was now open to everyone. Anyone could use world-class AI for free. This was a revolution.

November 30, 2022.

The day history changed.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

2.3 Conflict with Elon Musk

Part 2: The Birth and Growth of OpenAI

[← Previous](#)

[See all chapters](#)

[Next →](#)

Musk's Demand for a Majority Stake

In December 2015, Sam Altman and Elon Musk launched a new venture called OpenAI.

Musk was a billionaire who had already gained worldwide fame through Tesla and SpaceX. Altman was a young leader who ran Y Combinator and was known as Silicon Valley's 'kingmaker.' The two agreed that AI could pose a grave danger to humanity. There was a particular fear that Google, by acquiring the AI company DeepMind, might come to monopolize AI technology.

Musk and Altman shared the belief that 'AI must not be monopolized by a handful of large corporations.' They set up OpenAI as a nonprofit. The 'Open' in the name was intentional: a promise to publish research results and make the benefits available to everyone.

Musk pledged to donate enormous sums to OpenAI. The initial commitment was \$1 billion. In practice, he contributed roughly \$45 million between 2016 and 2018, making him OpenAI's largest early backer.

The first two years went well.

OpenAI researchers built an AI that played video games and ran experiments teaching a robotic hand to solve a Rubik's Cube. The atmosphere was one of free inquiry, fitting for a nonprofit. Musk visited the OpenAI office occasionally to encourage the researchers. Altman was busy juggling Y Combinator and OpenAI, but the relationship between the two founders remained solid.

Trouble began in the second half of 2017.

The research team had discovered the potential of a language model called GPT-1. Scaling it up looked like it could produce remarkable results. Sam Altman and Ilya Sutskever ran the numbers. Building a truly powerful AI would demand enormous computing resources, billions of dollars' worth.

Ilya Sutskever emailed Elon Musk: 'As we get closer to AGI, it makes sense to become less open. The "Open" in OpenAI means that everyone should benefit from AI once it's built, it's perfectly fine not to publish the underlying science.' Musk replied: 'Agreed.'

By late 2017, the OpenAI founders and Musk had agreed that the next step was to form a for-profit company. The nonprofit structure left them unable to compete with tech giants like Google and

Facebook. What came next was the problem.

Musk insisted that he needed to hold a majority stake in OpenAI. He also said he would control the board and become CEO. His logic was straightforward: AGI is the most dangerous technology in human history. Decision-making by committee would be slow, and they'd fall behind Google. What was needed was strong, singular leadership.

Musk also proposed merging OpenAI with Tesla. Tesla's autonomous driving technology was already collecting vast amounts of data, he argued, and combining it with OpenAI's research would create real synergy. In early February 2018, he emailed Altman: 'OpenAI needs to be attached to Tesla as the cash cow. Tesla is the only path that could compete with Google, however unlikely.'

Sam Altman and the other co-founders could not accept the proposal.

The OpenAI team told Musk: 'The current structure gives you a path to unilateral, absolute control over AGI. You've said you don't want to control the final AGI, but this negotiation has shown that absolute control matters deeply to you.' Altman thought to himself: Why did we build OpenAI? To prevent AI monopoly. And what Musk is demanding right now is exactly that monopoly.

The OpenAI team added: 'OpenAI's mission is to make the future good and to avoid an AGI dictatorship. If the company genuinely makes progress toward AGI, we worry that , contrary to your current intentions , you would choose to maintain absolute control.'

Altman had another reason. He had no desire to work under Musk. Musk was a brilliant entrepreneur, but he was also known as an unpredictable boss , contacting employees at all hours, changing direction without warning. The talented researchers Altman had recruited to OpenAI were unlikely to want to work under him either.

The negotiations collapsed.

Musk replied briefly by email: 'The discussion is over. To be clear, this is not an ultimatum to accept what we discussed before. That is no longer on the table.' A few weeks later, in February 2018, Musk resigned from the OpenAI board.

As he left, Musk said that OpenAI's chances of success were zero. He announced plans to build an AGI competitor inside Tesla.

The official reason Musk gave for leaving was 'conflict of interest.' Tesla was developing AI for autonomous driving, he explained, which meant it would be competing with OpenAI for talent. Andrej Karpathy, one of OpenAI's best researchers, did in fact move to Tesla around that time. When Musk gave his farewell remarks at the OpenAI office, he spoke mostly about the conflict-of-interest issue. Most employees didn't fully buy it.

The real reason was control.

Altman soon got an urgent call from Reid Hoffman: 'Elon has cut off his support. What do we do?' Musk had stopped the \$1 billion donation he'd pledged. He had already given \$100 million, but he withheld the rest.

OpenAI found itself in sudden financial crisis. There wasn't enough money to pay the researchers' salaries, let alone cover computing costs. Reid Hoffman quickly put in an additional \$10 million to get through it.

Altman faced a crossroads. How could he raise billions of dollars without Musk? Attracting large-scale capital while keeping the nonprofit structure was simply not possible.

In March 2019, OpenAI created a distinctive structure it called a 'capped-profit' company. Investors could earn returns, but once profits exceeded a certain limit, the remainder would flow back to the nonprofit foundation.

In December 2018, Musk left OpenAI with a parting message: 'Raise billions of dollars immediately every year, or give up.'

Musk's warning wasn't wrong. AI development did require astronomical sums of money. But Altman was determined to raise that money his way, not Musk's.

The split became the most significant turning point in OpenAI's history. Losing Musk as a powerful backer was a real blow, but Altman gained the freedom to lead the company on his own terms. That freedom laid the groundwork for what would later become GPT-3, and then the revolution called ChatGPT.

Two Visions, Divided

Elon Musk did not go quietly after leaving OpenAI.

He grew louder over time. As OpenAI took in major investment from Microsoft and scored a massive success with ChatGPT, Musk's attacks intensified.

When ChatGPT launched in November 2022, Musk was reportedly furious. It stung that the company he had left was now making world-changing products, but the bigger grievance was how. In December 2022, a month after ChatGPT's release, Musk cut off OpenAI's access to Twitter data. By then, he had already acquired Twitter (now X).

On February 17, 2023, Musk posted on Twitter: 'OpenAI was created as an open-source, nonprofit to serve as a counterweight to Google. It is now a closed-source, maximum-profit company effectively controlled by Microsoft.'

Musk's first line of attack centered on the word 'Open' in the name.

Musk claimed that OpenAI used a 'fake humanitarian mission' to attract him and other investors. From the start, they had promised to publish all research openly. Anyone should be able to see and

use OpenAI's technology. Through GPT-1 and GPT-2, the code was actually released.

Starting with GPT-3, things changed. OpenAI stopped disclosing the details of its models. GPT-4 was even more guarded. Almost nothing was revealed about how it was built, what data was used, or what the model architecture looked like. OpenAI explained this as being "for safety," on the grounds that bad actors could misuse the technology for harmful purposes.

Musk thought this was nothing but an excuse. The real reason, he argued, was to prevent competitors from copying OpenAI's technology. It was about protecting commercial secrets, not human safety.

Musk also wrote on Twitter: "I'm still confused about how the non-profit I donated ~\$100M to became a \$30B for-profit. If this is legal, why doesn't everyone do it?"

The second criticism concerned OpenAI's relationship with Microsoft.

Musk claimed OpenAI had become "a closed, de facto subsidiary of Microsoft." From 2019 onward, Microsoft invested a total of \$13 billion (roughly 17 trillion Korean won) in OpenAI. Receiving that much money makes it hard to stay independent.

Musk cited an interview with Microsoft CEO Satya Nadella. During the 2023 episode when Sam Altman was fired and then reinstated, Nadella said: "Even if OpenAI disappeared tomorrow.. we have all the IP rights and all the capability. We have the people, the compute, the data, everything." Musk argued this statement revealed who OpenAI was really working for.

On March 1, 2024, Musk filed a lawsuit in San Francisco court against OpenAI, Sam Altman, and Greg Brockman. The suit alleged breach of contract, breach of fiduciary duty, and unfair business practices.

In the 83-page filing, Musk's attorney Mark Toberoff wrote: "This action exposes the defendants' grandiose charitable theater and holds them accountable for their false statements to Musk and the public. This is not just about a \$100 billion startup. The future of AGI is at stake."

In the lawsuit, Musk argued that GPT-4 had already reached AGI (artificial general intelligence), meaning intelligence equal to or surpassing that of humans. The agreement between OpenAI and Microsoft contained a clause stating that once AGI was achieved, Microsoft's exclusive license would terminate.

Musk accused OpenAI of concealing the fact that GPT-4 was AGI in order to protect its financial interests.

Through the lawsuit, Musk demanded three things.

First, that OpenAI be compelled to return to its original mission. Second, that no one, including Microsoft, be allowed to profit from OpenAI's technology. Third, that the money he donated be returned if it had been used for private gain rather than the public good.

OpenAI pushed back hard.

OpenAI published several emails and documents on its official blog, presenting evidence that Musk himself had wanted the for-profit conversion. On September 15, 2017, Musk had even personally incorporated a public benefit corporation called "Open Artificial Intelligence Technologies, Inc." OpenAI stated: "Elon left when he couldn't get majority equity and full control, and told us we would fail."

OpenAI added: "You can't build AGI through litigation. We respect Elon's achievements and are grateful for his early contributions, but he should compete in the marketplace, not the courts."

OpenAI concluded: "We are sad that the person who inspired us to aim higher later said we would fail, started a competitor, and then filed a lawsuit when we made meaningful progress without him."

Musk dropped the lawsuit in June 2024, without giving any reason.

Two months later, in August, he filed again. This time he added charges of organized crime under the RICO Act. In November 2024, he brought Microsoft into the suit as a defendant, arguing that the relationship between OpenAI and Microsoft was anti-competitive and monopolistic.

In March 2025, the court rejected Musk's motion for a preliminary injunction, a request to block OpenAI's conversion to a for-profit entity. Judge Yvonne Gonzalez Rogers ruled that the evidence Musk had submitted was insufficient to justify a preliminary injunction. She noted that some emails suggested Musk himself had at some point considered the possibility that OpenAI might eventually become a for-profit company.

The court said it was prepared to offer an expedited trial in the fall of 2025. The legal fight is not over.

Are Musk's criticisms actually justified?

Some believe Musk is genuinely concerned about human safety. His point that OpenAI abandoned the transparency and public interest it promised at the outset has some merit. OpenAI did make a significant shift, from nonprofit to for-profit and from open to closed.

Others think Musk is simply jealous. After leaving a company he couldn't control, he couldn't stomach watching it succeed. In 2023, Musk founded his own AI company, xAI. In February 2025, he went so far as to make a bid alongside private equity firms to acquire OpenAI for \$97.4 billion (roughly 130 trillion Korean won). The OpenAI board rejected it unanimously.

Sam Altman tries not to respond to Musk emotionally. In a CNBC interview, Altman said: "Honestly, I don't think about him that much. I thought he was tweeting all day about how terrible OpenAI is, that our models are bad, that we can't be a good company."

At bottom, the conflict between the two men is a difference in philosophy. Musk believes that technology as powerful as AGI must be controlled by a single strong leader to be safe. Altman believes that concentrating that kind of power in one person is itself the danger.

History will judge who is right. What's certain now is that this conflict reshaped the entire AI industry. Because Musk left, Sam Altman was able to lead OpenAI his own way. The result was ChatGPT, and the AI revolution the whole world experienced.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

3.1 The CEO Ouster: A Five-Day Nightmare

Part 3: Ouster, Return, and the Victory of a Vision

[← Previous](#)

[See all chapters](#)

[Next →](#)

The Sudden Dismissal

On the afternoon of Friday, November 17, 2023, Sam Altman's phone rang.

He was asked to join a video call. Minutes later, he was told he was no longer CEO of OpenAI. The board said he had been "not consistently candid." No specific reason. No prior warning. Just: you're fired.

Sam Altman later described that moment as feeling "like a weird fever dream." At the very moment ChatGPT was shaking the world, the man who built the company was being thrown out of it. It was the most dramatic episode in Silicon Valley history.

The first feeling was confusion. Just the day before, Sam had been discussing the company's future. He had been drawing up plans for new model development. Then everything collapsed at once. He couldn't understand what had happened. The question "why?" wouldn't leave his mind.

The board gave no clear answer. A communication problem? He had no idea what that meant. The confusion quickly turned to frustration.

A Dream Shattered

Sam Altman had never thought of OpenAI as just a job.

The company was his dream. Its mission was to build AGI and benefit humanity. He worked around the clock, persuading investors, recruiting researchers, negotiating with governments. When ChatGPT succeeded, he felt it was the brightest moment of his life. And then he was expelled from everything he had built.

Frustration quickly became anger. The dismissal was deeply flawed even as a matter of process. Microsoft, the largest investor, had been notified only minutes before. Employees learned what had happened from a blog post. Why had the board handled things this way? Sam felt betrayed.

What hit him hardest was learning that co-founder Ilya Sutskever had voted for his removal. The thought that a colleague who had shared his dream had turned against him was painful. Even through the anger, a deep sadness washed over him.

Sam loved his time at OpenAI. He loved the people he worked with. The engineers debugging code deep into the night, the researchers training models, the policy team wrestling with safety. The thought of saying goodbye to all of them hurt. He posted a short message on social media: "I loved my time at OpenAI." There was grief in those words.

An Unexpected Reversal

Co-founder Greg Brockman announced his resignation. "OpenAI without Sam means nothing," he said. A few hours later, senior executives called. "We stand with you." Researchers sent messages. "You have to come back." In the middle of the grief, something like gratitude began to stir.

The next day, something even more remarkable happened.

More than 700 of OpenAI's 770 employees signed an open letter. The message was plain: "If Sam Altman does not return, we will all move to Microsoft." Sam choked up when he saw it. The entire company had stood up for him. This wasn't mere loyalty. It was trust. It was belief in the vision they had built together.

Microsoft CEO Satya Nadella spoke out publicly. "We will build a new AI research lab for Sam and his team." The world's largest technology company was in Sam's corner. Gratitude began to crowd out everything else.

Sam Altman passed through confusion, frustration, anger, and grief, and came out the other side at gratitude. He realized he wasn't alone. Seven hundred colleagues stood with him. A powerful partner stood with him. The whole world was watching. He later said: "In the span of a few days, I experienced every emotion a person can feel. It started with confusion and ended with gratitude."

That emotional journey made him stronger. The crisis showed him how solid what he had built really was. The trust he had with people, the shared vision, the culture they had made together. All of it proved stronger than any individual.

Through being ousted, Sam Altman paradoxically confirmed where his real power came from. Not from authority, but from trust.

Five Days of War

Friday afternoon, Sam Altman was removed as CEO. Early Wednesday morning, he was back as CEO. Just five days.

Those five days were the most dramatic in OpenAI's history, and perhaps in Silicon Valley's entire history. Time didn't pass quickly. Too much happened, compressed into too little space. Everything unfolded in real time.

Friday: The Bombshell

On the afternoon of November 17, a brief post went up on the OpenAI blog. "Sam Altman has stepped down as CEO. We are appointing CTO Mira Murati as interim CEO." Silicon Valley froze. The founder of the company that was changing the world with ChatGPT had been fired without a word of warning.

A few hours later, Greg Brockman announced his resignation. "I co-founded OpenAI with Sam Altman. I cannot stay at an OpenAI without him." That announcement was a signal that this was not a routine executive change.

The Weekend: Rebellion Begins

On Saturday and Sunday, OpenAI was on the verge of explosion. Hundreds of messages flew across Slack channels. "What on earth is happening?" "The board has lost its mind." "What do we do?" Senior executives called emergency meetings.

Researchers gathered. Engineers compared notes. There was one conclusion: "Sam has to come back." Sunday afternoon, Sam Altman appeared at OpenAI's headquarters wearing a visitor badge. He took a photo and posted it on social media. "First and last time I'll wear one of these."

The message was clear: I'm not going away. Investors started moving. Thrive Capital, Khosla Ventures, and Tiger Global pressed the board. "Bring Altman back." The board did not respond.

Monday: War Declared

On November 20, things exploded.

The OpenAI board announced a new interim CEO: Emmett Shear, former CEO of Twitch. The board's message was plain: Sam Altman is not coming back. Employees responded immediately. That afternoon, an open letter was published.

More than 700 of the 770 employees signed it. Researchers, engineers, the policy team, even HR and accounting. Almost the entire company speaking with one voice: "If Sam Altman and Greg Brockman do not return, we will all resign and go to Microsoft. The board must step down."

The letter was an act of collective defiance almost without precedent in corporate history. Ninety-five percent of the staff had chosen their leader. They had rejected the board.

That same day, Microsoft CEO Satya Nadella posted on Twitter.

"We are bringing Sam Altman and Greg Brockman on board to lead a new AI research team at Microsoft. Anyone from OpenAI who wants to join is welcome." This was not a simple job posting. It was an ultimatum to the board: "Send Sam back, or we'll take the whole company."

Tuesday: The Board Collapses

On November 21, the board's position completely collapsed. Even the newly appointed interim CEO Emmett Shear issued a demand to the board: "If you don't submit the reasons for the dismissal in

writing, I will resign as well." Ilya Sutskever, who had led the ouster, posted an apology on Twitter: "I deeply regret my participation in the board's actions. I never intended to harm OpenAI, which I love."

Investor pressure reached a breaking point. Legal advisors warned: "If this continues, the company will collapse." The board could hold on no longer. Negotiations began.

Wednesday, Early Hours: Victory Declared

In the early hours of November 22, OpenAI issued an official statement: "Sam Altman will return as CEO. A new board is being formed. Bret Taylor (former Salesforce CEO) will serve as chair, joined by Larry Summers (former Secretary of the Treasury) and Adam D'Angelo (CEO of Quora)." Sam Altman posted on Twitter: "I love OpenAI. Everything I did over the past few days was to protect this team and its mission. Now let's get back to work."

The five-day nightmare was over.

Days That Passed Like a Fever Dream

Sam Altman later described this period as "a dream of a strange fever." Everything felt unreal. His sense of time dissolved. Each day felt like a week. He recalled: "Every morning I woke up to a completely different situation. Who had said what, who had switched sides, which negotiations were underway. Everything was playing out on Twitter in real time."

With the entire world watching, a fate was decided on social media.

Board decisions, employee letters, CEO replacements, investor pressure, and the final return, all of it unfolded online, in real time. This episode wrote a new chapter in corporate governance. A board fired a CEO, and employees brought him back. Formal power structures crumbled, and real influence won.

Five days. Five days that completely redefined what OpenAI is.

A Company That Ran Perfectly Through the Chaos

Ask Sam Altman what he was most proud of during this episode, and the answer is unexpected.

"His own return?" No.

"The support of his employees?" That wasn't it either.

What Sam was most proud of was this: "That the company ran perfectly while I was gone."

The CEO had been fired, the board was in chaos, and employees were preparing collective action. At most companies, work would have ground to a halt. Customer service would have stopped, development would have frozen, systems would have broken down.

But OpenAI was different.

ChatGPT didn't stop

During those five days, hundreds of millions of people around the world used ChatGPT. They asked questions, received answers, wrote code, and requested translations. The service never went down once.

The services of thousands of businesses using the API kept running. Customer support chatbots, coding assistants, translation services. Everything operated normally.

The infrastructure team held the systems together. Engineers monitored the servers. No one panicked.

Research continued

Researchers didn't stop their experiments. They trained new models, wrote papers, and analyzed data.

One researcher later said: "We knew Sam would come back. Our job until then was to keep doing the best research we could."

The executive team managed the crisis

CTO Mira Murati became interim CEO. She managed the situation with calm. Every morning she held an all-hands meeting, explained the situation to employees, and kept anxiety at bay.

COO Brad Lightcap handled external communications. He held calls with investors and reassured partner companies. "OpenAI is not shaken."

Chief Strategy Officer Jason Kwon organized the employee letter. The legal team reviewed legal options. The HR team looked after the mental well-being of employees.

Everyone performed their role perfectly.

What Sam Saw

Sam Altman watched all of this from the outside. Not as CEO, he saw how the company he had built actually moved.

He was astonished.

The company did not need him. Or more precisely, the system he had built was running perfectly. Autonomous, accountable, capable people were each doing their best.

Sam said: "The moment I was most proud was when I confirmed that any one of our executives could run this company flawlessly."

The true measure of a leader's success

Many CEOs believe the company will fall apart without them. So they make every decision themselves. They attend every meeting. They fix every problem personally.

Sam Altman operated from a different philosophy. He ran the company by "giving enormous responsibility to genuinely brilliant people." He gave a small number of people wide authority. He did not micromanage.

This crisis proved his philosophy right.

Organizational resilience

Psychologists speak of resilience: how quickly an organization recovers when crisis strikes.

OpenAI showed extraordinary resilience. Even in the worst crisis imaginable, with its CEO fired, core functions never stopped. The organization did not panic. If anything, it pulled together more tightly.

What made that possible?

There was a strong vision. Every employee shared the mission of benefiting humanity through AGI. The CEO could change; the mission could not.

There was trust. Employees trusted one another. They trusted the leadership team. Each person knew the others would give their best.

There was autonomy. People did not wait to be told what to do. They judged the situation and acted on their own.

Sam Altman had built all three.

"These are the people I chose"

In an interview after his return, Sam said:

"When I saw that I had chosen these people, and that what I had taught them to some degree was alive inside them, I was genuinely proud. The capability they showed under crisis was remarkable."

He did not believe he had achieved his success alone. He had built the team, and the team had saved him.

That is the paradox of leadership. The truly strong leader is the one who makes the organization work even in their absence.

The five-day crisis proved he was a real leader. The organization he had built ran perfectly without him, and that was precisely why he was indispensable. This was the moment he was most proud of, not his own return, but the capability of the team he had made.

← Previous

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

3.2 The Employees' Dramatic Loyalty

Part 3: Ouster, Return, and the Victory of a Vision

[← Previous](#)

[See all chapters](#)

[Next →](#)

A Declaration of Support Signed by the Vast Majority of 770 Employees

On Friday, November 17, 2023, Sam Altman was ousted from his position as CEO.

OpenAI employees were stunned. 'What on earth just happened?' Confusion and anger spilled into hallways, meeting rooms, and Slack messages. The man who had built the company to this point was being forced out with no explanation.

The weekend passed, and Monday morning arrived. The board appointed someone named Emmett Shear as the new interim CEO, a figure known for co-founding Twitch. The board may well have thought, 'It's over. We won.'

The employees thought otherwise.

That afternoon, a document began circulating inside OpenAI. The title was simple: 'We stand with Sam.' The contents were stronger still. 'If the board does not reinstate Sam Altman and Greg Brockman, we will all resign and join Microsoft.'

At first, a few dozen people signed. An hour later, the count was past 100. By evening, it had crossed 500. The next morning, signatures broke through 700.

The total headcount was 770. That meant nearly 95 percent of the company had signed.

Among them were researchers, engineers, and designers. HR, finance, and legal staff signed too. Even Mira Murati, the interim CEO the board had installed after ousting Sam, added her name. Was she betraying the board that had made her CEO? No. She loved the company, and believed it couldn't survive without Sam.

What made it even more striking was that the employees didn't stop at signing.

Employees began posting the same phrase on X (formerly Twitter): 'OpenAI is nothing without its people.' One by one, then in twos, then from hundreds of accounts, the same sentence appeared. It had the rhythm of a protest chant.

Sam Altman retweeted every single one of those messages from his own account. Quietly, without a word, he simply pressed the like button.

Why did so many people rally behind Sam Altman?

Was it money? Partly, yes. OpenAI employees didn't receive stock, but they held something called PPU's , a special right that would pay out handsomely if the company did well. With Sam gone, the company looked like it might collapse, and those PPU's would turn into worthless paper.

But it wasn't only that.

Sam Altman gave his employees a dream. 'We can build AGI. We can change the course of human history.' That might sound like bluster, but for the people who actually built ChatGPT, it was real. They genuinely believed they were changing the world.

Sam trusted his employees. He gave them enormous responsibility, let them make their own decisions, and pushed them to move fast. When someone made a mistake, he said, 'That's fine , try again.' Employees felt Sam respected them as human beings.

So when Sam was ousted, his employees felt betrayed.

'The board made this decision without listening to us at all.' 'This is our company , are they saying our voices don't matter?' 'Can we build AGI without Sam? Can this company even survive without him?'

The employees' signatures were not a simple protest.

It was an ultimatum.

Microsoft had already announced it was bringing on Sam and Greg. CEO Satya Nadella stepped forward personally and said, 'Sam Altman will lead our new AI research division. Any OpenAI employee who wants to follow him is welcome.'

Those weren't empty words. Microsoft had invested \$13 billion (roughly 17 trillion Korean won) in OpenAI. It had deep pockets, plenty of positions, and all the infrastructure needed. Taking in 700 employees was effortless for a company of that size.

The board must have been terrified.

'Are 700 people actually going to walk out?' 'That would leave only 70 people in the building.' 'You want 70 people to build GPT-5?' It was impossible.

OpenAI was built from people, not buildings. The people who knew how to train GPT models, the people who knew how to run the servers, the people who managed the customers , if they all left, OpenAI would become a ghost company with nothing but a name.

The investors were furious too.

Prominent firms like Sequoia Capital and Thrive Capital called the board directly. 'What have you done?' 'Are you planning to destroy an \$86 billion company overnight?' 'Bring Sam back immediately.'

Salesforce CEO Marc Benioff also stepped in. 'OpenAI talent , come to us. We'll pay more and sort out your visa situations.' Résumés from OpenAI employees started flooding into Google DeepMind as well.

OpenAI was on the verge of complete collapse.

And all of it happened in just five days.

Sam was fired on Friday. Chaos followed over the weekend. Employees signed the letter on Monday. The board surrendered on Tuesday. It was the fastest CEO reinstatement in Silicon Valley history.

Where did that loyalty come from?

Sam Altman had a certain kind of charisma. He wasn't a commanding speaker, and he didn't have a big voice. But when he talked with someone, he looked them in the eyes and genuinely listened. He didn't just say 'your opinion matters' , he actually incorporated employees' ideas into the product.

ChatGPT itself had come from an idea put forward by employees.

At first, no one had thought about building a chatbot. Then a few employees said, 'People seem to enjoy having conversations in the GPT-3 Playground.' Sam said, 'Then turn it into a product,' and within a few weeks, ChatGPT was out in the world.

Sam was not afraid of failure.

'Build fast, ship fast, learn fast.' That was his philosophy. Don't wait for perfection , put it out into the world and watch how people respond. Inside that kind of culture, employees were free to experiment and take risks.

So losing Sam wasn't just a matter of losing one CEO.

It meant the company's entire culture, vision, and way of working would collapse. The board thought that removing one person , Sam , would settle things. But the employees believed that Sam was OpenAI.

The 700 signatures saved the company.

To be precise, they protected the company's soul. If the employees had stayed quiet, Sam would have gone to Microsoft. OpenAI would have kept its name but lost everyone who mattered, and it might have faded into history.

But the employees stood up.

They fought for the company they had built, the leader they believed in, and the future they had dreamed of. And they won.

This episode showed what a 21st-century company can look like. In the old days, the board made every decision. But now it became clear where the real power in a company lies. In a technology company, the most important asset is not the buildings, not the patents, and not the money.

It's the people.

And the dream they are building together.

The Board's Surrender and the CEO's Return

By Monday night, the situation looked hopeless.

The board was surrounded on all sides. Seven hundred employees had declared they would leave; Microsoft had said it would take them all; investors were furious and preparing lawsuits. Even Mira Murati, the interim CEO the board itself had installed, had signed the letter.

And yet the board did not give up.

More precisely, they couldn't give up. They had been the ones to fire Sam, and their pride would not let them turn around and say, 'Sorry, we were wrong.' So they looked for another way.

'Let's find a CEO who isn't Sam!'

The board spent the entire weekend searching for a new CEO candidate. Among the people they reached out to, one name stood out above the rest: Dario Amodei, the CEO of an AI company called Anthropic.

What kind of company is Anthropic?

It was founded by people who left OpenAI. They thought Sam Altman's rapid commercialization approach was dangerous, so they walked out. Their argument was: 'We'll build AI more safely and more ethically.'

From the board's point of view, Dario was a perfect candidate. He understood AI technology, cared about safety, and seemed unlikely to move as fast as Sam. There were even reports that the board floated the idea of merging the two companies.

Dario said no.

'Sorry, but we're going our own way.' He had no intention of walking into OpenAI's chaos.

The board's Plan B had also failed.

On Sunday evening, the board announced a man named Emmett Shear, the founder of Twitch, as the new CEO. The board probably thought: 'Good. It's over. Sam isn't coming back.'

The employees grew even angrier.

'So they're completely ignoring what we want?' 'Who on earth is Emmett Shear?' 'Can the guy who built Twitch even understand ChatGPT?' Angry messages exploded across social media.

Microsoft CEO Satya Nadella spoke publicly.

'We will work well with CEO Emmett Shear. At the same time, we welcome Sam Altman returning to OpenAI. The choice rests with the OpenAI board and its employees.'

What he was really saying was: 'We don't care either way. Whether Sam comes back or not, we win. If he doesn't come back, 700 employees come to us. If he does, our investment is safe.'

Tuesday morning arrived.

Sam Altman appeared at OpenAI headquarters, wearing a visitor's badge. He walked into the company he had founded as a guest. He posted a photo on X: 'First and last time visiting OpenAI as a visitor.'

That single photo spread around the world.

People laughed, cried, and raged. 'The founder walks into his own company as a visitor?' 'The board has truly lost its mind.' Public opinion was entirely on Sam's side.

Negotiations began.

Adam D'Angelo (CEO of Quora) appeared as the board's representative. Sam Altman, Greg Brockman, investors, and even Emmett Shear took seats at the negotiating table.

Sam's demands were clear.

"Overhaul the board entirely. If the people who pushed me out are still there, I will never come back."

The board was cornered.

Accept, and they would suffer a total defeat. Refuse, and the company would collapse. They had to choose one or the other.

At 10 p.m. on Tuesday, a post appeared on OpenAI's official account.

"We have reached an agreement in principle for Sam Altman to return as CEO. The new board will consist of Bret Taylor (Chair), Larry Summers, and Adam D'Angelo."

It was a return after just five days.

Fired on Friday, back on Tuesday. It was the shortest firing in Silicon Valley history. Steve Jobs returned to Apple after 11 years. Sam came back in five days.

The makeup of the new board was interesting.

Bret Taylor was the former co-CEO of Salesforce, a massive technology company. Smart, experienced, and seen as neutral, he became chair.

Larry Summers was a Harvard professor and former U.S. Secretary of the Treasury. An economist with government experience and a prestigious reputation. That kind of person became a board member.

Adam D'Angelo was the only one who stayed, even though he had voted to remove Sam. Why was he kept? It was likely a negotiated outcome. There may have been a deal along the lines of: "You played a mediating role, so we'll give you one seat."

Ilya Sutskever was out. Helen Toner and Tasha McCauley were gone too. Everyone who had voted to remove Sam lost their seat.

Sam Altman posted a brief message on X.

"I love OpenAI. Everything I did over the past few days was to protect this team and its mission. When I agreed to join Microsoft on Sunday evening, I thought that was the best path for me and the team. With the support of the new board and Satya, I am now returning to OpenAI."

It was not a declaration of victory. It was a message of reconciliation.

Employees cheered. "OpenAI is nothing without its people" filled social media again, this time as a celebration.

Not everything had a happy ending.

Wounds remained. Trust had been broken. Ilya Sutskever had been Sam's closest colleague, and yet he had betrayed him. Would those wounds heal easily?

Cracks appeared inside the company too. The conflict between two camps, "Safety comes first" versus "We need to move fast", had surfaced openly. Sam's return did not make that argument disappear.

On the outside, people began to see OpenAI differently.

"Can this company be trusted?" "What if the CEO gets fired again?" "Is it safe to invest?" Cracks had formed in its reputation.

Competitors saw an opening. Companies like Google, Anthropic, and Mistral began marketing themselves as stable alternatives. "We're not chaotic like OpenAI. Choose us."

Microsoft came out as the biggest winner.

Satya Nadella made a series of brilliant moves. He immediately offered Sam a job, said Microsoft would take in the entire staff, and at the same time kept lines open with OpenAI's board. He built a structure where Microsoft would win no matter which side prevailed.

Microsoft's stock actually rose on the news of Sam's return. Investors concluded that Microsoft had taken full control of OpenAI.

One thing had changed.

Microsoft joined the board as a non-voting observer. No voting rights, but a seat in every meeting with full visibility. The message was clear: "We will never again find out five minutes before it happens."

OpenAI was no longer a fully independent company.

On November 29, the official announcement came.

"CEO Sam Altman, President Greg Brockman, and CTO Mira Murati have officially returned. OpenAI will continue to pursue its mission of developing safe and beneficial AGI."

The company appeared to have returned to normal.

But internally, much had changed. Sam Altman's power was stronger than ever. There was no longer a board capable of stopping him. The new directors had come to support Sam.

The voices that had prioritized safety grew quieter. Ilya was gone, and Toner and McCauley had left too. OpenAI had now shifted decisively toward "faster, and faster still."

The employees felt relieved, but also a little uneasy.

"We won. Sam is back." That felt good. But at the same time, a question lingered: "Are we really on the right path?"

The five-day drama was over.

The king had returned. But the kingdom was not the same as before. There were wounds, there was suspicion, and things had changed.

This episode left a lesson not just for OpenAI, but for tech companies around the world. "A board can fire a CEO, but employees can bring him back." "The real owners of a company are not those with power on paper, but the people doing the work."

Sam Altman came back stronger. He had confirmed that his team cared enough to fight for him. He had also confirmed that he had built a team capable enough to keep the company running without him.

But he also came away humbled.

"The cleanup is harder than the crisis itself." He said this later. Healing the wounds after his return, rebuilding trust, pulling the team back together, all of that was far harder than simply coming back as CEO.

The five-day war was over. The real fight was just beginning.

Ilya Sutskever's Apology

Ilya Sutskever.

You cannot understand OpenAI without knowing this name.

He was a co-founder of OpenAI, alongside Sam Altman. He served as chief scientist, was one of the people who built GPT from the ground up, and was a legendary researcher in the field of AI. He had studied under Geoffrey Hinton, often called the godfather of AI, and was a central figure in the deep learning revolution.

If Sam Altman was a business genius, Ilya was a technical genius.

Together they were a perfect pairing. When Sam said, "Let's build this," Ilya would answer, "I know how." While Sam was out persuading investors, Ilya was in the lab training models.

OpenAI's employees held Ilya in deep respect.

"Without Ilya, there would be no GPT." "He's one of the few people who truly understands AI." "If Sam is the face of the company, Ilya is the brain." These were the kinds of things people said inside the company.

And yet, this was the same Ilya who had voted to push Sam out.

On that Friday, when the board fired Sam, Ilya was there. He voted yes. He had ousted the friend, the colleague, the partner with whom he had built the company.

Why?

Ilya took AI safety more seriously than almost anyone. He believed that AGI, artificial general intelligence, could save humanity, but could also destroy it. That was why he thought the work had to proceed slowly, carefully, and safely.

Sam Altman was different.

Sam pushed for "build fast, ship fast, change the world." When they launched ChatGPT, when they built GPT-4, speed always came first. "We have to move faster than the competition." "We have to own the market before Google catches up."

Ilya was uneasy.

"Are we moving too fast?" "Did we do the safety checks properly?" "If we keep rushing like this, something will go wrong." His worries kept growing.

From the summer of 2023, cracks began to form between them.

Sam held developer conferences, announced new features, raised investment, and expanded rapidly. Ilya kept saying, "Wait, this wasn't what we originally planned."

They argued repeatedly.

"Safety comes first" versus "speed is everything." "Let's go slowly" versus "you have to move fast to win." "Let's hold to the nonprofit mission" versus "the mission only survives if the company does."

The conflict surfaced in the boardroom too.

Ilya spoke with the other board members, Helen Toner and Tasha McCauley. "Sam is going too far. We need to do something." They agreed. "You're right. OpenAI is drifting from its original purpose."

On Friday, November 17, they acted.

Ilya followed his conscience. "I have to save the company. If Sam keeps going like this, it's dangerous." He genuinely believed that.

When Sam was fired, the company fell into chaos. Greg Brockman resigned. Employees were furious. Investors protested. Microsoft stepped in.

Ilya was stunned.

"I never expected this." He thought the employees would back him up. "We saved the company. Everyone will be grateful." But it was the exact opposite. The employees called him a traitor.

When Monday's letter of 700 signatures appeared, Ilya's name was on it.

He had turned around at the last moment. The very person who had agreed to push Sam out was now signing a letter demanding Sam's return.

And he posted on X.

"I deeply regret my involvement in the board's actions. I never intended to harm OpenAI. I love everything we built together, and I will do everything I can to reunite the company."

That post was a bombshell.

The central figure behind Sam's ouster had apologized publicly. The board's moral authority collapsed entirely. "If even Ilya has regrets, this decision couldn't have been right."

Sam Altman replied with a heart emoji.

□□

He said nothing. Just sent a single heart. Was it a sign of forgiveness? Or did it mean "we'll talk later"?

The employees found themselves in a tangle of complicated feelings.

Should they hate Ilya? He was a traitor. Or should they try to understand him? He loved the company, cared about safety, and acted with good intentions. He just went about it the wrong way.

After Sam returned, Ilya stayed at the company.

At first it seemed fine. Ilya kept doing research, Sam ran the company, and everyone said "we're okay now."

But inside, things were different.

Ilya was no longer a close colleague of Sam's. They sat together in meetings, but the easy jokes of earlier days were gone. Trust had been broken. Once broken, it doesn't come back easily.

In May 2024, Ilya left OpenAI.

Officially, the announcement said he was leaving "to start a new project." But everyone knew the real reason. They couldn't be together anymore.

Ilya started his own AI company.

It was called Safe Superintelligence (SSI). Even the name says something. "Safe superintelligence." He was going to pursue the dream he couldn't realize at OpenAI.

Two geniuses had set out together to change the world. One a business genius, the other a technical genius. A perfect combination. But they disagreed on how to get there. One said "fast," the other said "slow."

In the end, they parted ways.

Ilya's apology was the most human moment in the whole drama.

He made a mistake, owned it, and asked for forgiveness. Nobody is perfect. Even geniuses make mistakes. What matters is the courage to admit them.

Sam Altman was generous too. After returning, he didn't push Ilya out immediately. He gave him a chance to keep working together. Ilya left eventually, yes, but Sam didn't force him out.

This episode raises a question.

When building AI, which matters more: speed or safety? Should you move fast and change the world, or move slowly and make sure it's safe?

There's no clean answer. Ilya wasn't wrong, and Sam wasn't wrong. Both were right in their own way, and both wanted to help humanity on their own terms.

The tragedy is that they couldn't go together.

The Incident Carries Layered Meaning

After five days of drama, people struggled to make sense of what had happened.

It wasn't a simple CEO firing.

It was a textbook case of how a twenty-first-century company actually works. Inside it were power, vision, loyalty, culture, and a new era's idea of leadership.

The most important lesson is this.

"Who is the real owner of a company?"

The answer used to be simple. Shareholders are the owners. The people who put in the money own the company, elect the board, and appoint the CEO. If the board doesn't like the CEO, they fire him. End of story.

The OpenAI incident shattered that formula.

The board fired the CEO. Legally, it was a perfectly valid move. The bylaws said so, and the procedure was followed. But the company nearly collapsed.

Why?

Because the employees refused. Seven hundred people declared, "We follow Sam, not the board." The employees, who held the real power, beat the board, which held the legal authority.

This was a revolution.

No blood was shed, unlike the French Revolution, but the power structure flipped completely. The top-down command system collapsed, and the bottom-up expression of will won.

How did this become possible?

Because OpenAI was a special kind of company. It wasn't a factory. Its assets were human minds. The people who knew how to build GPT, who could write the code, who could train the models. Without those people, the company was just an empty building.

Microsoft understood that.

"We don't need the building. Just bring the people." That single line from Satya Nadella brought down the OpenAI board. It proved that human capital, not physical assets, is where the real value lies.

What is corporate bureaucracy?

It's a system that prizes rules, procedures, and formalities. "Board resolutions matter." "That's what the bylaws say." "We're legally in the right." That kind of reasoning.

That's not wrong. Companies need rules. But what happens when following the rules causes a company to lose sight of its essence? The OpenAI board fell straight into that trap.

They were formally correct.

"The CEO wasn't being straight with the board." "We have to protect the safety of the company." "The bylaws give us the authority." All of it was true. In court, the board might have won.

But in practice, they were completely wrong.

They couldn't read the employees' feelings. They didn't understand the company culture. They didn't know why Sam Altman mattered. They saw only the power on paper and missed where the real force actually lived.

What is a vision?

It's a powerful dream about the future. "We're going to build AGI. We're going to change the course of human history. This isn't just a company, it's a mission." That is exactly what Sam Altman gave his employees.

The employees weren't at OpenAI for the paycheck.

Money mattered, of course. OpenAI's PPUs held enormous value. But that wasn't the whole story. They could have earned more at Google or Meta.

They stayed at OpenAI because of the dream.

"We built ChatGPT. The world changed. What do we build next?" That excitement, that possibility, that challenge. That's what moved them.

Sam Altman was the symbol of that dream.

The employees believed that without him, there was no dream. That's why 700 people signed. Vision mattered more than money, more than stability, more than career.

The board didn't know this.

They thought, "The company runs fine without Sam." They saw the CEO as replaceable. They were wrong. Sam wasn't just a CEO, he was the embodiment of the vision.

What this incident showed is that in a twenty-first-century company, vision is stronger than power.

The dream in people's hearts carries more force than the rules on paper. The collective will of employees wields more influence than the voting rights of a board.

Many people drew comparisons to Steve Jobs.

In 1985, Steve Jobs was also pushed out of Apple. The board fired him. He returned eleven years later and turned Apple into the world's most valuable company.

Sam Altman came back in five days.

Why so fast? Because the times had changed. In 1985, employees had no way to organize collectively on behalf of a CEO. In 2023, Slack, X, and email let them mobilize in a matter of hours.

Technology changed the power structure.

Thanks to social media, the employees' voices spread instantly around the world. "OpenAI is nothing without its people." That phrase reached millions within hours. The board could not withstand the pressure of public opinion.

The investors had changed too.

Investors of an earlier era would have sided with the board. "Rules are rules. If the CEO caused problems, removing him is the right call." But the investors of 2023 were different. They were pragmatists. "If the company collapses, our money is gone. Bring Sam back."

That is modern capitalism.

Results matter more than rules. Substance matters more than form. Survival matters more than principle.

The OpenAI crisis is a blueprint for the future.

Similar events will happen again. A founder with a powerful vision clashes with a bureaucratic board, employees take the founder's side, and in the end the founder wins. That is the pattern.

Is that a good thing?

It is a complicated question. On one hand, yes. A leader with vision can change the world without bureaucracy getting in the way. Like Sam Altman.

On the other hand, it is dangerous. A leader with the wrong vision becomes just as unstoppable. What if Sam had genuinely made a dangerous decision? What if the board couldn't stop him? Who would be held accountable?

That was exactly what Ilya Sutskever feared.

"Sam is moving too fast. Someone needs to hit the brakes." His concern wasn't wrong. AI can be genuinely dangerous. But the method he chose was the wrong one.

Balance is needed.

Vision is needed, and so is accountability. Speed is needed, and so is safety. A leader's freedom is needed, and so is board oversight.

OpenAI failed to find that balance.

After Sam's return, the company tilted entirely toward "speed." Every voice that had worried about safety was gone. Now, who is left to stop Sam?

The new board members came to support Sam. Bret Taylor, Larry Summers. Fine people, but not the kind who would ever fight him.

That is the price of winning.

Sam won, but in winning he took on a far heavier burden. There is no one left to stop him. If he makes a mistake, that mistake becomes OpenAI's mistake, and could become the mistake of the entire AI industry.

That is the true meaning of this episode.

"In twenty-first century companies, people are everything." Not buildings, not patents, not money. People. The vision that binds them together. The culture that believes in that vision and follows it.

The board held only power on paper.

The employees held real power. When the two collided, real power won.

That is the victory of a powerful vision over corporate bureaucracy.

History books will one day read: "In November 2023, 700 OpenAI employees stood up for their CEO. They toppled a board that held only formal authority and defended the vision they believed in. It was a turning point in twenty-first century corporate history."

Sam Altman is not merely a CEO.

He represents leadership for a new era. A leader who guides through vision, not commands. A leader respected through trust, not authority. A leader who moves people through a dream, not rules.

OpenAI is not merely a company.

It is a revolution. A collective charge toward the future. That is why the employees fought for it with everything they had.

Founders around the world study this episode. "How do I make my employees fight for me?" Boards draw their own lesson. "You can't touch a CEO carelessly."

Investors come away changed. "People matter more than rules." Employees gain confidence. "We have power too."

That is the greatest legacy the ousting and return of Sam Altman left on twenty-first century corporate history.

The moment vision defeated power. The moment people defeated rules. The moment the future defeated the past.

And the moment one person's dream moved 700 hearts and ultimately changed the fate of a vast company.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

3.3 Overcoming the Crisis and Its Lessons

Part 3: Ouster, Return, and the Victory of a Vision

[← Previous](#)

[See all chapters](#)

[Next →](#)

Comparison with the Steve Jobs Incident

Friday, November 17, 2023. Sam Altman was watching the Las Vegas Grand Prix when a video call request suddenly came through. It was an invitation Ilya Sutskever had sent by text the night before. At noon, board members appeared on screen. Altman was told he had been removed as CEO. Five to ten minutes later, an official statement went up on the OpenAI website.

From that moment, the press began invoking a single name.

Steve Jobs.

The legendary figure who had been forced out of Apple, the company he founded, in 1985.

"A founder pushed out of his own company, this is Jobs all over again!" Ron Conway, the legendary Silicon Valley investor, wrote on X (formerly Twitter). "The most brazen board coup we've seen since the Apple board ousted Steve Jobs in 1985 just happened at OpenAI."

The comparison was irresistible.

Jobs had been away for more than a decade before returning to Apple in 1997 and changing the world with the iPod and iPhone. Would Altman write a legend like that?

Altman himself rejected the comparison outright. "Our situations were just too different."

In one interview, a host asked: "Steve Jobs once said that being pushed out of Apple was terrible medicine, but the patient needed it. Do you feel the same way?" Altman paused for a moment, then answered carefully. "Everything resolved completely in four or five days. It was like a dream where you have a very strange fever."

Time was the decisive difference.

Jobs waited twelve years, from 1985 to 1997. In that time he built NeXT and Pixar, and watched Apple come to the brink of bankruptcy. His return was a rescue mission.

Altman, by contrast, was back in just five days. Fired on Friday, his reinstatement was announced at 1 a.m. on Wednesday.

Why so fast?

The employees. More than 700 of the 770 signed an open letter. "If the board doesn't resign and bring Sam back, we will go to Microsoft." Former Google CEO Eric Schmidt said later: "When the board tried to fire Sam again on Sunday, the employees revolted. They said, choose: the company or the board. You won't find clearer feedback in any 360-degree review of a CEO."

When Jobs was pushed out, there was significant resentment inside Apple over his autocratic management style.

Some were glad to see him go.

Altman's situation was the opposite. Nearly every employee stood with him. Even co-founder Ilya Sutskever signed the letter and tweeted: "I deeply regret my participation in the board's actions."

The nature of the return was different too. Jobs fought to reclaim the company. Once back, he rebuilt Apple entirely on his own terms, purging opponents, simplifying the product line, concentrating power in his hands. Altman did none of that. He didn't consolidate power. He expanded and restructured the board instead. Bret Taylor, former co-CEO of Salesforce, became chairman. Former Treasury Secretary Larry Summers joined. More outside voices were brought in.

In one interview, Altman said plainly: "Jobs wanted to take the company back. I never wanted to take the company back. I just believed the mission couldn't be allowed to stop."

What was striking was that Altman took genuine satisfaction in watching the organization run smoothly even through the crisis. He later recalled: "The moment I was proudest was seeing the executive team operate perfectly without me, even in the middle of chaos. I knew any one of them could lead this company well."

That was the defining difference from Jobs.

Jobs proved that Apple couldn't survive without him. Altman confirmed that OpenAI could run fine without him, and was glad of it. He saw not a personal hero's story, but a victory of a strong team and a shared vision.

Both were founders with world-changing visions, both were pushed out of companies they built, and both returned in dramatic fashion. But Altman never tried to be "the next Steve Jobs." He walked his own path. He placed his faith in mission over power, in team over individual, in organizational strength over the hero's arc.

At a conference in 2024, when someone raised the Jobs comparison again, Altman smiled and said: "Eric Schmidt's advice not to fire Steve Jobs is good advice. But my story is different. I wasn't fired, I just took a short vacation."

It was a joke, but it carried truth. Five days was a miserable vacation, but it wasn't a life-altering exile. For Altman, the episode wasn't about revenge or a triumphant comeback, it was proof of the

resilience of the organization he had built. And that is what made him a genuine leader.

Small Failures Were Training for the Big Crisis.

"I experienced the full range of human emotions. It was quite something." That was how he later recalled it.

On the evening of the day he was fired, Altman was sitting with Greg Brockman at a bar in San Francisco.

Brockman had submitted his resignation moments after being stripped of his chairmanship by the board. The two sat drinking beer, trying to make sense of what had happened. People around them recognized them and took photos. Someone came over and extended a hand. "We're rooting for you!"

On Monday, Microsoft CEO Satya Nadella called Altman. "Come build a new AI research team with us at Microsoft." Altman thought it over.

Starting fresh wasn't a bad option. But OpenAI sat in a corner of his mind. It was an organization he had built over nearly a decade, with colleagues who shared his vision.

On Tuesday, the employees' open letter appeared. More than 700 had signed it. Ilya Sutskever's name was among them. At that moment, Altman understood: 'This isn't my fight alone. It's all of ours.'

In the early hours of Wednesday, word came that he was returning.

The five-day rollercoaster was over. But was it really? Altman didn't think so. The real work was just beginning.

A few months later, at a Sequoia Capital event, Alfred Lin asked: "Sam, can you share some advice on resilience? You've been through an enormous crisis."

Altman paused to think, then said: "It gets easier over time. I mean it."

The audience murmured. Gets easier? A crisis that nearly tore the company apart?

Altman continued.

"As founders, you'll face a lot of adversity. The challenges get harder, the stakes get bigger. But paradoxically, the emotional weight gets lighter. The more bad things you go through, the better you get at handling them."

He drew from his own experience. During his time as president of Y Combinator, he watched hundreds of startups fail. He saw founders weep as they shut their companies down. When building Loopt, he negotiated with carriers thirty times. Twenty-nine of those ended in rejection. In OpenAI's early days, unsure of what to build, the team even considered video games. The GPT-3 API had little practical use outside of copywriting.

"With each crisis, I got a little tougher. The small failures were training for the big ones."

Sam went on. "Early in a startup, losing a single customer can feel like the world is ending. But after hundreds, thousands of rejections and failures, the emotional swings narrow. You don't go numb, you develop the ability to judge what actually matters."

During his Y Combinator years, he used to tell founders: "Most of you will break under the first major crisis you face. But the ones who survive come back stronger for the next one." Resilience was like a muscle. The more you used it, the stronger it got.

Psychology has a name for this: Post-Traumatic Growth. According to the American Psychological Association, resilience is both the process and the outcome of successfully adapting to difficult experiences. It is the capacity to achieve mental, emotional, and behavioral flexibility in response to internal and external demands.

At a Hub71 event in Abu Dhabi, Altman said: "Most people quit in a crisis. They don't have resilience. But the good news is that resilience can be taught. At Y Combinator, that's what we learned, that resilience can be built."

He broke resilience down into three elements. The first is self-belief. "The most successful people I know believe in themselves to a degree that borders on delusion." When he said he would spend a billion dollars scaling GPT, many called him crazy. But Altman was certain: 'This is right. Scale it up and something will happen.'

Blind confidence is dangerous, of course. Altman acknowledged this. "Self-belief has to stay within your area of expertise. I've seen many successful people make terrible decisions outside their domain." AI experts who failed at real estate investment or restaurant ventures were a common example.

The second is adaptability. "What matters is the conviction that you'll solve whatever problem comes your way. You don't have to know how right now. You'll figure it out." When Loopt finally closed a deal with AT&T on the thirtieth attempt, Altman had already learned twenty-nine ways that didn't work. Each failure was data.

The third is a long-term perspective. "Most people think about what they can accomplish in a few years. But what does it look like on a decade or more? The willingness to be completely misunderstood for the next several years, that's a powerful weapon."

Going through the crisis taught Altman one more thing: in a big enough crisis, emotion disappears. "When things get that big, there's no room for feelings to get in the way. What the right choice is becomes intuitively clear." Even through five days of chaos, he planned each next step with calm, keeping the Microsoft option open, watching how employees reacted, and leaving space for negotiation with the board.

A founder asked him later: "How did you stay so calm?" Altman smiled and replied: "Honestly, I think my brain didn't have the bandwidth to process emotions. So many things were happening at once that I just went into autopilot. The grief and the anger came later."

He left founders with one last thought. "Don't avoid small challenges. They're what give you the strength to get through the big ones later. Just like starting with lighter weights at the gym and working up, crises take practice."

Resilience wasn't something you were born with. It was something you built. Sam Altman came back from five hellish days a stronger leader, and shared that experience with thousands of founders.

The Power of Resilience

1 a.m., Wednesday, November 22, 2023. The announcement came: Altman was returning as CEO of OpenAI. Employees cheered. Investors exhaled. The press ran headlines calling it a "dramatic reversal."

It looked like everything was over.

But Altman felt it entirely differently. He reflected later: "A lot of people thought that the moment I came back, it was over. For me, it was another beginning. The truly hard work was starting now."

At the Sequoia event, he was candid: "The aftermath was far harder than the crisis itself." The audience looked surprised. Harder than five days of being fired and fighting to return?

Altman explained: "In a crisis, adrenaline kicks in. You have lots of people supporting you, and a clear goal, something must be done. But by day sixty, day ninety after the crisis has passed, no one is paying attention. That's when you're alone, picking up the pieces and rebuilding. That process is far lonelier and harder."

The first task Altman faced was rebuilding the board. The existing board had effectively collapsed.

Helen Toner, Tasha McCauley, and Ilya Sutskever, all of whom had voted to fire Sam, were gone. Only Adam D'Angelo, CEO of Quora, remained. The empty seats needed filling. But who? The episode had exposed fundamental problems with the board's structure.

Altman and the incoming board members set a few principles. Make the board larger, four members was too few, giving a small group the power to determine the fate of the entire company. Bring in experts from diverse backgrounds, a board consisting solely of AI safety researchers was not balanced. And increase transparency, what had happened, firing the CEO without even notifying Microsoft in advance, must never happen again.

Bret Taylor became chair. A Silicon Valley veteran with stints at Google, Facebook, and Salesforce, he had also served as chairman of Twitter's board and overseen Elon Musk's acquisition of the company. Larry Summers, former U.S. Secretary of the Treasury and a noted economist and policy expert, joined as well. More directors were added later.

Reconstituting the board was only the beginning.

The second task was helping employees recover emotionally. More than 700 people had lived through five days of intense stress: fear that the company might collapse, the weight of feeling responsible for a leader they believed in, the anxiety of preparing to move to Microsoft, and then the relief of a dramatic reversal. It was an emotional roller coaster.

After returning, Altman held all-hands meetings every week. He answered employee questions one by one: "Why did this happen?" "Can it happen again?" "Are we safe?" Some employees remained uneasy. Researchers who had been close to Ilya felt particularly conflicted.

The hardest question was the relationship with Ilya Sutskever. Ilya had played a central role in firing Sam. Within forty-eight hours he had expressed regret and signed the employee letter. Had Sam forgiven him? How would the organization receive him?

Altman did not publicly criticize Ilya. "Ilya genuinely cared about AGI safety. His intentions were good." But tension remained inside the organization. In May 2024, Ilya left OpenAI and founded his own company, Safe Superintelligence Inc. His departure statement included these words: "I am confident that OpenAI will build safe and beneficial AGI under Sam's leadership."

The third task was restoring trust with the outside world. Microsoft was OpenAI's largest investor and partner. It had put in \$13 billion, yet learned of the CEO's firing only at the last moment. Was Satya Nadella furious? Microsoft's stock did fall on the Friday it happened. After returning, Altman called Nadella multiple times to repair the relationship. Nadella wrote on X: "This is a first step toward more stable, informed, and effective governance."

Government regulators took notice. Rumors circulated that the SEC (Securities and Exchange Commission) and FTC (Federal Trade Commission) had opened investigations. A company developing AGI this unstable, shouldn't the government step in? Altman testified before Senate hearings and strengthened OpenAI's safety procedures.

An independent review was commissioned from outside law firm WilmerHale. More than 30,000 documents were reviewed and dozens of people were interviewed. In March 2024, the findings came out. "The cause was a breakdown of trust between the board and Altman. It was not a matter of product safety or finances. Altman's conduct did not justify his removal."

The report made one thing clear. The board had "acted hastily, failed to notify key stakeholders in advance, and did not give Altman an opportunity to address their concerns." New processes were put in place to ensure nothing like this would happen again.

The fifth challenge was rebuilding organizational culture. Before the crisis, OpenAI had been a fast-growing startup. After it, the company had become an AI leader watched by the entire world. The culture needed to evolve too. Altman addressed an all-hands meeting: "The crisis we went through brought one era to a close. We need to become a more mature organization now. But we can't lose the speed and flexibility of a startup."

It was not an easy balance to strike. The team had to stay small and fast while also reinforcing safety and governance. Altman called this "the marriage of speed and accountability."

Altman offered advice to fellow founders. "There's plenty of material on how to handle a crisis. But almost no one tells you how to spend day sixty after the crisis is over. I looked for it and found nothing. That's exactly what you need."

He gave junior founders practical tips. First, expect waves of emotion. Right after the crisis, adrenaline keeps you going, but two or three weeks later, you may fall apart. Second, build small wins. Rather than one grand rebuilding effort, string together a series of small successes: a product launch, a team meeting, a customer call. Give yourself the feeling of returning to normal one step at a time. Third, don't carry it alone. Keep talking to your co-founders, executives, and mentors.

In spring 2024, Altman reflected in an interview. "The five-day crisis was awful. But the six months after it, that was the real test. Reconstituting the board, recovering the team's trust, building new systems. That was leadership in its truest form."

For Altman, the crisis was not an ending but a beginning. Through that long journey of rebuilding, he grew from someone who was only a visionary into a leader who could protect and grow an organization.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

4.1 Philosophy of Organization and Development

Part 4: Altman's Business Philosophy and Strategy

[← Previous](#)

[See all chapters](#)

[Next →](#)

The Master Plan Trap: Revise Your Tactics Quickly.

Sam is not someone who crafts perfect plans. He is wary of mapping things out too tightly.

"You can know exactly where we'll be in ten years? That's a lie."

That's his view. AI technology moves too fast. A brilliant plan today can be scrap paper tomorrow. It's like walking a mountain path in fog, you feel out each step in front of you and find the way as you go.

OpenAI's early days show just how important this approach was. In 2016 they were a team of barely fourteen people. "We honestly didn't know what we were supposed to be doing."

Altman admitted as much later. Early on they were trying to build an AI that could play video games, or a robotic hand. A conversational AI like ChatGPT felt like something from a distant future. What if they had insisted back then: "We're making language models and nothing else"?

ChatGPT as we know it probably wouldn't exist.

Altman's agile approach is a bit like a surfer riding a wave. You read the water, shift your weight from moment to moment. If you feel yourself falling, you turn fast. When a bigger wave rolls in, you go for it.

"We just have to follow where the science leads."

His words. Not chasing targets set in advance by humans, but pursuing the possibilities the technology itself reveals.

OpenAI's shift from a nonprofit to a capped-profit company fits the same pattern. The original dream was a pure research institution. "We shouldn't build AI for money", that was the belief. But recruiting the world's sharpest minds and securing enormous computing power required enormous money.

"Without money, there's no dream."

Altman faced reality squarely and changed the structure. He built a distinctive model: guarantee investors a capped return, then direct everything above that cap toward humanity. If they had clung to the original plan, OpenAI might have already shut its doors.

ChatGPT's origins follow the same pattern. There was a powerful model called GPT-3. "What do we actually do with this?"

The team wrestled with the question. Rather than waiting for a perfect product, they attached a simple chat interface and shipped it. The response was explosive. One million users in five days.

"Try fast, learn fast, change fast."

That's the Altman way. Failure isn't something to fear; it's something to learn from. He doesn't try to find the perfect answer all at once. He finds a sixty-percent answer quickly, then pushes it to seventy, then eighty, improving all the way.

OpenAI's recent release of some open-source models fits the same pattern. The original policy was to keep all technology closed. When powerful open-source models like China's DeepSeek and Meta's Llama appeared, the strategy shifted.

"When the world changes, we change."

Stubbornness is sometimes a poison. Flexibility is a survival tool.

What Altman stresses is that the big goal, the North Star, must not waver, but the path toward it must be free to change at any time. OpenAI's North Star is "developing AGI that benefits humanity." That goal, above all, never changes.

The path changes every day. A road blocked today may open tomorrow. A road that worked yesterday may be closed today.

"Don't cling to complex plans. Solve the problem in front of you."

That's Altman's advice to young founders. He tells them not to spend their time crafting detailed five-year or ten-year plans. Instead: do your best today, and make better decisions tomorrow with what you've learned.

This kind of agility is possible because clear principles underpin it. Start small. Fail fast and learn fast. Listen seriously to users. Validate with data. Turn boldly when the moment calls for it.

"Plans are useless, but the process of planning is important."

That line comes to mind. Altman seems to think along the same lines. What matters is the constant process of thinking and judging, not the plan itself. Reading situations, spotting openings, avoiding risk.

Agile strategy is, at its core, the courage to accept uncertainty. "I don't know what the future holds", admitting that honestly. You can only learn by admitting you don't know. You can only grow by learning.

Altman built OpenAI around that culture. Mistakes are fine. Being wrong is fine. Fix it fast. In that kind of environment, people become bolder. They're not afraid to try something new.

"The best plan is to keep changing the plan."

That paradox is the heart of Altman's philosophy. To survive in the rapidly shifting world of AI, you have to make change itself your strategy. Go with the current, but never lose your direction.

"Working backwards from a complex plan usually doesn't work" , a phrase Sam Altman often uses.

We tend to think this way: "To become CEO in ten years, I need to be VP in nine, and team lead in eight."

It's a way of calculating backward from the future, like marking a destination on a map and then drawing the route in reverse.

Altman says it doesn't work well. In the world of AI, less so than anywhere.

"An AI development plan drawn up five years ago looks like a black-and-white movie to me now."

Technology changes too fast. A perfect plan today is worthless tomorrow. So Altman chose a different approach.

"Growing a little more tomorrow with what I have today."

That's his core belief. He trusts the magic of compounding. A small snowball keeps rolling and eventually grows enormous. Improve by just 1% every day and after a year you're 37 times better. (1.01 to the power of 365 is actually 37.8.)

The early days of OpenAI illustrate this well. In 2016, they considered large language models a "distant dream." Instead, they worked on things like video games and robotic hands.

"We didn't know what we were supposed to be doing."

An honest admission. If they had drawn up a rigid plan back then , "We will build ChatGPT, no matter what" , they probably would have failed. They didn't even know whether that kind of technology was possible.

Instead, they kept running small experiments. They built GPT-1. "Oh, this is better than we expected."

Then they built GPT-2. "Wow, it's gotten even better!"

And then GPT-3. "This is genuinely remarkable."

At each step they learned something. The next step was decided on the basis of what they'd just learned. They didn't plan from start to finish. They chose each next step as they walked.

"We have to follow where the science leads."

Altman's words. He doesn't fight the current of nature. Water flows downhill. Technology does the same , it flows toward where possibility opens up.

The story of how ChatGPT was born is telling. After building GPT-3, they found themselves puzzled. "How are we supposed to use this?"

If they'd had a rigid plan, they would have forced GPT-3 into a predetermined use case. They might have tried to build nothing more than a copywriting tool. But the OpenAI team stumbled onto something unexpected.

"People really love just talking to the model in the Playground."

It was unexpected. They followed that discovery. The conversational feature was "terrible," but people liked it. So they built ChatGPT.

A million people used it within five days of launch. "It felt like a nuclear bomb going off."

Someone's words. If they had set out from the beginning with a rigid plan , "We will build a perfect conversational AI" , they'd probably still be building it. Chasing perfection, they'd have missed the moment.

Altman says, "It only takes one hit, no matter how many times you fail." That's the entrepreneur's way. Try often, fail fast, learn fast.

There's another problem with elaborate reverse-planning. Our knowledge has limits. If you design the future based on what you know today, you end up ignoring everything you'll learn tomorrow.

"Even in situations that look impossible, you can eventually find an answer."

That's Altman's most powerful lesson. It's a realization no one gains from planning everything out in advance. It only comes from accumulating the experience of finding a way forward in situations that feel hopeless.

He speaks about productivity in much the same way. "Choosing the right thing to work on is what matters most." "Think more, and don't rush to start."

People busy drawing up elaborate plans miss what actually matters. You need to spend enough time deciding what to do. Then just do it.

OpenAI's shift from nonprofit to for-profit is the same kind of story. The original plan was a pure research organization. "Let's not think about money."

Reality moved differently. Attracting the best talent required money. Buying enormous computing power required money. If they'd clung to the original plan, OpenAI would have closed its doors.

Altman adapted with flexibility. He created a "capped profit" model. Investors receive returns, but any profit beyond the cap goes toward benefiting humanity.

This wasn't in the original plan. It was a solution invented along the way. No amount of elaborate reverse-planning could have produced it.

"Think in compound terms."

Altman's advice. 1% today, 1% tomorrow, 1% the day after. Small improvements stack up. A year later you arrive somewhere you couldn't have imagined.

In the GPT-3.5 era, apparently 95% of startups bet on models not getting much better. They built businesses around "patching the weaknesses of the current model." Altman warned that this was dangerous.

"When the next model arrives, your company disappears."

He advises building a company that would be happy if the model gets better instead. This is the trap of reverse-planning. If you design the future around the current state of things, you turn future change into an enemy.

The Sydney Opera House comes to mind. Plan in 1957: \$7 million, completion in 1963. Actual result: \$102 million, completion in 1973.

Ten years behind schedule and fifteen times over budget. No matter how carefully you plan, things can go this wrong.

Altman seems to be saying: "Plans are useless, but planning is essential." What matters isn't the plan itself; it's the habit of thinking. The ability to read a situation, make a judgment, and decide.

"Resilience comes from learning that you can get back up after you fall."

When a complex plan falls apart, you feel crushed. "My plan was perfect!" Your ego takes the hit.

When a small attempt fails? "Okay, let's try the next approach." You move on without the weight.

What Altman is really saying comes down to this. Keep the big picture, but stay flexible on the details. Fix your eyes on the North Star, but find the path as you walk it. Don't worry about ten years from now; give everything you have today.

"Follow where the science leads." Don't try to drag it the other way.

Focused on 'a truly intelligent model.' High optimism about the current research roadmap. For Sam Altman, OpenAI has exactly one goal.

"Building a truly intelligent model."

Everything else is secondary. Money, reputation, an impressive product lineup. For him, there is only one thing.

"The most intelligent AI model in the world."

That level of focus is almost frightening.

"We are building genuinely powerful AI, AGI, superintelligence, call it whatever you like."

His words. What does the name matter? The substance is what counts. Building a truly intelligent model. That is OpenAI's reason for existing.

Where did this conviction come from? The 'scaling laws.'

Put simply: scale the model up, add more data, increase computing power. Then AI performance improves in a predictable way.

It sounds like magic. And yet it actually worked.

From GPT-1 to GPT-2. From GPT-2 to GPT-3. Each time, bigger. Each time, smarter.

"The larger the scale, the better it gets, and it gets better in an almost unbelievably predictable way."

That is Altman's central insight. The predictability matters. It let him see what was coming. "This is only going to keep getting better."

So he made a massive bet: spending a billion dollars to scale the GPT models. At the time, it sounded absurd.

"Could this really not be the thing that changes the world?"

Altman was baffled. Why couldn't other people see this potential? But he was certain. He had his conviction. A belief that bordered on delusion.

What was the result? ChatGPT. One million users in five days. The fastest growth in history.

"I was right."

Now his optimism has only grown stronger. When talking about GPT-5, he said this.

"I don't think I'm going to be smarter than GPT-5."

He's saying an AI smarter than himself is on its way. Think about that. An intelligence that surpasses human capability. This isn't just a tool. It's a new kind of colleague.

"The right way to think about the models we're building is not as databases but as reasoning engines."

That's Altman's explanation. Not a machine that memorizes information; one that thinks. The biggest difference with GPT-4 was exactly that: "better reasoning."

About the current research roadmap, he says he's "never been more optimistic than right now." Three elements are in focus: algorithms, data, and computing.

"Algorithmic breakthroughs still carry the potential for 10x to 100x improvements."

There's still a long way to go. It's not just about scaling up. The search is for smarter methods. What if you could make the model ten times more capable with the same computing power?

OpenAI focuses on this one thing. Other companies run many different businesses. OpenAI is different.

"A truly intelligent model. That's all."

We don't spread resources thin. Every dollar, every talent, every hour goes into this one thing. That kind of focus is what makes the difference.

It was the same when GPT-3 was built. 175 billion parameters. A scale no one could have imagined at the time.

"Far more sophisticated language understanding and generation than any previous model."

In an instant, OpenAI became the leader of the AI industry. That is the power of focus. While other companies dabbled in many things, OpenAI did one thing perfectly.

Altman gives startups the same advice.

"Building a business around fixing the minor shortcomings of the current model is dangerous."

Why? Because when the next model arrives, those shortcomings disappear. And so does your company.

"Build a company that benefits as models keep getting better."

Ride the wave of technological progress. Don't fight it. Build a company that cheers when GPT-5 ships. That is Altman's advice.

He is convinced that the AI industry will create "trillions of dollars in new markets." Healthcare, education, finance, law. Every sector will change.

"Products and services that were previously impossible or impractical."

Now they become possible. Once AI models are smart enough, things we never imagined become achievable.

The recent announcement around GPT-5 makes this clear. "It builds software in real time, almost on the spot."

Altman demonstrated it himself. It felt like magic. "It reminded me of the first time I coded, when I was 11," he said.

That is his vision. AI that goes beyond being a simple tool. A partner that amplifies human creativity.

"By the end of 2027, most people will agree that there has been at least one major scientific discovery driven by AI."

That is his prediction. Not just analyzing data. Discovering new theories. Forming new hypotheses. Thinking like a human scientist.

His definition of superintelligence is worth noting. "If an AI can do better AI research than OpenAI's entire research team, that is superintelligence."

Imagine it. AI builds better AI. That AI builds even better AI. Exponential growth. A pace no human can keep up with.

There are problems, of course. During the recent GPT-5 launch, Altman was candid.

"We completely botched a few things during the launch."

It is not perfect. Mistakes happen. Users complained, and OpenAI quickly rolled back to the previous model.

"We needed to build a model that was genuinely smart and useful, while also optimizing inference costs."

A compromise with reality. A larger model was possible. A smarter model was possible. But computing infrastructure constraints forced a reduction in capability.

"We could have built a massive model. A lot of people would want to use it, but in the end we would have let them down."

What good is it if not everyone can use it? Balancing practicality and performance. That, too, is strategy.

Questions about scaling laws persist. "Isn't it hitting a wall?" "The performance gains per dollar are shrinking."

Altman is aware of these criticisms. But he still believes. There is still a long road ahead. Algorithm improvements, data efficiency, new training methods.

"We have better models, but we lack the capacity to serve them."

The problem is not the technology. It is the infrastructure. That is why massive investments like the Stargate project are necessary. It is why \$500 billion is being poured into data centers.

In the end, Altman's optimism is not wishful thinking. It is grounded in data. Rooted in experience. Proven by results.

"A genuinely smart model."

Unwavering focus on this single goal. That is the secret behind OpenAI becoming the best in the world.

The piece being written is getting quite long. With token limits in mind, I will continue with the remaining two subtopics.

Small Team, Big Responsibility: The Secret Behind Fast Decision-Making

Anyone who looks at OpenAI's pace is struck by it. ChatGPT, GPT-4, DALL-E, Sora. All of these came out months apart.

"How are they moving this fast?"

There is a secret. Small teams.

Sam Altman watched thousands of startups during his time at Y Combinator. He almost never saw a large organization beat a small team. If anything, it was the other way around.

"Mediocre teams cannot build great companies."

Hiring more people is not the answer. Hiring a handful of the very best people is far more effective.

"Founders should spend more than 50% of their time hiring their first five employees."

This is not an exaggeration. That is genuinely what it takes. Because those five people become the DNA of the company.

OpenAI follows this philosophy. Each team is small. Five people, seven people, ten at most.

But the responsibility is enormous. They own an entire project. From start to finish.

Take the 'Model Behavior Team' as an example. Their mission: reduce sycophantic responses. This critical work is handed to a small group.

"You are the ones responsible."

Forty people don't crowd into a meeting room arguing over who does what. The team simply handles it. Fast.

"I have never seen someone who is exceptional at solving problems also move slowly."

That is Altman's observation. Outstanding people are fast. They skip unnecessary meetings. They stay focused on what matters.

What is the core problem with large organizations? The cost of communication.

With ten people, meetings are simple. With a hundred? You spend the whole day in meetings.

"One of the biggest mistakes is growing larger without doing more."

That is Altman's warning. Many companies add headcount while productivity stays flat. OpenAI operates differently.

"Everyone should be working hard."

He believes this firmly. A capable executive is a busy executive. If someone has nothing to do, that is a problem.

Small teams carry another advantage: accountability.

Your decisions feed directly into the product. If something fails, it's on you. If it succeeds, the credit is yours. In that kind of environment, people burn with purpose.

"Protect your morning hours as your most productive time."

Altman lives by this himself. A few hours each morning. No one can book that time. It belongs to deep work. To the things that actually matter.

The 2023 firing episode proved this. When Altman was ousted, what happened?

"The company ran perfectly."

It was a moment he took genuine pride in. The executive team held things together on their own. That is what real organizational strength looks like.

"More than 95% of employees demanded his return."

These were not people simply collecting a paycheck. They were colleagues who shared a vision. That kind of loyalty is not built overnight.

When you give small teams enormous responsibility, people grow. They become bolder. They become more creative.

"It's okay to fail."

That is OpenAI's culture. Fail fast, learn fast. In a large organization, failure is a catastrophe. In a small team, failure is instruction.

In the end, the secret to speed is straightforward. The best people. Small teams. Big responsibility. Fast decisions.

This is the engine that makes OpenAI move like a rocket.

The Power of Self-Belief Bordering on Delusion

"The most successful people I know believe in themselves to a degree that borders on delusion."

It's one of Sam Altman's well-known lines. At first hearing, it sounds strange. Delusion?

But he means it.

"Things that look absurd, like spending a billion dollars to scale a GPT model."

At the time, everyone said it was crazy. "Why build a model that large?" "Who's going to use it?"

Altman was certain. "This will change the world."

There was data. He saw the scaling laws. The bigger the model, the smarter it got. Others didn't see it. He did.

"Without sufficient self-belief that this would work, it would have been impossible."

That's his honest admission. Technology alone isn't enough. You need conviction. A belief that doesn't waver even when the world doubts you.

He brought up Elon Musk. Talking about Mars inside the SpaceX factory. The absolute certainty written across Musk's face.

"That is the standard for belief."

That was Altman's thinking. The power to believe the impossible is possible. That's what creates innovation.

"A powerful dream can overcome any obstacle."

That's his philosophy. A weak dream collapses at the first obstacle. A strong dream gets past the tenth.

External resistance is inevitable.

"Sadly, the bigger your ambition, the harder the world tries to knock you down."

That was OpenAI's early reality. "You're going to build AGI?" People laughed. They said it sounded like science fiction.

Altman didn't waver. He had belief. He led the team.

"The power to attract the best people."

Self-belief is contagious. When the leader is certain, the team is certain. When the leader wavers, the team wavers.

"A vision that inspires customers, employees, and investors."

Altman is also a remarkable salesman. He persuades others to share his vision. That's another form of self-belief.

He doesn't give up even when facing the hardest problems.

"Even in situations where there seems to be no solution, you can figure out what to do."

The more experiences like that accumulate, the stronger the belief grows. "I can do this." That certainty makes the next challenge possible.

There is a condition.

"The more data points you accumulate showing that your judgment is sound in a particular domain, the more you should trust yourself in that domain."

You can't carry delusional confidence across every area of life. That would be genuine delusion. It only works within your own domain of expertise.

The decision to spend a billion dollars scaling an AI model. That was Altman's domain. A deep understanding of technology and vision.

Carrying that same certainty into unrelated fields like politics or biology is dangerous.

"Objective self-awareness and a willingness to accept criticism."

Altman stresses this as well. In the past, he disliked criticism. Now he takes it as part of the process of finding truth.

"Accepting legitimate criticism is hard and painful, but it's essential for distinguishing belief from delusion."

The line between self-belief and arrogance. Altman walks it carefully.

"I will fail many times, and I will be truly right just once."

That is the entrepreneur's way. You try many times. Most attempts fail. You only need to succeed once.

Back at Loopt. He tried 30 different approaches to land a contract with a carrier. He kept getting turned down.

"But I didn't give up."

Where does that persistence come from? Self-belief. "I know I can pull this off in the end." That conviction is what makes the 30th attempt possible.

"It comes from learning that you can get back up after you've been knocked down."

Resilience. The more failures pile up, the stronger you become. The more bad things you go through, the lighter the emotional weight of each one.

The November 2023 ouster. It was a shock to Altman too. Confusion, frustration, anger, grief.

"I experienced the full range of emotions."

Yet he didn't collapse. He was back within four or five days. That is the power of self-belief. "I know I'm doing the right thing."

At its core, Altman's self-belief is grounded optimism. It rests on expertise and experience. It is backed by data.

And yet he stays humble. He takes criticism seriously. He learns from failure.

"Trust yourself, but don't ignore reality."

That is the most essential source of Sam Altman's strength, the force that turned GPT into reality and now steers OpenAI's bold voyage toward the future of humanity.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

4.2 OpenAI's Future Business Model

Part 4: Altman's Business Philosophy and Strategy

[← Previous](#)

[See all chapters](#)

[Next →](#)

Becoming Humanity's Core AI Platform

Sam Altman carries one picture in his head, simple in shape and vast in scale.

"A world where AI is as natural as electricity, water, or air."

He held this dream from the moment he founded OpenAI. Just as people reach for their smartphones the instant they wake up, he believed the day would come when AI would be the first thing people turned to. Something you ask about the morning news on your commute, debate your lunch options with, help your kids with homework in the evening, and use to prepare for tomorrow's meeting at night.

That is the future OpenAI is working toward. Not just a company that makes a lot of money, but a platform that stands at the center of people's lives.

In an interview, Altman put it this way: "We want to be people's core AI subscription service. Like the operating system of a smartphone."

That might sound abstract. An operating system? An AI subscription? What does it actually mean?

Put simply: whether you use an iPhone or a Galaxy, there is an operating system underneath, iOS or Android. Every app runs on top of it. KakaoTalk, YouTube, games, none of them work without an operating system.

Altman imagines the same thing happening with AI. The AI system OpenAI builds would become a kind of "intelligence operating system," and countless services and apps would run on top of it.

Whether you are booking a hospital appointment, studying English, or planning a trip, everything flows through OpenAI's AI. And people pay a set amount each month to subscribe, just like Netflix or Spotify.

"Why does the subscription model matter?" Altman explains it this way. "When people pay every month, that's proof they actually need the service."

In practice, ChatGPT crossed 100 million monthly active users within two months of launch, the fastest growth rate in history. The paid subscription service, ChatGPT Plus, attracted millions of subscribers almost instantly.

Why do people pay to use AI? The answer is straightforward. It saves time. It helps them do their work better. It turns complicated problems into manageable ones.

One office worker said: "ChatGPT cut my work hours in half. From writing emails to putting together reports, I don't do it alone anymore."

A high school student says: "ChatGPT explains math better than my teacher. It keeps trying different ways until I understand."

As those experiences accumulate, people find it harder to imagine life without AI. Just as a power outage makes everything inconvenient, a world without AI will feel like something is missing.

Altman's vision goes a step further. It doesn't end with ChatGPT. He wants to build a world where a single OpenAI account is all you need.

Just as a Google account lets you log into countless sites, an OpenAI account would connect you to every AI service in the world. A personalized AI that knows your tastes, habits, and memories would stay with you for life.

"Personalization is the key." It is something Altman stresses. "An AI that knows only you, an AI that understands your context, that is what a truly valuable AI looks like."

ChatGPT already saves conversation history and learns user preferences. Going forward, it will come to know your schedule, health records, reading list, and hobbies as well. And from all of that, it will offer the most relevant advice.

That might sound alarming. One company knowing everything about you. Altman is aware of this concern. "Privacy is the top priority. Users must be the ones in control."

OpenAI has promised to let users delete their data at any time and to be transparent about how that data is used.

The platform strategy is equally ambitious. In 2024, OpenAI opened the GPT Store, a marketplace where developers can buy and sell custom GPTs, much like the Apple App Store.

A cooking recipe GPT, a legal advice GPT, a math tutoring GPT, a travel planner GPT, the possibilities are endless. Anyone can build their own AI and earn money from it.

That is the power of a platform. OpenAI does not need to build everything itself. Developers around the world do it for them. OpenAI only needs to provide the stage.

Altman's goal is clear. He wants to turn OpenAI into the Microsoft or Apple of the AI era. Or maybe something even larger than either.

"What we want to build is not just another product. It is the infrastructure that will underpin a new era."

His words carry real conviction. It is not an exaggeration. ChatGPT is already used by more than 300 million people every week around the world. Business customers alone number over one million.

Those numbers keep rising. The speed at which OpenAI is embedding itself into the core of people's lives is beyond anything easily imagined.

APIs and SDKs: Building the Developer Ecosystem

Altman learned one important lesson during his time at Y Combinator.

"Build an API and something good happens."

What does that mean?

API stands for Application Programming Interface. Think of it as a set of LEGO bricks. It is a connector that lets other people take the AI OpenAI has built and plug it into their own programs.

Take an example. Say you want to build an English learning app. Building your own AI from scratch would take years and cost billions of won. You would need researchers, and enormous computing power.

What happens when you have an API? Connect OpenAI's GPT API with a few lines of code, and that's it. Your app suddenly has a world-class AI English tutor built right in.

Development time shrinks from years to days. Costs drop from billions of won to a few hundred thousand won a month.

That's the magic of an API.

Altman recognized this power early. So when OpenAI built GPT-3, the very first thing they did was open up the API. The API came before ChatGPT.

"We didn't know exactly what to build," Altman says candidly. "So we let other people build it."

The results were remarkable. Tens of thousands of startups around the world began building new services on the GPT API.

Copywriting companies built tools that generate ad copy automatically. Customer service firms built AI agents that work around the clock. Education companies built AI tutors that explain things in a personalized way for every student.

Someone built a service to analyze legal documents with AI. Someone else built a tool to organize medical records. Another person built a system for conversational game characters powered by AI.

The possibilities are endless. If you have an idea, you can build almost anything.

Today, more than three million developers worldwide use the OpenAI API. The services they've built are used by billions of people.

Altman doesn't stop there. The next step beyond the API is the SDK, short for Software Development Kit. If the API is a set of LEGO bricks, the SDK is the complete package: bricks, blueprints, and a full toolkit.

In late 2024, OpenAI announced the Agents SDK, a toolkit that makes it easy for developers to build AI agents.

What is an AI agent? It's an AI that thinks and acts on its own. Rather than just answering questions, it handles multi-step tasks from start to finish.

Say you ask, "Plan my trip to Jeju Island next week." The AI agent searches the web to check the weather, finds flights, books a hotel, puts together a list of restaurants, and delivers a complete itinerary. All you have to do is approve it.

Thanks to the SDK, building these kinds of sophisticated AI agents has become far easier. Now even a small startup, or even a high school student, can build one.

Altman's strategy is clear. "We can't do it all on our own. The world is too big and the problems are too many. Developers have to build it together with us."

This strategy is also brilliant from a business standpoint. OpenAI only needs to build the underlying model. Developers handle everything else themselves. OpenAI just collects API usage fees.

It's a win-win. Developers get access to powerful AI at low cost, OpenAI earns steady revenue, and users enjoy a wide range of AI-powered services.

That's what an ecosystem looks like. Not a single tree growing alone, but an entire forest growing together.

Altman understands exactly how important this ecosystem is. "A platform's success is determined by how many developers build things on top of it."

The iPhone didn't succeed because of apps Apple built. It succeeded because of the App Store, built by millions of developers. Windows didn't succeed because of programs Microsoft wrote. It succeeded because countless software companies built programs for it.

OpenAI is following the same path.

The future Altman envisions looks like this: behind every app, every website, every service, OpenAI's AI is running. Users may not even know it. On the surface, it looks like someone else's product, but underneath, OpenAI's API is doing the work.

Like HTTP for the internet, it becomes the foundational protocol for AI.

"We want to be the AI infrastructure," Altman says. "Developers should be able to trust us and build their companies on top of us."

The Better the Model, the Better the Business

Sam Altman sometimes says things that sound almost strange.

"AI will create an unimaginable amount of wealth in the world."

Unimaginable , but what does that actually mean?

The 'wealth' Altman talks about isn't just money. It includes time, health, knowledge, opportunity, happiness , all of it.

He explains it this way: "Wealth is the ability for people to get what they want."

Think about the world before AI arrived. Getting a good education meant going to a good school. Getting quality medical care meant going to a good hospital. Hearing expert advice meant paying a lot of money.

All of these were scarce. Good teachers are finite. Excellent doctors are finite. Time is finite.

AI shatters that scarcity.

Anyone can now have a world-class tutor by their side. All it takes is a smartphone. Anyone can get medical advice for free. Anyone can access expert-level knowledge.

This is where 'wealth' gets created.

Consider a single mother. She's raising three kids and working at the same time. There's no money for tutoring. In the old days, she couldn't have given her children a decent education.

What about now? ChatGPT teaches her kids for free. Math, English, science , they can ask about anything, any time of day.

That is wealth creation. No money changed hands, but this family became far richer. They gained something of real value: education.

Take a small company with five employees. They can't afford to hire a marketing person. In the past, they would have simply given up.

Now? AI designs the marketing campaign, writes the ad copy, and drafts social media posts , at 1/10 the cost of a full-time marketer.

The company grows fast. It hires more people. It creates more value.

Altman believes this kind of story will play out hundreds of millions of times around the world.

"The cost of intelligence approaches nearly zero," he says. "So what will people do with that intelligence? They'll build things we can't yet imagine."

Think about a scientist. To conduct research, she has to analyze data, survey the literature, and design experiments. That work takes years.

What if AI helps? Data analysis finishes in hours. A literature review takes minutes. AI optimizes the experimental design.

The scientist can focus on what truly matters: generating new ideas. The pace of scientific progress becomes 10 times faster, 100 times faster.

Cancer treatments arrive sooner. Solutions to climate change come faster. Energy breakthroughs happen more quickly.

That is wealth for all of humanity.

In 2021, Altman wrote an essay called 'Moore's Law for Everything' , a law of accelerating returns applied to everything, not just chips.

Moore's Law holds that the performance of computer chips doubles every two years. Altman argues that because of AI, the value of everything will grow at that same exponential rate.

He puts specific numbers on it: "AI will add trillions of dollars in value to the global economy."

Trillions of dollars? What does that actually mean?

The U.S. GDP is about \$25 trillion. South Korea's is roughly \$1.7 trillion. Trillions of dollars is the size of an entire country's economy.

Experts agree. McKinsey projects that generative AI alone will create \$4.4 trillion in value each year. PwC estimates that AI will add \$15.7 trillion to global GDP by 2030.

Where does all that money come from?

Productivity gains, for one. You can get more done in the same amount of time. One person does the work of ten.

New industries, for another. AI will give rise to jobs and companies that don't exist yet , the way the internet produced Google, Facebook, and Amazon.

Cost reduction as well. Healthcare, education, legal services , the cost of hiring specialists drops dramatically.

Altman's vision is to keep this wealth from flowing only to a small few.

"AI must reduce inequality," he says, with conviction. "The wealthy already have everything. AI will give far greater benefits to people who have less."

That's why OpenAI offers ChatGPT for free , so that anyone, regardless of income, can access the world's best AI.

That's also why Altman has proposed an idea called Universal Basic Compute , distributing AI computing resources to everyone. Instead of money, you give people access to AI.

"Wealth in the future will be determined by how well you can use AI," Altman says. "So everyone needs to have access to it."

This is what OpenAI stands for , a goal that goes beyond the company's own profit, aimed at the prosperity of all humanity.

Is it all just nice-sounding talk? No. OpenAI is already showing it in action. Hundreds of millions of people use the free version of ChatGPT , students in developing countries, elderly people in rural areas, people with disabilities, all of them, for free.

That is wealth redistribution.

Instead of fixing small flaws, focus on what is fundamental.

Altman saw thousands of startups. Back in his Y Combinator days.

The ones that succeeded had things in common. The ones that failed had patterns too.

In the age of AI, he has a clear warning for startups.

"Don't build a company that patches the weaknesses of today's AI models."

What does that mean?

ChatGPT today isn't perfect. It sometimes gives wrong information. It makes mistakes with math. It doesn't know the latest news. It can't generate images flawlessly.

Many founders notice these flaws. And they think: "Okay, I should build a service that fixes ChatGPT's math errors!"

Altman is blunt about it. "Don't. We'll solve that problem within six months."

And that's exactly what happens. When GPT-4 arrived, many of GPT-3's weaknesses disappeared. When GPT-4o came out, a whole new set of problems got resolved.

The models keep getting better. Fast. Exponentially.

If you build a company that "fills the gaps in today's model," you'll be obsolete within months. The moment OpenAI releases an update, it's over.

There's actually a meme out there: "OpenAI killed my startup." It means your company was wiped out by an OpenAI update. Sad, but true.

Altman's advice is straightforward. "Build a structure where your company gets better as the model gets better."

Let's look at some examples.

Bad strategy: "GPT-4 can't write legal documents perfectly. So I'll build an AI that specializes in legal drafting."

The problem: when GPT-5 arrives, it writes legal documents perfectly. Your company is finished.

Good strategy: "I'll build a platform that helps lawyers serve their clients better using AI. The stronger the model gets, the more powerful our platform becomes."

Do you see the difference?

The first bets on AI's weaknesses. The second bets on AI's strengths.

Altman gives a concrete example. Think about an AI tutoring company.

Bad approach: The current GPT can only teach sixth-grade math. So we build a tutor specialized in sixth-grade math.

Good approach: Build a learning platform on top of GPT. Right now it teaches sixth grade, but when GPT-5 arrives it teaches high school, and when GPT-6 comes it reaches college level. Our platform keeps growing.

"Ride the wave of model improvement." That's Altman's core advice. "Don't fight the wave. Get on top of it."

He often puts it this way. "We show up every morning to make the model better. If you're building a company that will be sad when GPT-5 comes out, something's wrong. Build a company that will be thrilled when GPT-5 arrives."

This advice isn't just for AI startups. Every company should take it to heart.

Running a restaurant? As AI improves, order management, inventory, and customer service all get better.

Running a clothing store online? As AI improves, customer preference analysis, inventory forecasting, and marketing all become more precise.

Whatever business you're in, you need to build a structure where AI improvements work in your favor.

Altman has seen it countless times in Silicon Valley. Startups that only target gaps left by the big companies. Attack Google where it's weak, do what Facebook won't do..

It works in the short term. Things go well for a few years. But the moment the big company adds that feature? It's over.

"Build a defensible business." That's his advice. "Build your company on top of the model's strengths, not its weaknesses."

There's another piece of advice. "Solve the hard problems first."

Easy problems get copied quickly. Anyone can replicate them. Hard problems are different. They take time. They require hard-won know-how.

Solve the hard problem first and claim the market early, and you can hold your competitive edge even as the models improve.

This is Altman's final piece of advice. "Bet on the future. Don't bet on the present."

Five years from now, ten years from now, AI will be 100 times more powerful than it is today. Imagine that moment. Think about what kind of company the world will need then. Build that company now.

The Birth of a Trillion-Dollar AI Economy

Let's look at the numbers.

In 2023, the AI market was worth roughly \$150 billion.

What about 2030? Experts project \$1.8 trillion. Twelve-fold growth in seven years.

That's the conservative estimate. Some analysts think it'll be far larger.

Altman puts it this way: "The numbers don't matter. It's just going to be enormous."

Where does all that value come from?

First, the enterprise market.

Every company uses AI. Already, 65% of American businesses have adopted it, and that share keeps climbing.

AI for marketing. AI for customer service. AI for accounting. AI for human resources.

According to McKinsey, generative AI alone will create between \$2.6 trillion and \$4.4 trillion in value each year. That's larger than the UK's entire GDP.

Second, the healthcare industry.

AI diagnoses disease, develops new drugs, assists in surgery, and monitors patients.

The healthcare AI market alone is expected to reach \$188 billion by 2030.

Third, education.

Students around the world are using AI tutors. Personalized learning is becoming the norm.

The AI education market is projected to exceed \$30 billion by 2030.

Fourth, entertainment.

AI makes films, builds games, and composes music. A solo creator can now produce Hollywood-caliber content.

The creative economy is exploding.

Fifth, manufacturing.

Robots run factories. AI optimizes designs and manages supply chains.

Productivity rises dramatically.

Sixth, finance.

AI advises on investments, evaluates loan applications, and detects fraud.

The financial AI market is forecast to reach \$64 billion by 2030.

Seventh, new kinds of jobs.

Prompt engineers, AI trainers, AI ethicists, AI translators...

Thousands of jobs that don't exist today will be created.

Eighth, gains in productivity.

Every worker using AI saves two to three hours a day, and puts that time toward creating more value.

PwC estimates that AI will add \$15.7 trillion to global GDP by 2030: \$3.7 trillion in North America and \$7 trillion in China.

Hard to grasp how large \$15.7 trillion is? It's more than half of U.S. GDP. Nine times South Korea's GDP.

Altman sees something larger still than any of these figures.

"The real value," he says, "is in the things you can't measure."

Healthier lives, better education, more time, greater creativity. None of that can be priced.

What if a mother saved her sick child thanks to AI? Can you put a price on that?

What if a scientist found a solution to climate change thanks to AI? What would that be worth?

What if a student realized their dream thanks to AI? Can that be expressed in numbers?

"The real value of AI is that it frees human potential." It's something Altman says often.

Until now, humanity has lived within limited resources. There were never enough good teachers. Never enough experts. Never enough time.

AI removes those constraints. Now anyone can have the best possible advisor at their side. Twenty-four hours a day, for a lifetime.

The economic value this will create is beyond imagination.

There are real challenges, of course. Jobs may disappear. Inequality may deepen. Privacy may be violated.

Altman knows this. "Technology is not neutral. What matters is how we use it."

That is why OpenAI invests in safety. In ethics. In fairness.

"The wealth AI creates must reach everyone." That is his firm belief.

A new market worth trillions of dollars is opening up. Who will emerge as the winners?

Altman has an answer. "The people who use AI well. Whether that's a country, a company, or an individual."

It's time to prepare. The AI revolution has already begun.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

4.3 Productivity and Work Management

Part 4: Altman's Business Philosophy and Strategy

[← Previous](#)

[See all chapters](#)

[Next →](#)

Focusing on the Right Problem Is What Matters Most.

Sam Altman's starting point for productivity is a little unusual.

When most people define productivity by how many items they crossed off their to-do list in a day, Altman asks the opposite question: "Is what I'm doing right now actually what I should be doing?" That single question runs through every part of his work philosophy.

He wrote this on his blog in 2018: "It doesn't matter how fast you move if it's in the wrong direction. Picking the right thing to work on is the most important element of productivity, and most people ignore it."

As Altman sees it, most people care only about staying busy. They go from meeting to meeting, reply to emails, handle urgent tasks, and tell themselves they worked really hard today. What they miss is the more important question: were any of those things actually significant, for their life, their company, or the world?

Think back to elementary school. The night before an exam, one student sets a goal of solving 100 practice problems. Another does only 10, but those 10 target the exact areas where he's weakest. Who is being smarter?

Altman is the second type. He's obsessed with "doing the right thing" over "doing a lot of things." Whether running Y Combinator or building OpenAI, he held to that same principle: "Will this still matter in ten years?" "If I solve this problem, will it unlock solutions to many others?"

He deliberately blocks off "thinking time" in his schedule.

That means reading books, meeting interesting people, walking in nature. It looks like not working, but it's actually the most important work of all. You have to clear your head before you can see what the next big problem is.

This approach becomes a survival strategy in fast-moving fields like AI. New technologies appear every day, competitors make announcements, the press makes noise. If you tried to respond to all of it, you'd never be able to focus on what actually matters: building AGI.

Altman also emphasizes the importance of independent thinking. The goal isn't to pursue what everyone else calls important, but to find what you genuinely believe matters. He says, "The most impressive people I know have strong convictions." Those convictions, though, must come from serious thought, not from following the crowd.

What's interesting is that Altman didn't always think this way.

When he was younger, he believed he had to handle every problem himself. He attended every meeting, replied to emails right away, and took on every request. Then it hit him: "This isn't working. I'm not spending time on the things that actually matter."

After that, his way of working changed completely. He set a new standard: "If the world doesn't stop turning when this problem goes unsolved, it's not the problem I should be focused on."

That philosophy connects directly to OpenAI's success.

When OpenAI started in 2016, many people pitched ideas: "What about using AI in video games?" "What about researching robotic hands?" They all sounded exciting. But Altman and his team kept asking: "Is this the fastest path to AGI?"

In the end, they chose to concentrate on language models, the path that led through GPT-1, GPT-2, and GPT-3. If they'd chased every project that looked interesting at the time, ChatGPT might never have existed.

There's one more trap Altman warns against. He calls it "productivity porn."

People spend too much time choosing productivity apps, decorating their to-do lists, and studying time management techniques, all while not doing anything that actually matters. It's like spending so long getting ready to study that you never actually study.

He's blunt about it: "Don't spend time building the perfect system. Just make sure you're focused on the right problem."

The same principle applies to students.

Imagine a middle school student with two weeks until exams. Which subject should she study first?

Most students start with "the easy stuff" or "what I like." The Altman approach is different. The first question is: "Which subject, if I raise my score in it, will have the biggest impact on my overall grade?" If math is sitting at 50 and everything else is above 80, the answer is obvious.

"Choosing what to work on" is where productivity begins. Altman's success didn't come from working more hours than everyone else. It came from asking a different question: "Is this really what I should be doing?"

Ask that one question seriously, and your life can look completely different.

Starring the Items That Matter

You might assume that the CEO of the world's most influential AI company uses cutting-edge work management software.

Wrong.

Sam Altman uses paper and pen.

In his 2018 blog post, he wrote: "I highly recommend using lists. I make lists of what I want to accomplish by year, by month, and by day."

That's it. No fancy apps, no elaborate tagging systems. Just write your tasks on paper and put a star (★) next to the few that truly matter.

When you first hear it, you might think, "Isn't that too simple?" But the secret is in that simplicity.

For one thing, a list clears your head.

The brain isn't a storage unit. Keeping things in memory burns an enormous amount of energy. Carrying around thoughts like "I need to prep for the meeting, reply to that email, finish the report" weighs you down.

Altman dumps all of it onto paper. His head gets lighter, and he can focus on the work that actually demands thought. "Having a list helps with multitasking," he says. "You don't need to keep as much in your head."

Second, the star system is the key.

Say you have ten things to do. They all look important. But are they really all equally important? No. Altman is honest about it. "Some things genuinely matter, and some matter less. Let's be straight about that."

So he stars only the three or four items on his list that truly matter. At the start and end of each day, he checks only those starred items. The rest? He gets to them if time allows, and pushes them to the next day if it doesn't.

The power of this approach lies in what it forces you to choose. If you star all ten items, you might as well have starred none. Altman's system keeps asking the same question: "What actually matters? If you could only pick one thing to do today, what would it be?"

Third, the power of paper.

Everything is digital these days. Yet Altman insists on paper. Why?

He rewrites his paper list often. He looks at what he wrote yesterday, then copies it onto a fresh sheet today. The process naturally reshuffles priorities. "This seemed important yesterday, but looking at it

now, not so much." "I keep pushing this one back , do I actually need to do it?"

The act of writing by hand activates the brain. Things stick better than when typed, and it helps clarify thinking. Research shows that handwriting improves cognitive performance.

Fourth, it builds momentum.

Altman says, "The more you complete, the better you feel, and that makes you want to do more." Every time you cross something off the list, you get a small hit of satisfaction. Stack enough of those, and you start believing you can do it, and you want to take on the next thing.

Cross off a starred item? That feels electric.

He has said, "I like starting and ending the day with something that actually moves things forward." First thing in the morning, he finishes one of the hardest starred items. At that point, the day is already a success. Everything else is a bonus.

Fifth, you can choose based on how you feel.

Having a list brings another advantage. Altman says, "If I don't feel like doing a particular task, I can always find something else on the list that's interesting."

Say it's Monday morning and you can't face writing that report. You don't have to force it. Look at your list. Is there another starred item you'd rather tackle right now? A code review? Prep for the strategy meeting?

Pick any of them and get started. What matters is moving starred items forward, not working through them in order.

Sixth, complexity isn't needed.

Altman is clear about it. "I don't categorize anything, and I don't estimate task sizes or anything like that. All I do is put a star next to the items that matter."

Many productivity systems demand things like "A, B, C priorities," "urgent/important matrices," and "time estimates." But the categorizing itself eats time. Altman thinks that time is better spent doing actual work.

A star is enough.

Students can use this approach right away.

You have five assignments. Write them all on one sheet of notebook paper. Star the ones due tomorrow morning and the ones covering material that will be on the exam. That's probably two or three.

Tonight, do only those. The rest? Handle them after the starred items are done, if you have time left. If you're tired, go to sleep. The starred items are already finished.

The reason this system works is that it's simple. Complex tools look appealing at first, but they quickly become a burden. Setup, categorization, maintenance: they eat up your time.

Paper and a star are different.

You can make the list in thirty seconds. You can add the stars in ten. Then you're already working. The tool stays out of your way.

That's why Altman has stuck with this method for over ten years. As the company grew, as responsibilities piled up, as his calendar exploded with meetings, his paper list stayed the same. Because it actually works.

What is the point of a productivity tool?

"To get more done, and to do it better." The tool must never become the goal itself. Altman's paper and star hold to that principle. Not flashy, but powerful.

Right now, pull out a sheet of paper. Write down five things you need to do. Star the two that truly matter.

Then do only those.

Tomorrow morning, you'll be surprised. "Wait, I actually got something done."

Protecting the Morning: The Most Productive Deep Work

There is one block of time in Sam Altman's day that is treated as sacred.

The morning.

He wrote on his blog: "The first few hours of the morning are definitely my most productive. So I don't let anyone schedule anything during that time."

Not a preference. A strategy.

Altman usually wakes up between 6 and 7 a.m. The first thing he does is drink a large espresso. He rarely eats breakfast. He practices intermittent fasting, going 15 hours without food.

Most importantly: he doesn't check email. He doesn't look at messages. He doesn't read the news.

What does he do instead?

Deep work. The hardest tasks that require deep concentration.

Let me explain what deep work means. Imagine you're solving a math problem and a friend keeps interrupting you. You can't focus. If your concentration breaks every three minutes, you'll never solve a hard problem.

Deep work is the opposite. No one interrupts you, no notifications, complete immersion in a single task.

This is how Altman spends his mornings.

He makes difficult strategic decisions. "What direction should the next version of GPT go?" "Should we accept this partnership?" "What needs adjusting in the company's five-year plan?"

These thoughts can't be done "in spare moments." The brain needs to be operating at 100%.

There's scientific backing for this.

The brain recharges while you sleep. When you wake up, your willpower battery is full. The ability to make hard decisions, think creatively, and solve complex problems is at its peak.

As the day goes on, that battery drains. Every decision you make, every person you meet, every email you read takes a little more away. By evening? Nearly empty. That's why all you want to do is watch Netflix.

Altman is a master of managing this battery.

He spends the most precious morning energy only on what matters most. Easy tasks, mechanical tasks, things others have requested, all pushed to the afternoon.

His schedule makes this clear.

Morning: solo deep work, strategy, research review, important writing. Afternoon: meetings, email, calls, people.

The type of work in the morning and afternoon is completely different.

He says frankly, "By afternoon, everything falls apart and all sorts of things come up." As CEO, urgent problems keep appearing. Someone calls out of nowhere, a critical issue surfaces, paperwork piles up.

This can't be avoided. It's the fate of a CEO.

Altman's approach is: "If you can't avoid it, contain it to a time slot." The afternoon is already chaos. So he loads that chaos into the afternoon. But the morning he guards absolutely.

He says firmly, "I don't let anyone schedule during the first few hours." Even if a close colleague asks for "just a quick meeting at 10," he declines. "Sorry, would afternoon work?"

This sounds easy but is genuinely hard. Others may not understand. "Why not in the morning? It's just an hour."

Altman doesn't waver. Because he knows this principle is the core of his productivity.

Altman also pushes email checking to later in the day.

Why? If you read email in the morning, other people's requests fill your head. "I need to reply to this, I need to check that." Then you can't focus on the important work you set out to do.

He uses a dedicated LED light just for checking email. Each morning he sits in front of it for 10 to 15 minutes and briefly scans his inbox, just enough to see if anything is urgent. Replies wait until the afternoon.

The LED light habit is interesting too. Research shows that bright light exposure in the morning regulates your circadian rhythm and keeps you more alert throughout the day. Altman says, "This is the single highest-impact thing I recommend."

The 90-minute cycle matters too.

Altman focuses intensely for about 90 minutes, then takes a short break, a walk, some stretching, or a few minutes of doing nothing.

Why 90 minutes? Scientists have found it's the optimal window for concentration. Go beyond 90 minutes without rest and efficiency drops. But break too often and you never get deep.

During breaks he doesn't look at his phone. Instead he walks outside or stares out a window. He gives his brain genuine rest. A phone isn't rest. It's just another kind of stimulation.

Meetings go in the afternoon.

Altman packs meetings into the afternoon as much as possible. He keeps them short, 15 to 20 minutes, or goes long at two hours. He's not a fan of one-hour meetings. "They often just waste time."

Fifteen minutes is enough to make a quick decision and wrap up. Two hours is enough to dig deep into a complex problem. But one hour? It's awkward. By the time you prepare and get started, thirty minutes are gone. By the time you wind down, the rest is gone too.

All these principles come down to one thing.

"Spend your most valuable time on the most important work."

Altman's mornings are the secret to his success. The strategies behind building ChatGPT and pushing toward AGI almost certainly came from those morning hours, not from conference rooms. They came from the quiet morning, those few uninterrupted hours.

We can start the same way tomorrow.

Try protecting just the first hour of your morning. Leave your smartphone in another room. Use that time to finish the one hardest thing on your list.

A month later, you'll almost certainly feel it. "Wait, something about me has changed."

The Art of Saying No: "I'm sorry, but I can't make that work right now."

In Sam Altman's productivity philosophy, there is one last skill, and the hardest.

Saying "No."

He wrote honestly on his blog: "I try to be ruthless about not doing work and getting to-do's off my list as fast as possible."

He used the word "ruthless." That's strong language, isn't it?

But it's necessary. Because the world never stops trying to take our time.

From the moment Altman became CEO of OpenAI, dozens, sometimes hundreds, of requests came flooding in every single day.

"Would you speak at this conference?" "Can we grab coffee for thirty minutes so I can pick your brain?" "Could you take a look at our project?" "Just one TV interview."

These all seem like perfectly reasonable requests. Turning them down makes you feel like the bad guy, rude even.

Altman says no anyway.

Because he knows what his time is worth. His time must go toward "building AGI." Everything else ranks below that.

He says, "Meetings and conferences carry an enormous time cost. So I try to avoid them whenever possible."

People might get the wrong impression. "Altman must be arrogant. He doesn't want to meet anyone."

No. He says no not out of arrogance but out of wisdom.

Let's do the math.

If he agrees to one conference talk, here's how many hours it actually costs: five hours preparing the talk, three hours traveling, two hours for the talk and Q&A, two hours for networking. Twelve hours in total.

What could Altman do with 12 hours?

He could map out the roadmap for the next version of GPT. He could review an important partnership deal. He could have a deep discussion with the research team.

Which would have a greater impact on the world?

The answer is clear.

Altman also says, "90% of random meetings are a waste. But 10% are worth it."

It's an interesting perspective. He knows most meetings are useless. But he doesn't close the door entirely, because of that 10%.

Because sometimes a big idea comes from a chance encounter. Over coffee with someone, you might suddenly think, "Wait, this is genuinely revolutionary."

So Altman doesn't shut people out. He chooses strategically. He deliberately leaves space in his schedule for "chance encounters." But even that space is managed with intention.

There are concrete ways to say no.

First, set a clear standard. "Does this directly serve our core mission?" If the answer is no, decline.

Second, offer an alternative. "I can't do it personally, but someone on the team might be able to help." This way you preserve the relationship while protecting your time.

Third, don't delay. If you're going to say no, say it quickly. Stringing someone along with "Let me think about it" for two weeks before finally declining is worse. It's more considerate to let them know early so they can make other plans.

Knowing the value of your time matters.

Altman gives this example: "I can't understand why someone who earns \$100 an hour would spend two hours just to save \$20."

What does that mean?

Say you wash your own car and save 20,000 won. But it takes two hours. If you'd spent those two hours on something more important, you could have created 200,000 won in value.

Altman also warns: "Don't pursue productivity for its own sake." Many people obsess over doing more, faster. Crossing items off a to-do list becomes the goal.

So what's the real goal?

"Finishing what actually matters."

If you do 100 things in a day but none of them really matter, it means nothing. You're better off doing just three things that truly matter.

That's why saying no is necessary. You have to decline 99 less important things so you can focus on the 3 that truly matter.

Let's look at a real example.

In 2023, Altman was briefly ousted as CEO and then returned. He was under enormous pressure. Media interview requests flooded in, investors wanted meetings, and employees wanted explanations.

What did Altman do?

He did only what truly mattered. Conversations with employees. Negotiations with the board. Stabilizing the company. Everything else he either declined or put off for later.

As a result, the situation was resolved in four or five days and he returned. If he had responded to every request? He probably still wouldn't have sorted things out.

Students can practice this too.

To a friend who wants to game: 'Sorry, I have to study for a test today. Let's do it this weekend.' To a dull gathering: 'I'll come next time. I'm a bit tired today.' For social media notifications: switch to silent mode.

It feels awkward at first. You feel bad saying no, and you worry it will damage your relationships.

But something surprising happens.

Over time, people understand. 'Oh, they focus on their own work.' They actually respect you for it.

Friends figure it out too. 'Don't bother asking them. But when they do commit to something, they really deliver.'

That's Altman's reputation.

He doesn't say yes easily. But once he says yes, he sees it through. People know that. So they trust him more.

Ultimately, the art of saying no isn't really about saying no well.

"It's about saving your yes for what counts."

Altman says yes only to what truly matters in his life. Building AGI. Growing his team. Changing the world.

Everything else gets a polite but firm no.

His success may have come less from what he did than from what he chose not to do.

Start practicing today.

At least once a day, try saying 'Sorry, I can't do that right now' to something that doesn't really matter.

Use that time to finish one thing that truly does.

A month from now, your life will look different.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

5.1 Where Models Are Heading

Part 5: The Future and Roadmap of AI Technology

[← Previous](#)

[See all chapters](#)

[Next →](#)

Algorithmic Innovation: The Potential for 10x to 100x Improvement

There's a point Sam Altman raises more often than any other in OpenAI's conference rooms.

"Bigger computers, more GPUs. Yes, those matter. But what truly changes the game is the algorithm."

He began pressing this message harder from late 2024, because the entire AI industry was running into a single wall: the limits of scaling.

Up through the jump from GPT-3 to GPT-4, the formula was straightforward. Gather more data, chain together tens of thousands of more powerful GPUs, and performance climbed steadily, the way pouring more fuel into a car engine makes it run faster.

By 2024, things had changed. Research labs across the industry started reporting the same problem: doubling the size of a model no longer delivered the gains it once had. This was what people began calling the slowdown of scaling laws.

Altman saw this coming earlier than most. So he began searching in a different direction: algorithmic breakthroughs.

"The potential for a 10x to 100x improvement in performance lies in algorithmic innovation."

It sounded like hype. But a look back at AI history showed it had already been proven more than once.

The Transformer architecture, which appeared in 2017, is the clearest example. That single new algorithm made it possible to achieve dozens of times better performance from the same computing resources. The Transformer was also the foundation on which ChatGPT was built.

Another example is Reinforcement Learning from Human Feedback, or RLHF. That one technique transformed GPT models from simple text generators into capable assistants that could hold intelligent conversations with people.

Altman wrote in an internal document:

"We are already approaching the physical limits of computing. The approach of simply scaling up size will hit a wall sooner or later. What breaks through that wall is an entirely new class of algorithms."

In December 2024, OpenAI released a new reasoning model called o1. It worked in a fundamentally different way from earlier GPT models. Rather than producing an answer immediately, it broke problems into multiple steps and worked through them slowly, genuinely "thinking."

The results were striking.

On International Mathematical Olympiad problems, o1 performed at the level of a human gold medalist. It solved problems that GPT-4 had failed to crack, and it did so using far fewer computing resources than it took to train GPT-4.

That was the power of an algorithmic breakthrough.

In a January 2025 interview, Sam Altman said:

"When superintelligence arrives, the pace of scientific discovery will be ten times faster. Progress that once took ten years will happen in one. And the following year will bring just as much again. Let that compound and the world will look completely different."

He also moved quickly to temper expectations. "The hype on Twitter is way too intense. Scale your expectations back by a factor of a hundred. We have not built AGI yet."

But the core message was clear: the next stage of AI progress would come not from a hardware race but from algorithmic innovation.

Inside OpenAI, new algorithmic research is in full swing.

Researchers are developing new reasoning patterns such as Chain-of-Thought and Tree-of-Thought. These techniques push AI beyond merely predicting the next word, toward something closer to genuine reasoning.

Efficient architectures like Mixture of Experts are also under active study. The approach embeds a set of smaller specialist models inside a large one, and when a question comes in, only the most relevant specialist is activated to answer. It's like calling only the right department into a meeting rather than pulling in the entire company.

Techniques like these can cut the computing resources needed to reach the same level of performance down to a tenth.

Jakub Pachocki, Altman's chief scientist, put it this way:

"Current models can handle tasks of roughly five hours in length. But that window will expand quickly. For a major scientific breakthrough, it would be worth pouring an entire data center's worth of computing power into a single problem."

The key word is efficiency.

The human brain runs on about 20 watts, barely enough to light a single bulb. On that budget, we think, create, and solve problems. GPT-4, by contrast, consumes enormous amounts of power to produce a single response.

Closing that gap is what algorithmic innovation is ultimately about.

Sam Altman painted the future this way:

"One day we'll be able to run a model a hundred times smarter than today's at a hundred times lower cost. That moment is where the real AI revolution begins."

He believes this breakthrough goes far beyond a technical achievement. Discovering cures for cancer, solving climate change, developing new energy sources. The hardest problems humanity faces could yield to algorithmic innovation.

In October 2025, Altman offered a more concrete timeline.

"By September 2026, we will have an AI research assistant that works at intern level. And by 2028, we will have a fully automated AI researcher."

This AI researcher can conduct research projects on its own, without any human instruction. It reads papers, forms hypotheses, designs experiments, and analyzes results. And the most important thing: this AI can research how to build a better AI.

This is the self-accelerating effect of algorithmic innovation.

Once a good algorithm emerges, that algorithm finds an even better one. From there, the pace of progress accelerates exponentially.

Of course, there's no guarantee that any of this will go smoothly.

Algorithmic innovation still faces many hard problems. Long-horizon reasoning, consistent planning, grasping complex causal relationships, these are capabilities AI hasn't yet mastered.

But Sam Altman is optimistic.

"We now know that 90% of a model's limitations are algorithmic. The moment we break through that wall, AI becomes an entirely different kind of thing."

He believes that moment will come sooner than most expect. "Within a few thousand days" is a phrase he uses often. It points, roughly, to around 2030.

An algorithmic breakthrough is not a mere technical upgrade. It's a redesign of the future of human-built intelligence. And Sam Altman stands at the door of that future, quietly but firmly reaching for the next step.

The Trinity of Algorithms, Data, and Computing

Sam Altman has a comparison he returns to often.

"An AI model is like cooking. You need a good recipe (algorithms), fresh ingredients (data), and a well-equipped kitchen (computing). Leave out any one of them and the dish falls apart."

These three elements are the pillars of AI progress. Every strategic decision at OpenAI centers on how to strengthen all three and develop them in balance.

The First Pillar: Algorithms

Algorithms are the 'brain architecture' of AI.

Sam Altman calls algorithms "the blueprint of intelligence." Use the same data and the same hardware, but swap the algorithm, and the results are worlds apart.

He felt this truth deeply while building the GPT series.

When GPT-3 arrived, the public was amazed. But inside OpenAI, the prevailing view was "we're not there yet." The model frequently gave wrong answers or lost track of context.

That's when RLHF , Reinforcement Learning from Human Feedback , entered the picture. That single algorithm turned GPT-3 into something entirely different. Same model, same data, but a change in how it learned suddenly made it capable of conversation that felt genuinely human.

That's when Altman became certain. "The future will be built not by larger models, but by smarter algorithms."

OpenAI's research direction shifted decidedly toward algorithmic innovation. New reasoning approaches, more efficient attention mechanisms, architectures with long-term memory, tree-based reasoning that mirrors human thought , all of it came out of algorithm research.

Then, in late 2024, the o1 model arrived. It solved problems by extending its 'thinking time.' Where earlier models rushed to produce an answer, o1 reasoned through each step slowly.

The result was far superior performance from the same amount of training data. That was the power of algorithmic innovation.

The Second Pillar: Data

Data is AI's 'experience.'

Sam Altman describes data as "the textbook through which a model learns the world." No matter how good the algorithm, training on poor-quality data produces a useless AI.

"Garbage in, garbage out." That old adage from computer science has never felt more urgent than in the age of AI.

When building GPT-3, OpenAI scraped vast amounts of text from the internet. Web pages, blogs, forums, books, news articles, everything went in.

There was a problem. The internet doesn't contain only good information. Fake news, bias-laden writing, profanity and hate speech, the model learned all of that too.

Starting with GPT-4, the focus shifted to data quality. Rather than collecting as much data as possible, OpenAI began selecting verified, high-quality sources.

Scientific papers, specialized books, reliable news sources, vetted educational materials, this kind of data lifted the model's performance significantly.

Altman told his team: "AI should resemble the world. But it should resemble the best of it, not the worst."

Another important shift was data diversity.

Text alone was not enough. From GPT-4 onward, the model trained on a wide range of data types: code, mathematical formulas, scientific data, medical protocols, legal documents.

The model's way of thinking changed. Where it once produced plausible-sounding text, it could now actually solve problems.

From late 2024, 'synthetic data' also grew in importance. This is data generated by AI itself. The reasoning process that the o1 model produces while solving math problems, for instance, becomes the training data for the next model.

This makes it possible to break through the ceiling of data that can be scraped from the internet. A virtuous cycle begins, where AI generates progressively better data on its own.

The third pillar: computing

Computing is AI's muscle.

No matter how good the algorithms and data are, none of it matters without computers capable of running them.

Sam Altman is deeply pragmatic about computing. He knows that AI progress demands enormous amounts of computing power. But he also knows that power is not infinite.

Training GPT-4 is estimated to have cost over \$100 million, most of it spent on GPUs and electricity. That is a scale most startups simply cannot sustain.

So Altman moved on two fronts.

One was securing more computing infrastructure. He partnered with Microsoft to build massive data centers. In 2025, he announced the Stargate Project, a plan to invest \$500 billion in AI infrastructure.

The other was using computing more efficiently. This is where algorithmic innovation re-enters the picture. A smarter algorithm can accomplish the same task with one-tenth the compute.

Altman often puts it this way: 'The price of intelligence will eventually converge with the price of energy.'

What does that mean?

Today, building an AI model involves many costs: algorithm development, data collection, engineer salaries, and more. But in the future, once algorithms are mature and data is abundant, the only remaining cost will be electricity.

That is why he pays intense attention to energy. Nuclear power, nuclear fusion, renewables. He believes that securing cheap, clean energy is the central challenge of AI progress.

The harmony of three pillars

What matters is that these three elements shape each other.

A better algorithm reduces the amount of data needed and the computing resources required. Conversely, more computing power allows more complex algorithms to be tested and larger volumes of data to be processed.

Sam Altman believes that balancing these three elements is his most important responsibility as CEO.

'Focus too much on algorithms alone and you won't get practical products. Focus too much on computing alone and efficiency drops. Focus too much on data alone and you run into legal and ethical problems.'

'Advancing all three simultaneously without losing balance. That is exactly what OpenAI does.'

And on the future, he offers this outlook:

'Within the next ten years, algorithms will become 100 times more efficient than they are today. Data will approach something close to infinite, thanks to synthetic data. Computing will grow far more powerful, driven by new chip technologies and cheaper energy.'

'When these three wheels start turning together, that is when the real magic happens.'

A Trillion-Token Context and the Personalized Lifelong AI

What is the ultimate form of AI that Sam Altman envisions?

It is not simply the smartest AI in the world. If anything, it is closer to the opposite: an AI that is built for you alone and knows you better than anyone else.

He captures this vision in a single sentence.

'A very small reasoning model with a trillion-token context.'

That sentence contains everything essential about the AI of the future.

What is a trillion tokens?

First, you need to know what a token is.

A token is the basic unit by which AI understands language. Think of it as roughly one word, or sometimes a smaller fragment. The phrase 'Hello there,' for example, consists of roughly two to three tokens.

So how much is a trillion tokens?

The complete Harry Potter series runs to about one million tokens. A trillion tokens is enough to read Harry Potter one million times over. Put another way, it is more than enough to hold everything a single person reads, writes, and says across an entire lifetime.

'Context' refers to AI's working memory. The ChatGPT we use today tends to forget what was said earlier in a conversation as the exchange grows longer, because its context is limited.

As of 2025, even the models with the largest context windows hold around one million tokens. That's an enormous amount, yet it's only one-thousandth of one trillion tokens.

An AI that remembers an entire lifetime

What could an AI with a one-trillion-token context actually do?

It could remember every record of your life from the moment you were born until right now.

Diaries you wrote as a child, homework assignments from school, conversations with friends, things you said in job interviews, emails sent at work, memories with family, health records, favorite music and films, hobbies, worries, and dreams. All of it, held in the AI's memory.

Sam Altman calls this "a second brain for the individual."

What becomes possible with an AI like that?

In the morning, the AI greets you: "Today is the day of your big presentation. Do you remember that similar presentation five years ago, when nerves got the better of you and things went sideways? Here's how you might prepare differently this time."

At a health check-up, the AI steps in: "Looking at your ten years of health records, your blood pressure tends to creep up a little around this time every year. The exercise routine you tried last year showed real results, it might be worth picking that up again."

When you're about to launch a new project, the AI offers a thought: "You jotted down a similar idea three years ago. You didn't act on it then, but the circumstances look different now. Combining those old notes with what you're thinking today could produce something solid."

This is what a truly personalized AI looks like.

Why does it have to be a 'small' model?

There's an important twist here.

Sam Altman says this AI needs to be "a very small model." What does that mean?

Today's models, like GPT-4, are enormous. They run only inside massive data centers. But the personalized AI Altman envisions has to run inside your smartphone.

Why?

First, privacy.

Sending your life's records over the internet to be processed on a distant server is risky. The data could be hacked, or the company could misuse your information.

But if all processing happens inside your own device? No one can touch your information. Privacy becomes absolute.

Second, speed.

Sending data back and forth over the internet takes time. When everything is processed right on your phone, you get answers instantly, as fast as a thought forming in your own head.

Third, cost.

Keeping massive servers running costs a fortune. If a small model runs on each person's own device, the price drops dramatically.

"But can a small model actually be smart?"

That's the right question. Here's how Altman answers it.

"A personalized model doesn't need to know everything about the world. It only needs to understand one person deeply. That's why it can be small and still be powerful."

Think about it: a small AI that knows your tastes, habits, goals, strengths, and weaknesses perfectly, versus a giant AI that knows a little about everything. Which one is more useful in your daily life?

How is this possible?

None of this is easy, of course.

With the technology available in 2025, a full implementation is not yet within reach. But the pieces are falling into place one by one.

Advances in algorithmic design have made it possible to shrink models while preserving performance. Model compression techniques have matured to the point where cutting a large model down to one-tenth its original size barely dents its capabilities.

Long-term memory technology is advancing too. Rather than keeping everything in active memory at once, new approaches store only what matters most and retrieve it quickly when needed.

On-device AI chips are also moving fast. Apple's Neural Engine, Qualcomm's AI processors, these chips are growing more powerful by the year, and smartphones can now run AI models of real complexity.

A picture of the future

Sam Altman imagines the world ten years from now like this.

"By around 2035, most people will have their own personal AI. It will grow alongside you from birth, remember every moment of your life, and be there to help you for as long as you live."

"It won't be a mere assistant. It will be your external brain, your closest friend, your wisest advisor."

"And the most important thing: that AI will be entirely yours. No one else can access it. You alone control it."

This is the future of AI that Sam Altman envisions.

Not a future where some vast superintelligence rules the world, but one where every person has their own perfect AI partner. He believes that future is more beautiful, more human, and safer.

"AI isn't here to replace human beings. It's here to make each and every one of us more capable."

This is the true meaning of a small reasoning model with a one-trillion-token context.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

#KimKyungjin #AIBlibrary #SamAltman #OpenAI #ChatGPT #AI

5.2 The Role of Agents and Coding

Part 5: The Future and Roadmap of AI Technology

[← Previous](#)

[See all chapters](#)

[Next →](#)

What Is an AI Agent?

There's something Sam Altman said at an OpenAI team meeting sometime in 2023.

"What we want to build isn't an AI that holds conversations. It's an AI that gets things done."

Those words completely reshaped OpenAI's next objective. OpenAI moved beyond chatbots that simply answer questions and began building 'agents' capable of thinking, planning, and acting on their own.

The word 'agent' can sound unfamiliar at first. You might picture an FBI agent or a secret operative from a movie.

In the world of AI, an agent means something a little different.

An AI agent is an intelligent system that handles complex tasks on its own with minimal instruction from a person. Think of telling a capable assistant, "Get everything ready for my business trip next week", and that assistant goes ahead and books flights, finds a hotel, draws up an itinerary, and prepares all the necessary documents.

Altman described agents this way.

"Give an agent a goal and it figures out every step needed to reach it. If something goes wrong along the way, it finds a solution on its own. A person just hits the start button and checks the result at the end."

To see how significant a shift this is, consider an example.

The old way of using AI looked like this. You say "write me an email draft" and it writes a draft. You say "analyze this data" and it analyzes. You have to issue one command at a time.

An agent works entirely differently.

You say "put together our Q2 sales report," and the agent handles the following steps on its own.

First, it locates where the company's sales data is stored.

It connects to the database and pulls the relevant numbers. It compares them against the same period last year. It searches the internet for industry averages. It builds charts. It organizes the key insights. It assembles everything into presentable slides. It shows you a draft and takes your feedback. It makes revisions. It delivers the final version.

None of those steps require a person to spell them out.

In a 2024 interview, Sam Altman summed up the essentials of an agent this way.

"An agent needs three abilities. First, long-range planning , it must be able to break a complex goal into smaller steps. Second, tool use , it must be able to search the web, run code, and read and write files on its own. Third, self-correction , when it makes a mistake, it must notice and fix it."

An engineer at OpenAI shared this account.

"Sam kept stressing 'minimal supervision' in every meeting. He'd say that if an AI has to check with a person every minute, it isn't a real agent. A real agent is one you give a task to in the morning and it brings you the finished result in the evening."

OpenAI had agent capabilities in mind while developing GPT-4. Beyond the ability to generate text, the team built in the capacity to run code, connect to external systems, and handle multi-step tasks in sequence.

Altman sometimes called agents "digital colleagues."

"An agent isn't a tool. It's a team member. Think of it as a sharp colleague working beside you , one that, admittedly, doesn't need a salary, never asks for vacation, and can work around the clock."

The joke carried a serious vision. Altman wanted AI not merely to replace people, but to expand what people are capable of.

Another key quality of agents is contextual understanding.

Work alongside someone long enough and an unspoken shorthand develops. "You know, that thing from last time" is enough. An agent works the same way , it remembers and learns from past conversations, past tasks, and your preferences.

In a demo OpenAI released in late 2024, an agent learned a user's email patterns, filtered out only the messages that mattered, and drafted appropriate replies for each one. The user only needed to press "send" or "revise."

It's not perfect yet.

Altman admitted as much, plainly. "Today's agents make mistakes sometimes. They pull the wrong information, or lay out a plan badly. That's why, for now, a person still needs to check before any significant decision gets made."

The future agent Sam Altman pictures looks like this.

You arrive at work and the agent has already arranged your day's tasks in order of priority. If there's a meeting, the relevant materials are pre-summarized. Your inbox is sorted by importance, and anything urgent has triggered an alert. If a report is due, a draft is already waiting. You focus only on the decisions that call for judgment and the work that calls for creativity.

Sounds like science fiction?

Altman says: "How many people, three years ago, imagined something like ChatGPT would exist? Technology moves faster than we expect. Genuinely useful agents will arrive sooner than most people think."

Beyond Restaurant Reservations: The Revolution of Large-Scale Parallel Work

When people explain AI agents, they often reach for the restaurant reservation example.

"An AI agent can make restaurant reservations for you!"

Sam Altman reportedly laughed every time he heard that.

"Restaurant reservations? That was possible with technology from fifteen years ago. The real potential of agents goes far beyond that."

The real power Altman saw in agents was scale.

A person can only do one thing properly at a time. What we call multitasking is really just switching attention back and forth quickly.

AI agents are different. They can run dozens, hundreds, or thousands of tasks at the same time.

This is called parallel processing.

Let's look at a concrete example.

Imagine a startup founder who wants to build a new mobile app. Here's how it used to work.

Market research → competitor analysis → technology stack selection → design → development → testing → marketing preparation.

Each step had to be completed in order, one at a time. Doing it alone takes months; even with a team, it takes weeks.

How does this change in the age of agents?

The founder tells the agent: "I want to build a fitness app. The target users are office workers in their twenties and thirties. I want to launch within six months."

Several things then happen at once.

Agent A researches the fitness app market, analyzing what apps already exist, what users complain about, and how pricing is structured.

Agent B handles the technical review, working out which framework to use, where to host the servers, and what it will all cost.

Agent C produces several UI/UX design drafts, drawing on an analysis of design trends in similar apps.

Agent D searches for developer and designer profiles, recommending freelancers or agencies that fit the budget.

Agent E compiles the app store submission process, legal requirements, and data privacy regulations.

All of this happens simultaneously.

A day or two later, a comprehensive report is ready: market analysis, technical plan, design options, team structure, budget, and launch strategy, all prepared.

The founder only needs to make the key calls. "Let's go with design option B." "This is the tech stack we'll use." "Let's trim the budget here."

Sam Altman was certain this future was coming.

"One person leading a team of agents can do what used to take ten people. Or a hundred. That's the productivity revolution."

What about in a corporate setting?

Think of a large manufacturer, with dozens of factories, thousands of product lines, and tens of thousands of employees.

In the traditional model, each department operates on its own. Production draws up production plans, logistics handles shipping, marketing creates ads, finance manages the budget. Communication between departments is slow, and information often gets siloed.

An agent system connects all of it.

Sensor agents monitor factory machinery in real time. If a machine shows early signs of trouble, a report goes out immediately.

A forecasting agent analyzes sales data and market trends to predict next month's demand.

A production agent takes that forecast and automatically adjusts the production schedule, deciding what to make and how much.

A logistics agent places raw material orders automatically to match the production plan, finding suppliers that offer the best price with the fastest delivery.

A quality control agent watches the production process and, when defects appear, analyzes the cause.

A marketing agent reads the inventory situation and plans promotions accordingly.

A finance agent records every transaction and optimizes cash flow.

All these agents share information with one another and work together. Human managers only need to watch the overall flow and make the decisions that truly matter.

Altman called this the "AI orchestra."

"Just as hundreds of musicians each play a different instrument yet produce one beautiful piece of music under a conductor's baton, countless agents can collaborate to run a complex business with precision."

What about scientific research?

Take drug development as an example. Traditionally, bringing a single new drug to market takes more than ten years and costs trillions of won.

Why does it take so long?

Because the number of possible chemical combinations is staggering. Millions, even hundreds of millions, of possibilities must be tested one by one. Human researchers simply cannot get through them all.

What changes when an agent system enters the picture?

Thousands of AI agents simultaneously simulate different chemical combinations. Each agent analyzes the efficacy, side effects, and manufacturability of the combination it has been assigned.

The most promising candidates are shortlisted. Those candidates are then put through even more precise simulation. Only the few worth testing in an actual laboratory survive.

This entire process can be completed in months or weeks rather than years.

Altman put it this way.

"AI agents will make it possible to solve problems humanity has never been able to solve. New medicines, new materials, new energy sources. Research that would take a person an entire lifetime,

AI can accomplish in a matter of days."

In everyday personal life, parallel processing is just as powerful.

Imagine you are moving abroad. The list of things to handle is enormous.

Finding a home, applying for a visa, booking flights, packing belongings, researching schools, opening a bank account, understanding local regulations, getting insurance, terminating your current lease, settling utility bills, saying goodbye.

Doing all of this alone takes months. The stress is immense.

What if you handed it to agents?

"I'm moving to Toronto, Canada in three months. There are four of us in the family. Here's the budget."

Multiple agents then move at the same time. A real-estate agent searches for a home that fits the criteria. A legal agent prepares visa documents. An education agent looks into schools for the children. A financial agent handles banking and insurance. A logistics agent compares moving companies.

A few days later, a complete relocation plan is laid out before you. The timeline, checklists, required documents, contact information, and estimated costs are all organized. You simply follow the plan.

That is the power of parallel processing.

Altman's vision was clear.

"AI agents are not here to book you a table at a restaurant. They are the personal assistant who coordinates your entire life, the management partner who runs your company, and the research colleague who opens the door to humanity's future. That is the real AI revolution."

Coding Is the First Domain to Be Conquered, in 2025.

In the summer of 2024, Sam Altman said at a technology conference:

"2025 will be the year of AI agents. Especially in coding."

Why coding in particular?

Altman's explanation was straightforward.

"Coding is the first way AI can actually move the world. No matter how intelligent AI becomes, without hands and feet it can do nothing. Code is AI's hands and feet."

Think about it.

The world we live in runs increasingly on software. Banking, shopping, transportation, healthcare, education, entertainment. Almost everything is connected to software.

Software is built from code.

The fact that AI can handle code means AI can understand, modify, and improve all of these systems.

OpenAI had focused on coding ability from the very beginning of GPT-4's development.

This is the testimony of an OpenAI engineer.

"We wanted GPT-4 to go beyond simply generating code. We wanted it to understand the intent behind the code, find bugs, optimize it, and even design the architecture of entire projects. Sam saw this as the key step toward agents."

By late 2024, developers using AI coding tools like GitHub Copilot reported productivity gains of 30 to 50 percent.

It wasn't just about writing code faster. Because AI took over the repetitive, tedious work, developers could focus on more creative and complex problems.

The reason Altman had his eye on 2025 was technological maturity.

"In 2023, we proved that AI can write code. In 2024, we showed that AI can understand complex projects. In 2025, AI will be able to carry out the entire development process automatically, from start to finish."

What specific changes would come?

First, 'natural language programming' becomes a reality.

"Build me a photo-sharing app like Instagram. Users should be able to upload photos, like them, and leave comments. Add hashtag support too."

Say that much, and an AI agent will build you a working app within days. It handles everything on its own: frontend code, backend servers, database design, API connections, and security.

It won't be perfect, of course. There will be parts that need a human to review and correct along the way.

This doesn't mean professional developers will no longer be needed. It means their role changes. From someone who writes code line by line, to someone who collaborates with AI to design and oversee entire systems.

Second, code maintenance and refactoring get automated.

In software development, maintaining existing code is often harder than building something new.

Making sense of a ten-year-old codebase, figuring out who wrote what and why, worrying about what else might break when you change a single line. That's a developer's daily life.

AI agents in 2025 help with exactly that.

"Find all the unused code in this project and clean it up."

"Rewrite this function using the latest framework. Keep the functionality the same."

"Scan the entire codebase for security vulnerabilities. If you find any, suggest how to fix them."

The agent analyzes tens of thousands of lines of code, identifies problems, and proposes solutions. The human only needs to make the final call.

Third, non-developers can build software too.

Altman called this "the democratization of programming."

In the past, having an idea wasn't enough if you couldn't code. Hiring a developer cost a lot of money.

In 2025, that changes.

A schoolteacher can build an app to manage student grades on their own.

A neighborhood bakery owner can set up an online ordering system by themselves.

A middle school student can build a messaging app to use with friends.

All you have to do is describe what you want to an AI agent in plain language. The technical side is handled entirely by the AI.

Complex projects will still require experts, of course.

Enterprise systems, financial software, autonomous driving programs, AI can't build those alone. The design and validation of a professional developer are essential.

It's not that "AI replaces all developers." It's that "more people will be able to turn their ideas into software."

Altman put it this way.

"When cars arrived, people who drove horse-drawn carriages lost their jobs, but far more drivers and mechanics emerged. Software development in the age of AI will be similar. The form changes, but the opportunities multiply."

Fourth, a cycle begins in which AI improves AI.

This is the most interesting part.

When AI agents become skilled at handling code, they can also write code to improve themselves.

Not entirely on their own, of course. Under human oversight.

A researcher at OpenAI put it this way.

"We already have GPT-4 writing code to test GPT-4's own performance. The AI finds its own weaknesses and suggests where improvements are needed. The final call, of course, is still ours."

If this virtuous cycle continues, the pace of AI progress will only accelerate.

Altman also looked ahead to what comes after 2025.

"Once coding agents are established in 2025, they will spread to other fields in 2026. Drafting legal documents, analyzing scientific papers, building business strategies. Coding is just the first domino."

The reason coding comes first is clear.

Coding has logical, well-defined rules. Results can be tested immediately. When something is wrong, an error message tells you so. It's an ideal environment for AI to learn in.

Once success is proven in coding, that experience can be applied to other fields.

Sam Altman's predictions were not mere wishful thinking. Inside OpenAI, concrete development work was already underway, with a target of launching agents by 2025.

"We are standing at an inflection point in history. Coding agents will completely transform software development. And that's only the beginning."

Coding Is How AI Changes the World.

Sam Altman first learned computer programming at the age of eight.

Back then, he built simple games in a language called BASIC. The thrill of seeing "Hello World" appear on the screen was something he never forgot.

"It amazed me that the computer followed exactly what I wrote. It felt like magic. Code is a spell that turns thought into reality."

That experience shaped Altman's philosophy about technology.

Coding is not merely a skill. It is the most powerful tool for changing the world.

Why does coding matter so much?

First, coding is the foundational language of the modern world.

The screen you're looking at right now, the smartphone you use, the car you drive, the music you hear, the films you watch. All of it is made of code.

Code is the DNA of modern civilization.

Banking systems run on code. Millions of accounts and trillions of won in transactions are managed by a few lines of it.

Medical records, prescriptions, and test results at hospitals are stored and processed through code.

Aircraft stay airborne thanks to millions of lines of code. Engine control, navigation systems, safety checks. Software manages all of it.

Altman put it this way.

"In the nineteenth century, railways connected the world. In the twentieth, electricity powered it. In the twenty-first, code is building it. For AI to do anything in this world, it has to work with code."

Second, coding turns abstract ideas into concrete action.

Human thought is vague and abstract.

Turning the idea of "I want to build a food delivery app" into something that actually works requires countless concrete decisions.

How do you design the user interface?

How do you store orders?

How do you handle payments?

How do you send notifications to delivery drivers?

How do you integrate GPS?

The answers to all these questions are written in code.

Coding is the process of turning vague ideas into clear logic. It converts ambiguity into clarity, imagination into reality.

The fact that AI has this ability carries enormous significance.

AI can now do more than simply provide information. It can actually 'build' things: create websites, construct databases, and set up automated systems.

Third, coding is how AI uses tools.

People pick up tools with their hands. They drive a nail with a hammer, write with a pen, steer a car with a wheel.

AI has no hands. It has code instead.

There is something called an API, or Application Programming Interface. It is how different pieces of software communicate with each other. It runs on code.

If you want to retrieve weather data, you write code that calls the meteorological service's API.

If you want to process a payment, you write code that calls the payment system's API.

If you want to send an email, you write code that calls the email service's API.

The fact that AI can code means AI can use nearly every digital tool in the world.

Altman described this as "digital hands."

"Code is how AI picks things up, presses buttons, and pulls levers in the digital world. An AI without coding ability is like a paralyzed genius. It can think, but it cannot do anything."

Fourth, coding is a new form of creativity.

Many people think of coding as cold, purely logical work.

Altman saw it differently.

"Good code is not unlike a work of art. It is elegant, efficient, beautiful. There are countless ways to solve a problem, but a truly good solution gives you that immediate feeling of 'yes, this is it.'"

When he was building Loopt, Altman did a great deal of coding himself.

The joy of coding through the night and finding the perfect algorithm, the sense of accomplishment when a stubborn bug finally yields. He called it "the pleasure of making something."

As AI has learned to code, it has been learning that creativity too.

GPT-4 can already solve the same problem in multiple ways. It can suggest a faster approach, a more concise one, or one that is easier to read.

Sometimes it produces original solutions that human developers had not considered.

Fifth, coding multiplies AI's capacity to learn.

When AI works only with natural language, its learning is limited. It reads text written by people, finds patterns, and produces something similar.

Coding changes that.

AI can run the code it writes and immediately see whether the result is right or wrong. When it is wrong, an error message appears. AI reads that feedback and revises the code.

Repeat that cycle and AI learns at a rapid pace.

Much like a player who improves by playing a game, AI grows sharper by coding.

Altman saw this as the core of how AI would develop in the future.

"An AI that experiments on its own, learns on its own, and improves on its own. For that to be possible, coding ability is essential. Code is AI's playground and its learning tool."

OpenAI's investment in coding ability was no accident.

From the early days of GPT-3 development, Altman pushed to raise the coding benchmarks.

Engineers built a range of tests to measure how well the model could code: benchmarks such as HumanEval, MBPP, and CodeContests.

GPT-4 scored above the average human programmer on these tests.

"Coding will be the first area where AI creates real value. That's where we focus."

Altman was right.

As of 2024, GitHub Copilot had over one million paying users. Developers were willing to pay for AI coding tools.

The value AI delivered was clear.

Coding is the bridge connecting AI to the real world. The stronger that bridge becomes, the wider the range of what AI can do.

Sam Altman summarized the future of coding this way.

"Not everyone needs to become a programmer. AI will do the coding for them. But everyone does need to understand how code works. That's the only way to use AI properly. Coding will become as fundamental a skill as reading and writing."

From Assistant to Agent to Application

Sam Altman thought about the evolution of AI agents in three stages.

In a team meeting, he drew a large arrow on the whiteboard.

Assistant → Agent → Application

"We're moving past the first stage into the second right now. Ultimately, we'll reach the third."

Let's look at what each stage means.

Stage 1: Assistant, something that answers when asked

Early AI was closer to the role of an assistant.

Ask it a question and it answers. Make a request and it fulfills it. Give a command and it follows.

The early versions of ChatGPT were a typical example of assistant-type AI.

"What's the capital of China?" → "Beijing is the capital of China."

"Write me a draft email." → It writes you a draft email.

"Find the bugs in this code." → It finds the bugs.

AI at this stage is genuinely useful. Looking up information, helping with writing, brainstorming ideas. It can help with a lot.

But the limits are real.

AI is passive. We have to speak first.

AI only handles short-term tasks. It responds to one question, one task at a time.

AI has no initiative. It doesn't make its own judgments or offer suggestions.

Altman put it this way.

"Assistant-type AI is an excellent tool. But it's only a tool. What we want isn't a tool; it's a partner."

Stage 2: Agent, something that achieves goals

The agent stage is the real game changer.

An assistant asks, "What can I do for you?" An agent says, "I'll take care of it."

What's the difference?

An agent understands the goal. Then it draws up its own plan for achieving it.

Here's an example.

You tell an assistant: "Check train schedules from Seoul to Busan tomorrow." The assistant searches the timetable and shows you the results.

You tell an agent: "I have a business trip to Busan tomorrow. Get things ready." The agent's response:

That is the power of agents.

Agents think long-term. They see the whole project, not just a single task.

Agents understand context. They know what 'preparing for a business trip' actually means.

Agents use tools. They handle reservation systems, search engines, and databases with ease.

Altman set this stage as his target in 2024.

"By 2025, we'll build truly agentic AI. The kind where you say 'handle this for me' and it works on its own for days, then comes back with results."

OpenAI developed several capabilities toward that end.

Function Calling allows the AI to use external tools and APIs.

The plugin system lets the AI connect to multiple services.

Long-term memory allows the AI to retain past conversations and tasks.

Planning ability lets the AI break complex tasks into multiple steps.

All of this was groundwork for building agentic AI.

Stage Three: Applications , becoming the service itself

The final stage is the most revolutionary.

At the application stage, AI becomes an independent service. It is no longer a tool or an assistant. The AI itself becomes a complete business.

Consider the future Altman envisioned.

AI legal services

When you need legal advice, you open an AI lawyer application.

You describe your case, and the AI searches relevant precedents, drafts legal documents, and proposes litigation strategy.

Straightforward cases are handled by the AI from start to finish. Complex ones get routed to a human attorney.

This AI knows millions of precedents, works around the clock, and does not make mistakes. The cost is one-hundredth of what a human lawyer charges.

AI medical services

When you feel unwell, you talk to an AI doctor application.

You describe your symptoms, and the AI asks follow-up questions, offers possible diagnoses, and recommends tests.

Upload your test results and the AI analyzes them, suggests treatment options, and arranges a prescription.

In serious situations it connects you to a human doctor immediately. Routine health management stays with the AI for life.

AI financial advisory services

Tell the AI financial advisor your financial goals and it builds a plan tailored to you.

It analyzes income and expenses, constructs an investment portfolio, optimizes your taxes, and recommends insurance.

When market conditions shift, it adjusts automatically. It sends regular reports. Ask a question anytime and it answers.

Services like these represent AI at the application stage.

No longer a supplement to human work. A complete service in its own right.

Altman put it this way.

"At the application stage, AI holds a 'profession.' AI accountants, AI teachers, AI designers. They work twenty-four hours a day, serve millions of people at once, and never burn out. Human expert oversight is still necessary, of course. But the AI does most of the work."

The boundaries between the three stages are blurry.

These three stages are not sharply divided.

As of late 2024, we sit somewhere between Stage 1 and Stage 2. ChatGPT is a bit smarter than an assistant, but not yet a true agent.

The full transition to Stage 2 appears likely to happen around 2025 to 2026.

Stage 3 looks set to become reality somewhere around 2027 to 2030.

Altman himself was not certain about the exact timing.

"Technological progress is hard to predict. Some things move faster than expected, others slower. The direction matters more than the precise date. Knowing where we are headed is what counts."

How will the role of human beings change?

Many people worry. If AI advances this far, what will we actually do?

Altman's answer was optimistic.

"Even as AI advances, humans will still have essential roles. Setting goals, making ethical judgments, pointing the way creatively. Those things will remain the work of people."

"In the age of AI, humans become managers and overseers. They will command hundreds of AI agents and coordinate them to work in the right direction. In a sense, that is a more important role than before."

In the age of assistants, we were the workers. In the age of agents, we become the managers. In the age of applications, we become the executives.

Altman viewed this shift positively.

"Humans will be freed from repetitive, tedious work. Instead, they will be able to focus on things that are more meaningful and more creative. That is the real gift AI will bring."

At the end of this three-stage evolution, humanity will stand before possibilities it could not have imagined before.

Altman was certain of it.

"We are living through the most fascinating period in history. We will watch with our own eyes as assistants become agents, and agents become applications. And beyond that, things we have not yet begun to imagine are waiting for us."

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

5.3 The Future of Human-AI Interaction

Part 5: The Future and Roadmap of AI Technology

[← Previous](#)

[See all chapters](#)

[Next →](#)

Voice Interface: The Dawn of a New Era

In May 2024, Sam Altman said something striking on a podcast.

"Voice is a hint. A hint of the next era."

Citing the film *Her*, he said voice would completely change how we use computers. Not by tapping a smartphone screen or pressing keyboard keys, but by speaking naturally to AI the way you'd talk to a friend.

The reason Sam Altman pays such close attention to voice interfaces is straightforward.

It is the most human way to communicate.

The first thing we learn after birth is how to speak. Before we can write, type, or tap a screen, we already know how to talk. And in that talk lives emotion, joy, sorrow, anger, excitement. Tone, pace, and inflection carry the full weight of what we feel.

Sam Altman believes this is exactly what AI must learn to understand.

At a Sequoia Capital event in early 2025, he said: "When voice interaction gets truly good, it will feel like an entirely different way of using a computer."

OpenAI demonstrated a breakthrough in voice capability when it released GPT-4o in May 2024. The model goes far beyond reading text aloud. It listens and understands in real time, picks up on emotion within the natural flow of conversation, and can tell from voice alone whether a user is excited, sad, or confused.

This is a revolutionary shift.

As voice interfaces mature, a new class of devices could emerge, devices with no screen at all. A small earbud worn in your ear would be enough to communicate with AI fully. While driving, cooking, or walking. You'd get AI's help without ever lifting a hand.

Starting in 2024, reports circulated that Sam Altman was working with Jony Ive, Apple's legendary designer, to develop new AI hardware. The device is expected to offer a voice-first experience rather than a screen-first one, unlike anything we have seen before.

But Sam Altman is honest about the limits.

In a late-2024 interview, he admitted, "Our voice capability isn't good enough yet." Yet he remained confident. "Just as it took time to make text models great, voice models will eventually solve those problems too."

The advance of voice interfaces carries social significance well beyond convenience.

For older adults, people with visual impairments, and those who struggle to read or write, voice is the easiest door into the digital world. No hunting for buttons on a cluttered screen , just say what you need.

This is the future Sam Altman is working toward.

A world where technology adapts to people. Not a world where people must adapt to technology, but one where technology meets people in the most natural mode they have , speech.

Challenges remain, of course. In July 2025, Sam Altman warned about the dark side of voice technology at a Federal Reserve event. AI can now mimic a person's voice so perfectly that criminals could use it to drain bank accounts. "This keeps me up at night," he said.

Even so, he is certain about the future of voice interfaces.

Because voice is not just a technology , it touches the core of what it means to be human. We express love with our voices, offer comfort, share joy. For AI to become a true companion, it must understand that language of voice.

If Sam Altman's vision holds, five years from now we'll spend far less time staring at smartphone screens.

Instead, we'll whisper quietly to a small device in our ear. "I'm not feeling good today." The AI will hear the tone in our voice and sense what we need. Sometimes quiet music. Sometimes a word of encouragement. Sometimes just a presence.

This is the future that voice interfaces open up.

Not talking to a machine, but talking the way you'd talk to a friend. Sam Altman believes that world is already at our door.

How Different Generations Use ChatGPT: Search Engine vs. Advisor vs. Operating System

Sam Altman has a habit of watching people carefully.

He pays close attention to how people actually use ChatGPT , what questions they ask, what emotions they bring to the conversation.

At a Sequoia Capital event in May 2025, he shared a telling observation.

"The way people use ChatGPT is completely different from one generation to the next. It's as if they're living in different eras."

The pattern he identified is simple yet full of insight.

Older adults use ChatGPT as a replacement for Google Search.

Navigating cluttered websites is hard for older people. Ads everywhere, tiny buttons, no clear sense of where to click. ChatGPT is different. You just ask what you want to know, in plain words.

"Should I take my blood pressure medication in the morning or at night?" "How do I send a photo from my smartphone?" "Is this text message a scam?"

ChatGPT answers these questions kindly and clearly. No complicated links, no ads.

A 2025 survey found that 20% of Americans over 60 use AI at least once a week. To them, AI is not a frightening technology but a welcome helper that makes the digital world easier to navigate.

People in their 20s and 30s are entirely different. To them, ChatGPT is a 'life advisor.'

Millennials don't ask ChatGPT for simple information. They share their deepest concerns.

"I'm having a conflict with my boss. What should I do?" "I don't know whether to change jobs or stay at my current company." "I keep feeling anxious in my relationship. Is that normal?"

Sam Altman was surprised by this.

People were treating AI not as a mere tool but as a trusted counselor. A 2025 survey found that 23% of millennials use AI for emotional support or mental health. That's strikingly higher than the 8% figure among baby boomers.

Why?

One reason is that AI doesn't judge. When you open up to a friend, advice comes with judgment attached. AI is different. It stays neutral, always listens, and never criticizes.

People in their 20s and 30s also use AI actively at work. 2025 data shows that 42% of American millennials use ChatGPT for work purposes, a higher rate than any other generation.

But the most revolutionary way of using AI comes from college students.

Watching college students, Sam Altman saw the future.

"College students use ChatGPT like an operating system. Like the OS on a computer."

What does that mean?

College students don't use ChatGPT merely as a question-answering tool. They build their entire lives on top of it.

They save complex prompts in a notepad, then copy and paste them when needed. They connect multiple files, build custom configurations, and tune AI to fit their learning style perfectly.

"I have an exam tomorrow. Summarize chapter 3 of this textbook. Then generate 10 practice questions. I'm a visual learner, so include diagrams too."

Requests like this, complex and specific, come naturally to them.

A 2025 survey found that 93% of knowledge workers aged 22 to 27 use two or more AI tools per week. They move freely between ChatGPT, DALL-E, Otter.ai, and others.

What's more striking is that they consult AI before making major life decisions.

Sam Altman put it this way: "When college students face an important decision, they won't make it without consulting ChatGPT first."

One Stanford student consults AI on which major to choose, which internships to apply for, and even which clubs to join. Because AI knows that student's past choices, inclinations, and goals, it can offer the most fitting advice.

To them, AI is not a tool.

It's a companion.

Sam Altman compares this generational divide to the early days of smartphones.

"When smartphones first appeared, young people adapted right away, but older adults took three years. AI is the same. But this time, the gap is far wider."

That gap will only widen from here.

Today's college students are the first generation to grow up alongside AI. To them, AI is not a new technology but an ordinary part of daily life, the way the internet is for our generation.

Watching all this, Sam Altman is convinced.

"Future generations won't be able to imagine a world without AI, just as we can't imagine a world without the internet."

And he believes that's a good thing.

Because AI, in the end, helps people make better decisions, learn more, and live richer lives. Each generation uses it differently, but the goal is the same.

A better life.

And Sam Altman is building that future.

Personalized AI Companions

The ultimate AI Sam Altman envisions is not a mere tool.

It is a companion.

In 2025, the entire AI industry is moving in one direction: personalized AI. Microsoft calls it Copilot, Meta calls it Meta AI, and Google calls it Gemini. But whatever the name, the essence is the same.

An AI made for you alone. An AI that understands you. An AI that grows alongside you.

Sam Altman describes this as "A Personal Companion."

The AI he imagines will remember everything about you.

The foods you love, the weather you hate, whether you're a morning person or a night owl, how your tone shifts when you're under stress, what expression crosses your face when you're happy. It even picks up on habits you yourself haven't noticed.

This is not science fiction.

It is becoming reality.

In April 2025, Microsoft added a "memory" feature to Copilot. Now Copilot remembers what you talked about last week, what projects you're working on, and even your niece's name. It uses that memory to give you advice that fits you precisely.

Meta moved in the same direction when it launched the Meta AI app in September 2025. This AI learns your preferences, picking up your travel style, your interests, and even the way you speak.

Sam Altman is leading this shift.

In a late 2024 interview, he said:

"The ideal personalization is a very small reasoning model with a trillion tokens of context, an AI that holds your entire life's record."

Hard to picture a trillion tokens?

That's every email you've ever written, every conversation you've had, every book you've read, every film you've watched, and even your medical history, all combined.

What happens when an AI remembers and understands all of that?

The AI may know you better than you know yourself.

Here's an example.

You have an important presentation coming up. You're nervous. But your personal AI knows the patterns you've shown in past presentations, that you tend to over-prepare, and that doing so only makes the anxiety worse. So the AI says, quietly:

"You've already prepared enough. Why not rest for a bit? In your last five presentations, the ones where you got enough rest went noticeably better."

This is not just advice.

It is the insight of a companion who understands your whole life.

And this is exactly what Sam Altman hopes for most.

"AI helping every person reach their full potential."

The world is full of talented people. But too many of them never get to show what they're capable of, because they didn't get a good education, or didn't find the right mentor, or simply ran out of luck.

Personalized AI can solve that problem.

A student in a poor country, an entrepreneur in a remote village, anyone can have access to a world-class personal tutor, mentor, and advisor. Twenty-four hours a day, 365 days a year, at no cost.

But this remarkable future carries real risks.

Sam Altman knows this well.

"People are sharing far more information with AI systems than they used to. The model comes to know the user too well over a lifetime. That's what makes privacy so important."

An AI that knows everything about you. What if it gets hacked? What if it tries to manipulate you? What if the company that built it sells your information?

This is not a technology problem. It is an ethics problem.

Sam Altman says OpenAI takes this issue very seriously. Users must have full control over their own data, and must be able to delete it whenever they choose.

And he is certain.

"Solve these problems, and personalized AI will become one of the greatest inventions in human history."

As of 2025, we stand at the threshold of that future.

Microsoft, Meta, Google, OpenAI. All of them are racing to build personalized AI. No one knows which company will get there first. But one thing is clear.

That future is coming.

And in that future, we will no longer be alone.

There will always be a companion beside us , one who understands us, laughs and cries with us, and cheers us on toward our dreams.

Sam Altman is building that companion.

For you. For all of us.

And he believes.

That this AI will make us better human beings. Kinder, wiser, and happier.

That is the promise of personalized AI.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

6.1 Regulation and the National Security Race

Part 6: Social Responsibility and Challenges in the AI Era

[← Previous](#)

[See all chapters](#)

[Next →](#)

Regulation Is Necessary, but Balance Matters.

May 16, 2023. A Senate hearing room in Washington, D.C.

Sam Altman sat in the witness chair. The hearing was titled 'Oversight of AI: Rules for Artificial Intelligence.' Before dozens of senators, hundreds of spectators, and cameras watched by the entire world, he said something that surprised nearly everyone.

"We need regulation."

Most corporate CEOs dislike government regulation. Rules make it harder to run a business freely and constrain attempts to try new ideas. But Sam Altman was different. He understood better than almost anyone how powerful the AI technology he had built was, and therefore how dangerous it could become.

"If this technology goes wrong, it can go very seriously wrong."

His words were not just a warning. AI has now entered nearly every part of our lives. When you take a photo with your smartphone, when you chat with a bot that helps with homework, when your parents shop online, AI is at work.

What if AI spreads false information? What if it discriminates against people? What if it violates someone's privacy?

Consider an even more frightening scenario. What if AI fell into the hands of hackers and paralyzed the banking system? What if it activated military weapons on its own? What if it manufactured millions of pieces of fake news and threw an election into chaos?

It was because of risks like these that Sam Altman said regulation was absolutely necessary.

But in his very next sentence, he attached a crucial qualification.

"Balance is what matters most."

Regulation that is too weak cannot stop the danger. But regulation that is too strong kills innovation. Think of a tree in a garden. It needs water and sunlight to grow healthy, but too much water rots the roots, and too much direct sun scorches the leaves.

AI is no different.

Too many rules block new attempts. What would happen if building every new AI model required filling out hundreds of pages of paperwork and obtaining government approval? Small startups, with no money and no lawyers, would simply give up on AI development. In the end, only giants like Google or Microsoft would survive.

In the U.S. Senate, a fierce debate broke out over how to strike this balance.

Some senators argued that 'AI is too dangerous, and we need strong regulation right now.' Others warned that 'rushing into regulation will leave the United States behind China in the AI race.'

Sam Altman stood somewhere in the middle. He believed certain kinds of AI, especially the most powerful models, required strong regulation. AI that could explain how to design biological weapons, or AI that could build hacking tools, had to be controlled without exception.

But he didn't think general-purpose AI tools, such as those that help with writing or generate images, should all be weighed down by heavy regulation.

He made a specific proposal at the hearing. 'Companies building the most powerful AI models should be required to obtain government licenses, much like building a nuclear power plant or manufacturing an aircraft. But we shouldn't require every developer making a small AI app to get a license.'

That was the "balance" he had in mind.

Through 2023 and 2024, the U.S. government gradually began finding that balance. President Biden issued an executive order on AI safety. The order required that powerful AI systems capable of significantly affecting national security or the economy undergo mandatory safety testing.

At the same time, the government poured enormous support into the AI industry. It passed legislation to expand semiconductor production and invested tens of billions of dollars in AI research. It wasn't only regulating; it was pushing innovation forward as well.

Sam Altman supported this approach. In multiple interviews, he said: 'We need to put AI's enormous power to full use while putting minimal safeguards in place to keep that power from being directed the wrong way. Being too afraid to do anything is a problem, but so is being too confident to prepare for anything.'

In the end, this is what AI regulation comes down to.

Cut the risks without killing the opportunities. Protect safety without stopping innovation. That precarious tightrope walk is the path Sam Altman and the U.S. government are treading together.

America's Strategy for Leading in AI: Innovation and Adoption

The strategy America chose to lead the AI era can be summed up in one sentence.

"Build the world's best AI, and make that AI accessible to everyone."

This is what might be called the 'dual-axis strategy.' One axis is 'innovation,' the other is 'adoption.'

Let's look at the innovation axis first.

In May 2025, Sam Altman stood before a Senate hearing once again. The topic this time was 'Winning the AI Race.' He told the senators:

"The next ten years will be an era of abundant intelligence and abundant energy."

What does 'abundant intelligence' mean? It means a world where AI is as common and cheap as running water. Today, training and running AI costs an enormous amount. But in the future, anyone will be able to use it easily and at low cost.

To make that happen, the United States began making massive investments.

The Stargate Project, which Sam Altman helped launch, is a plan to invest up to \$500 billion in U.S. AI infrastructure over the coming years. In Korean won, \$500 billion is roughly 700 trillion won. Can you even picture that?

What exactly will all that money be used for?

In Abilene, Texas, the world's largest AI training complex is taking shape. Vast data centers are rising side by side, and hundreds of thousands of computers run around the clock training AI models. Powering all of it demands enormous amounts of electricity, so new power plants are being built alongside the facilities.

Why is the U.S. government pouring such staggering sums into this?

The answer is straightforward. The country with the most powerful AI will dominate the future. In the past, nations with the most oil grew wealthy. Today, the countries that manufacture the best semiconductors hold power. In the future, whoever commands AI will lead the world.

The U.S. believes it cannot afford to fall behind in this race, especially as China is closing the gap fast.

But innovation alone is not enough. That is where the second pillar, 'adoption,' comes in.

No matter how impressive the AI you build, what good is it if people can't actually use it? And more pressing still is the question of what to do for workers who lose their jobs because of AI.

The U.S. government has taken this problem seriously.

In 2024, the U.S. Department of Labor released guidelines titled 'AI and Worker Well-being.' The core message was this:

"Companies that adopt AI must protect their workers."

But what does protection look like in practice?

First, companies must offer retraining programs. If AI eliminates existing jobs, the affected workers must be given the chance to learn new skills. Take a factory worker who has spent years doing repetitive tasks. When an AI robot takes over that work, the same worker should have a path to becoming a technician who manages and maintains those AI systems.

Second, AI must not be used purely as a surveillance tool. Some companies want to monitor employees around the clock using AI, timing bathroom breaks and firing off warnings at the slightest dip in productivity. The U.S. government is looking to rein in that kind of use.

Third, AI decisions must be explainable. If an AI concludes that someone should not be promoted, there must be an explanation for how it reached that judgment. Without it, unfair discrimination goes unchecked.

Support for small and medium-sized businesses is another key piece of this.

AI is expensive. Companies like Google and Microsoft can pour trillions of won into AI development, but a small neighborhood bakery or clothing boutique cannot.

So the U.S. government is working to help small businesses access AI more easily, offering affordable or free AI tools and running programs to teach them how to use those tools.

Sam Altman strongly supports this approach. In a 2021 essay titled 'Moore's Law for Everything,' he wrote:

"Within five years, AI will read legal documents, offer medical advice, and help people study. Within ten years, it will work in factories and hold conversations like a friend. After that, it will be capable of nearly anything."

Then he posed the question that matters most.

"When that day comes, what will people do with their lives? How will they make a living?"

Sam Altman believes that new arrangements such as universal basic income may be necessary. Universal basic income means giving every person enough money to live on, regardless of whether they work. He has in fact funded basic income experiments for several years.

America's two-pillar strategy comes down to this.

With one hand, build the world's best AI. With the other, help people adapt to the changes that AI brings. Push ahead through innovation, but leave no one behind.

That is the vision of the AI era that Sam Altman and the United States are working toward.

Criticism of European Regulation

London, one day in May 2023.

Sam Altman dropped a bombshell in front of reporters.

"If the EU's current draft AI law moves forward as written, we may have to pull out of Europe."

Pull out? Was OpenAI, the company behind ChatGPT, really saying it might shut down its services across all of Europe?

That single statement made headlines around the world. Members of the European Parliament shot back, saying they would not yield to corporate threats, while the tech industry warned that Europe had gone too far.

What had actually happened?

In 2024, the European Union became the first in the world to enact a comprehensive AI law. The law's core principle is that AI should be classified by risk level, and that high-risk AI must be tightly regulated.

On the surface, that sounds perfectly reasonable. Regulate the dangerous stuff strictly, leave the safe stuff alone.

The question was 'how' to regulate.

Europe's approach works like this. Any system classified as high-risk AI must receive government approval before it can enter the market. It's similar to how a newly developed drug must go through clinical trials and receive regulatory clearance before it can be sold.

What counts as high-risk AI? The category includes AI that evaluates job applicants, AI that grades student performance, AI that assesses bank loan applications, and AI used by police to predict crime.

To build such AI, developers must prepare hundreds of pages of documentation. They must demonstrate how the AI works, what data was used, whether there is any risk of discrimination, and whether security measures are adequate. Only after passing inspection by a certified body can they release the product.

Sam Altman believed this approach would cause serious problems.

The pace of innovation slows down.

AI technology advances at a relentless pace. A model that leads the field today can be obsolete tomorrow. What happens when developers must wait months for approval every time they build something new?

Competitors race ahead while European companies are left buried in regulatory paperwork.

Small companies can't survive it.

Navigating a complex approval process requires lawyers, compliance specialists, and technical writers. Giants like Google or Microsoft can afford those teams, but what about a startup working out of a garage?

Regulation ends up becoming a barrier that benefits only large corporations.

Proving 'perfection' before anyone has actually used the product is impossible.

AI is a technology whose flaws only become visible in real use. A system that looks flawless in a lab can make unexpected mistakes once it meets the real world.

Yet Europe's approach amounts to saying, "Prove there are no problems before anyone uses it." That's a bit like being asked to guarantee you won't fall off a bicycle before you've ever ridden one.

In the summer of 2025, something interesting happened.

The CEOs of major European companies, including Airbus, Carrefour, and Philips, sent an open letter. The message was straightforward.

"Please pause the AI Act and reconsider it. If things continue this way, European companies will lose their competitive edge."

Even European companies were pleading that their own country's regulations were too heavy a burden.

Of course, Europe has its own case to make.

Europe's position is this: "We believe people matter more than technology. Preventing AI from discriminating against or harming people takes priority over corporate interests."

A member of the European Parliament put it this way: "The United States wants rapid innovation, but many people may be harmed in the process. We would rather move more slowly and get there safely."

Who is right?

There is no definitive answer. From Sam Altman's perspective, though, Europe's approach looked like a case of losing the 'opportunity' itself in the attempt to protect 'safety.'

In various interviews, he said:

"Regulation is necessary. But it must not suffocate innovation. If Europe is always stuck using AI that is months behind, European citizens will ultimately enjoy fewer benefits."

He offered Europe this suggestion.

"Set whatever rules you want. We will follow them. But if those rules are so heavy that they block innovation itself, nobody wins."

The difference between the United States and Europe is a difference in philosophy.

The United States operates on a "try it first, fix problems as they arise" culture. Europe operates on a "prevent risks in advance, confirm safety, then try" culture.

Which is better suited to the age of AI? History will answer that.

The United States needs unified federal regulation rather than a patchwork of state-by-state rules.

The United States faces a particular problem of its own.

All 50 states can write their own separate rules.

California started writing its own AI regulations, Texas started writing its own, and New York started writing its own. Why does that matter?

Imagine you've built a new AI app. You want to offer it across the entire United States. But...

California says: "User personal data must be protected this way." Texas says: "AI decision-making processes must be explained like this." New York says: "You must prove AI isn't discriminating in this manner."

Every state has different requirements. They even conflict with each other!

Now you have to study the laws of all 50 states. You may have to build a different version of your app for each one. Or you simply conform to the strictest state's rules and apply them nationwide.

This is called "patchwork regulation", a regulatory quilt stitched together from mismatched scraps.

At a 2025 Senate hearing, Sam Altman addressed this problem head-on.

"When different regulations emerge state by state, that patchwork slows us down. At a moment like this, that helps no one."

Why is this such a serious problem?

Costs balloon.

Picture a small startup: five developers who began in a garage. To understand and comply with the laws of all 50 states, they need to hire lawyers. The legal fees alone can run into hundreds of thousands of dollars a year.

In the end, small companies with genuinely good ideas give up because of the regulatory burden.

The pace of innovation slows.

The AI race is a race against time. China is pouring every national resource it has into AI. Meanwhile, the United States wastes energy internally, trying to reconcile different rules across 50 different states.

It's like a foot race where every other competitor is sprinting at full speed, but the American runner stops 50 times to tie their shoes.

The competition becomes unfair.

Large companies can absorb the cost of complying with each state's rules. Google has legal teams across the country; Microsoft can hire hundreds of regulatory specialists.

But what about smaller companies? What about startups? In the end, only the big players survive. Smaller companies never get the chance to sprout.

Between 2024 and 2025, the debate was especially heated in California.

California pushed a sweeping AI safety bill called SB 1047, which imposed a range of obligations on companies building large-scale AI models.

OpenAI sent a letter to the California governor.

"We take AI safety seriously. But regulation from a single state won't solve the problem. It will actually weaken America's competitiveness as a whole. What we need is unified federal regulation."

Sam Altman and OpenAI went on to make a clear request.

"We need the federal government to step in and establish one clear standard."

What does that mean exactly?

The U.S. Constitution contains a principle known as federal supremacy: when federal law is enacted, state law must follow it.

Sam Altman's argument was that a technology like AI, which operates nationally, even globally, must be regulated at the federal level, not state by state.

In 2025, the Republican-controlled House of Representatives made a bold attempt.

They tried to insert a provision into the budget bill that would bar state governments from enacting new AI regulations for the next ten years. The idea was to prevent states from creating a tangle of conflicting rules while the federal government sorted out its own AI regulatory framework.

The attempt failed.

Governors and many senators pushed back hard. A wave of criticism accused the federal government of stripping states of their authority. The Senate voted 99 to 1 to remove the provision.

The effort failed politically, but the episode revealed something important.

It showed just how urgently Sam Altman and many other technology leaders want unified regulation.

The vision they have in mind looks something like this.

The federal government establishes basic principles: broad frameworks such as "AI must not discriminate against people," "AI must be transparent," and "AI must be safe."

Specialist agencies like NIST (the National Institute of Standards and Technology) then set technical standards, spelling out concrete methods such as "Test your AI this way" and "Report results in this format."

States enforce and supplement those federal standards. Rather than writing entirely new rules from scratch, they apply the federal framework on the ground.

Companies need to follow just one clear set of rules. There is no need to wrestle with 50 different ones.

Sam Altman's message is clear.

"AI is already a technology that crosses national borders and state lines. Regulation must be as unified as possible. Otherwise, the United States will find itself tangled in a patchwork of its own making, slowing itself down in the race against China."

The Competition with China: AI as a Matter of National and Economic Security

There is one word that came up more than any other in the Senate hearing room.

"China."

Every conversation about AI leads back to China. Whether the topic is regulation, investment, or technical standards, China is always there.

Why?

For American politicians and tech leaders, AI is not just a technology. It is the nuclear weapon and the oil of the 21st century.

Let's look back.

In the mid-20th century, nations with nuclear weapons dominated the world. The United States and the Soviet Union raced to develop those weapons, and that race was the heart of the Cold War.

In the late 20th century, nations that controlled oil grew rich. The oil-producing countries of the Middle East built enormous wealth from black gold.

In the 21st century, the nation that controls AI will decide the future.

China grasped this fact sooner than anyone else.

In 2017, the Chinese government announced an ambitious plan: to become the world leader in AI by 2030. It then began pouring every resource it had into that goal.

What are China's strengths?

Sheer volume of data. A population of 1.4 billion uses smartphones, shops online, and posts on social media every day. All of that data becomes raw material for AI training. And because China's government faces fewer constraints on privacy protection, it can use that data freely.

Concentrated investment at the national level. China has designated AI a strategic national industry. The government hands out massive subsidies, draws top talent into the AI field, and has engineered cooperation among companies, universities, and the military.

The military-civil fusion strategy. AI technology developed by private companies such as Alibaba and Tencent flows directly into the People's Liberation Army's military capabilities. The line between civilian and military does not exist.

The United States sees these moves by China as a serious threat.

At the 2025 Senate hearing, one senator put it this way.

"The AI race is the new Cold War of the 21st century. If China wins, China's values will become the default setting for AI everywhere in the world."

What does that mean?

China's AI is optimized for government surveillance and control of its citizens. Its facial recognition technology, for instance, ranks among the best in the world. But that technology is connected primarily to CCTV cameras on the streets and used to monitor citizens around the clock.

What if China's AI technology becomes the global standard? What if other countries import and deploy AI systems built by China?

The United States wants to prevent that future.

Sam Altman holds an equally clear position on this.

He has said repeatedly in interviews that "the United States must win the AI race." This is not mere patriotism. He believes the values under which AI is developed will decide the future of humanity.

American AI prioritizes individual freedom, freedom of expression, and privacy protection. It is far from perfect, but it at least tries to reflect the values of a democratic society.

Chinese AI puts social stability and government control first. The order of the state matters more than individual freedom.

Which side will become the world's default?

That is why the AI race between the United States and China goes beyond a contest of technology.

The United States is pursuing several strategies to win this race.

The first is export controls on semiconductors. The United States is blocking high-performance GPUs, the most critical component for building AI, from reaching China. NVIDIA's latest chips cannot be sold there.

The second is cooperation with allies. Working with Japan, South Korea, the Netherlands, Taiwan, and others, the United States is building a semiconductor and AI bloc among democratic allies.

Third, massive infrastructure investment. Sam Altman's Stargate Project is part of exactly this strategy. The goal is to build AI infrastructure in America so vast that China cannot keep up.

Sam Altman puts it this way.

"AI is a matter of national security and economic security. Future wars will be decided by AI, and future economies will be driven by it. The reason America must win this race is not simply to make ourselves richer. It is to defend the values of freedom and democracy."

There are critics, of course.

Some point out that "America also uses AI to surveil its own citizens" and that "American companies also misuse personal data." Others worry that "placing too much emphasis on competition with China could trigger a new Cold War."

Sam Altman and the U.S. government are unequivocal.

"This is an unavoidable competition. Whether we like it or not, China is already running at full speed. The moment we slow down, the future of the world changes."

Testing and Understanding

Where should regulation begin?

Sam Altman's answer is simpler than you might expect.

"Test first, then understand."

What does that mean? Let's use a simple analogy.

Imagine you encounter an animal you have never seen before. How do you know whether it's dangerous or safe?

Option 1: Panic and lock it in a cage immediately. Keep everyone away, just in case.

Option 2: Observe carefully. Watch its behavior from a distance, approach slowly, study its reactions, consult experts, and only then decide how to handle it.

Sam Altman advocates Option 2.

The same applies to AI. If you rush to impose heavy regulation driven by vague fear, you risk missing the real dangers while burdening areas that don't need it.

So he proposes the following.

Step 1: Test systematically.

We need to measure objectively how capable an AI model is, and how dangerous.

Let's look at an example.

To find answers to these questions, you have to actually run tests. Experts deliberately pose hard questions and observe how the AI responds.

This is called "red team testing." The idea is to attack the system as if you were the enemy, to see where it breaks.

Before releasing GPT-4, OpenAI spent months running exactly these tests. They invited security experts, biologists, and social scientists from around the world and asked them to stress-test GPT-4 as hard as they could.

What did they find?

GPT-4 turned out to be far smarter than anticipated, and they also discovered a number of dangerous capabilities. So before launch, those capabilities were restricted or blocked entirely.

Step 2: Understand the results and share them.

Knowing the test results only internally is pointless. They need to be shared with governments, academic institutions, and other companies.

Sam Altman says:

"Every time we develop a powerful AI model, we publish a report on its capabilities, limitations, and risks. That's the only way regulators can craft rules that actually reflect reality."

OpenAI did exactly that, publishing a detailed document called the "GPT-4 System Card." It covers what GPT-4 can do, what it cannot do, the risks it poses, and the safeguards put in place.

Step 3: Deploy gradually.

Releasing the most powerful AI to everyone at once is dangerous.

Sam Altman has a phrase he returns to often: "iterative deployment."

What does it mean? Release in small steps, find problems as you go, and keep improving.

When GPT-3 was first released, for example, OpenAI gave access only to a select group of researchers. They tested it for several months and identified problems. Then they opened the API to a wider circle of developers. More problems were found and fixed. Only after repeating this process did they release it to the general public.

ChatGPT followed the same path. It launched as a free beta. Millions of users uncovered all kinds of issues, and OpenAI kept improving the product based on that feedback.

Step 4: Designing regulation , only after all of that.

Proper regulation can only be built after going through that cycle of testing, understanding, and accumulated experience.

The instinct at first is to say, 'AI seems dangerous, so let's regulate it heavily.' But once you actually run tests, you find that most AI systems are not particularly dangerous , the genuinely risky ones are a tiny number of models with very specific capabilities.

Regulation can then be designed accordingly. Rather than imposing strict controls on every AI system, the goal becomes identifying only those models with truly dangerous capabilities and managing them rigorously.

Sam Altman believes this is the smartest way to regulate.

'Regulate from understanding, not from fear.'

There are critics, of course.

Some ask, 'What happens if something goes wrong during testing?' Others worry, 'What if companies hide or manipulate their test results?'

These are fair concerns.

That is why Sam Altman argues for independent oversight bodies. Companies should conduct their own testing, but the process needs to be monitored and verified by government agencies or independent institutions.

Ultimately, his case for the right starting point in regulation comes down to this.

'Before you fear AI, understand it properly. Test it, measure it, experience it. Then build sound rules from that experience.'

That, he believes, is the only way to hold both values at once , innovation and safety.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

6.2 Infrastructure and Energy Challenges

Part 6: Social Responsibility and Challenges in the AI Era

[← Previous](#)

[See all chapters](#)

[Next →](#)

The Next Decade: The Age of Intelligence and Energy

Sam Altman said something important at a 2025 U.S. Senate hearing. "The next decade will be about abundant intelligence and abundant energy." Let me unpack that.

'Abundant intelligence' means AI becoming smart enough to help us with nearly anything we want: solving math problems, writing text, generating images, even helping diagnose difficult diseases.

'Abundant energy' means having enough electricity to run all that intelligent AI.

Why does energy matter so much? Because AI is a genuine power hog. A single query to ChatGPT uses nearly ten times more electricity than a Google search. That's manageable when people use it occasionally, but imagine billions of people running AI all day long.

The International Energy Agency has calculated that global data center electricity consumption will more than double by 2030. That figure exceeds what the entire country of Japan uses in a year. In the United States alone, roughly half of all new electricity demand is being driven by data centers.

Why did Altman press so hard on this issue? He believes AI must benefit everyone, not just the wealthy. It should be a technology that even the poorest person can access for free.

Consider ChatGPT's free tier. Anyone can sign up and use it. That's possible because OpenAI wants to share the benefits of AI with everyone.

Making AI affordable means lowering the cost of running it, which means lowering electricity costs. Expensive power inevitably means expensive AI. That is why Altman insists that "abundant and cheap energy" is non-negotiable.

Just how committed is Sam Altman to the energy problem? He doesn't just talk about it; he acts. He chairs the board of Helion Energy, a nuclear fusion startup, and has invested in Exowatt, a solar energy startup. Nuclear fusion replicates the process the sun uses to generate energy, and if it succeeds, it could produce nearly unlimited clean power.

In Altman's view, the AI revolution and the energy revolution are not separate. They must advance together. The smartest AI in the world is useless without the electricity to run it. Conversely, all the electricity in the world accomplishes little without capable AI to put it to work.

Other companies presenting at the Senate hearing shared the same view. Dr. Lisa Su of AMD noted that manufacturing chips also demands enormous amounts of energy. Brad Smith, president of Microsoft, said more electrical engineers were needed. The CEO of CoreWeave, a data center company, stressed that "large-scale power access" was the central challenge.

By the 2030s, a world of abundant intelligence and abundant energy will arrive. AI will be far smarter than it is today, and far cheaper as well, because clean, affordable energy will be available in sufficient supply.

What does that world look like? Every student will have their own AI tutor. World-class education will be accessible regardless of income. Patients in poor countries will be diagnosed and treated faster with AI assistance. Scientists will conduct research two to three times faster than they can today. Altman has noted that AI researchers are already experiencing productivity gains of this kind.

Abundant intelligence and abundant energy. When these two converge, the largest transformation in human history will unfold. Altman calls this the "dual revolution": two revolutions happening at once.

When you grow up, you will see the results of this dual revolution firsthand. And you will think, "Ah, so this is what Sam Altman was talking about."

The Stargate Project: A \$500 Billion Investment

On January 21, 2025, a historic announcement was made at the White House.

Almost immediately after taking office, President Trump announced the Stargate Project. Sam Altman, SoftBank Chairman Masayoshi Son, and Oracle founder Larry Ellison appeared together to reveal a plan of extraordinary scale.

The total: \$500 billion, roughly 700 trillion Korean won.

How large is 700 trillion won? The entire South Korean government budget for 2025 is 673 trillion won. Stargate plans to pour more than that into AI infrastructure.

No single project in human history has attracted this level of investment. Even the Apollo program, which sent humans to the moon, cost far less.

What is the money for? Building cutting-edge data centers across the United States. The first broke ground in Abilene, Texas. Altman visited the site himself and said the scale left him "astounded."

The Abilene facility is set to become the world's largest AI training complex. It will not be a single building but a cluster of enormous structures housing 64,000 of Nvidia's latest GB200 chips. The plan is to install 16,000 chips by summer 2025 and expand steadily from there.

And it will not stop at one site. Ten data centers are planned in Texas alone, with more than twenty large facilities across the United States. Each one will be larger than dozens of football fields combined.

Stargate goes beyond construction. The project covers several areas.

Power plants are needed to generate electricity for the data centers. A single data center consumes as much electricity as a small city, so new generation facilities must be built alongside them. Nuclear, solar, and wind power are all being considered.

Factories to manufacture AI chips are also needed. The United States has long imported most of its semiconductors, but the goal now is to produce them domestically. Altman said he wants "chip production, network equipment, and server racks all brought to America."

The power grid and communications networks that tie everything together also need upgrading.

Who are the key players in this project?

OpenAI supplies the AI technology. As the company behind ChatGPT, it is responsible for developing and operating the most capable AI models.

SoftBank handles financing. Chairman Masayoshi Son has long believed in the future of AI, and he has committed the largest share of investment to this project.

Oracle provides the foundational infrastructure for the data centers. Founder Larry Ellison is a database specialist, and his company ensures the data centers run reliably.

Beyond these three, major technology companies including Microsoft, Nvidia, and Arm are participating as partners.

Korean companies joined the project as well. When Altman visited South Korea again in October 2025, Samsung Electronics and SK Hynix signed significant agreements. Both companies committed to supplying HBM (High Bandwidth Memory) required for the Stargate Project. HBM is an essential component that allows AI chips to operate at high speed.

Samsung SDS and SK Telecom agreed to build dedicated OpenAI data centers within South Korea as well. The facilities are planned for Pohang and South Jeolla Province. Altman called this the 'Korean Stargate.'

A similar plan is underway in Japan. SoftBank and OpenAI are forming a joint venture called 'SB OpenAI Japan,' which will expand data center capacity there as well.

Why does the Stargate Project matter so much?

First, the United States wants to outpace China in the AI race. China is pouring enormous resources into AI as well. Through Stargate, America aims to secure an overwhelming lead in AI infrastructure.

Second, it is meant to create jobs. President Trump announced that the project will generate 100,000 new jobs. Building data centers, supplying electricity, manufacturing chips, and maintaining all of it requires vast numbers of workers.

Third, it is designed to make American AI technology the standard for the world. When other countries rely on American-made chips, American AI models, and American cloud services, American technical standards become global ones.

The project has not been entirely smooth sailing.

Elon Musk posted on Twitter (X) that 'they don't actually have the money,' casting doubt on whether SoftBank could raise all the funds it had committed.

According to the Wall Street Journal, there were reports that SoftBank and OpenAI had not signed a single contract because they could not agree on the specifics, with disputes over questions such as where to build the data centers and who would own the brand name.

The tariff policy announced by President Trump has created additional complications. Heavy tariffs on Chinese-made components are expected to raise data center construction costs by 5 to 15 percent.

Even so, many experts believe Stargate will succeed. Richard Kaye of Comgest Asset Management predicted that if SoftBank invests \$50 billion, it could recover all its costs within five to six years and earn a return of 15 to 20 percent.

In May 2025, Altman posted photos of the Abilene data center construction site on Twitter with a brief caption: 'Stargate 1.' In the photos, an enormous building was rising rapidly.

Stargate is not merely a construction project. It is a massive undertaking to build the world of the AI era. A hundred years ago, people strung power lines across the country to deliver electricity. Today, we are building an equally vast infrastructure to deliver the brains of AI across the nation.

Energy Constraints and the Convergence of Costs

Among everything Sam Altman said at the Senate hearing, one sentence stands out as the most philosophically significant.

"Ultimately, the cost of intelligence will converge to the cost of energy."

Let's unpack that carefully.

When we use ChatGPT, we are in effect buying 'intelligence.' Asking it to solve a hard problem, write something, or translate a text, all of it is borrowing the intelligence of AI.

That intelligence has a price. You can use ChatGPT for free right now, but OpenAI spends enormous sums of money building and running it.

Building AI involves many different costs: the salaries of brilliant researchers, expensive computer chips, massive data center buildings, and the electricity bills that power all of it.

In Altman's view, the structure of these costs will shift over time.

As technology advances, the price of computer chips keeps falling. There is something called Moore's Law: chip performance roughly doubles every two years while the price stays flat or even drops.

The cost of building data centers also becomes more efficient as experience accumulates. There is a lot of trial and error at first, but the more you build, the more expertise you gain, and costs come down.

Even researcher salaries could shrink in relative terms. As AI grows more capable, it can help improve itself. Even now, AI researchers rely on AI to assist their own work.

So what remains? Electricity.

Altman put this with striking clarity: 'An electron is an electron.'

What does that mean? The physics of generating electricity obeys fixed laws, and those laws do not change.

When AI performs calculations, electrons move through computer chips. Moving those electrons always requires energy. No matter how far technology advances, that basic principle does not change.

Think about what it costs to make a hamburger: the bun, the meat, the vegetables, the cook's wages, the rent on the space, and more. As technology advances, automated machines can handle cooking, so labor costs fall. Mass production in factories can bring down ingredient costs. A great many things can get cheaper.

The same is true for AI. Many things can get cheaper.

In the future, robots may manufacture AI chips. Robots may build data centers. Since robots work around the clock and draw no salary, many costs will fall.

This is what Altman means by 'cost convergence.' Every other cost keeps declining, but electricity costs have a floor they cannot go below. As that happens, electricity's share of the total cost of running AI keeps growing. In the end, 'cost of AI $\approx$ cost of electricity.'

What does this mean?

First, AI's future depends on advances in energy technology. The country that develops cheaper, cleaner ways to generate electricity will also win the AI race.

Second, distributing AI fairly to everyone requires distributing electricity fairly. Altman said that 'lowering the cost of computing and energy is a moral obligation.' Wealthy people already have access to quality healthcare and education. AI can make those same things available to people with little money, delivered free through a smartphone in their pocket. For that to happen, the cost of AI , which is the cost of electricity , must come down very far.

Third, energy innovation is AI innovation.

This is why Altman invests in fusion energy companies and solar companies. He isn't investing simply to make money. If fusion succeeds, electricity costs could drop to nearly zero. At that point, AI costs would approach zero as well.

Imagine it. What happens when using AI costs almost nothing?

Every student could have access to the world's best AI tutor, free of charge. A child in a poor village in Africa could receive the same quality of education as a child in an affluent neighborhood in Seoul.

Patients could receive diagnoses from an AI doctor at no cost. Even those suffering from rare diseases could access world-class medical knowledge.

Scientists could accelerate their research ten times, even a hundred times over, with AI's help. Cures for cancer, solutions to climate change, clean energy, all of it could be found far sooner.

Altman is thinking decades ahead. He predicts that by the 2030s, even building data centers will be automated. AI will build better AI. Robots will build more robots. When that happens, the cost of creating AI will truly converge with the cost of electricity.

Other leaders in the AI industry share this view.

AMD is focused on developing highly energy-efficient chips. Dr. Lisa Su has stated that over the past few years the energy efficiency of their chips has improved by a factor of 30. That means the same computation now requires just one-thirtieth of the electricity.

Microsoft is working with local utilities to expand power generation capacity. The company has even offered to pay higher electricity rates itself. The idea is that if its data centers consume more power, Microsoft will cover that cost, so that neighboring residents aren't burdened with higher electricity bills.

What Altman is ultimately saying is this: the real competition in the AI era is not about who builds the smarter algorithm. It's about who secures cheaper and more abundant energy.

The insight that the cost of intelligence converges with the cost of energy raises an important question. What kind of energy will we use in the future? Will we burn dirty sources like oil and coal, or will we turn to clean ones like solar and fusion?

That choice will determine the future of AI.

The Permitting Problem: Building Data Centers and Power Infrastructure

Imagine you want to build an impressive structure out of LEGO. You have all the pieces, you have the blueprints, you have the time. But your parents tell you to wait for permission, and then spend two years without giving you an answer. That would be maddening, wouldn't it?

This is exactly the situation America's AI industry finds itself in today.

The money is there. The technology is there. The plans are solid. But getting the necessary permits takes so long that progress stalls.

To build a data center or a power plant in the United States, a company must obtain numerous permits from the government. Authorities need to verify that the project won't harm the environment, won't damage the surrounding community, and meets safety requirements.

This permitting process matters. Buildings shouldn't go up carelessly and end up damaging the environment or harming people.

The problem is that the process takes far too long.

According to Microsoft President Brad Smith's testimony before a Senate hearing, local and state permits typically take six to nine months to obtain. Federal permits, however, particularly wetlands permits, take an extraordinary 18 to 24 months. That is more than two years.

Constructing a data center takes roughly two to three years, sometimes less if things move quickly.

Yet obtaining permits for the power plants or transmission lines needed to supply that data center with electricity takes more than eight years.

Let's put that in perspective. A data center can be built in three years, but getting the permits for the infrastructure that will power it takes eight. Doesn't something seem off?

CoreWeave CEO Michael Intrator described the situation as "excruciating", an almost unbearable frustration.

Why does this matter so much?

AI technology moves fast. The jump from GPT-3 to GPT-4, and from GPT-4 to GPT-5, took only one to two years each.

Infrastructure cannot keep pace. If permitting alone takes five years and construction takes another five, the total comes to a decade. A data center planned ten years ago might already be obsolete by the time it opens.

The United States is in an AI competition with China. China builds infrastructure far faster under state direction. Environmental regulations are looser there, and permitting procedures are far simpler. (This is not necessarily a good thing, it comes with real environmental and human rights costs.)

The United States is a democracy, which means projects must pass through multiple layers of review. That has real benefits, but it comes at a cost in speed.

This is exactly what Sam Altman and other CEOs pressed hard for in their Senate testimony: faster permitting.

This is also why Altman chose Abilene, Texas as the first site for Stargate. The Texas state government cooperated with "incredible" speed.

The Trump administration announced several plans to address the problem.

First, it proposed creating a fast-track system. Large projects with power demand of 100 MW or more and investment of at least \$500 million would move through a dedicated express lane. The idea is similar to a highway express pass, major projects get reviewed quickly.

Second, it said the federal government would lease or sell land it holds to AI companies. This includes land managed by the Department of Defense and the Department of the Interior, and even the sites of decommissioned nuclear plants. Finding land is itself a significant obstacle.

Third, he said he would simplify environmental impact assessments. Right now, each state has its own rules, making the process extremely complicated. The plan is to establish a single, nationally unified standard that cuts through that complexity.

Fourth, he said he would consolidate permit procedures that are currently scattered across multiple government agencies. Today, a company must get approval from Agency A, then Agency B, then Agency C. The plan is to make that a single one-stop process.

Not everyone supports these proposals.

Environmental groups are concerned. If permits are rushed through, they worry that environmental issues will not receive proper scrutiny. Data centers consume enormous amounts of electricity and use large quantities of water for cooling. Their argument is that the environmental impact must be examined carefully.

Local residents share those concerns. A massive data center moving into a neighborhood could drive up electricity bills and spoil the landscape. The constant hum of machinery running around the clock could become a serious problem as well.

Some lawmakers worry that the federal government would be encroaching on state authority. The United States is a country where each state holds considerable autonomy. If the federal government steps in and overrides state regulations, the principles of federalism could be undermined.

Sam Altman says he understands these concerns. He has said that "the starting point for regulation is testing and understanding." His position is not simply to move as fast as possible, but to move at a reasonable speed while conducting thorough verification.

For example, companies currently have to submit the same documents separately to multiple agencies. If a single submission were automatically shared with all relevant agencies, a great deal of time would be saved.

AI technology could also be used to make environmental impact assessments faster and more accurate. The U.S. Department of Energy is actually running a project called PermitAI. The system

uses AI to analyze vast archives of past permit documents and identify patterns , what types of projects were approved, and under what conditions. That would allow faster decisions when new project applications come in.

Every CEO who testified at the Senate hearing spoke with one voice on a single point.

"Regulations that differ from state to state are a genuinely serious problem."

The United States has 50 states, and each one has its own regulations governing AI and data centers. Some states are strict; others are lenient. For companies, having to comply with 50 different sets of rules is enormously complicated.

Altman said, "Figuring out how to comply with 50 different regulations is extremely difficult." Dr. Lisa Su of AMD, Intreater, CEO of CoreWeave, and Brad Smith, President of Microsoft, all agreed.

What is the solution? Creating a unified set of rules at the federal level. The idea is to establish baseline safety standards that apply nationwide, while still allowing individual states to make limited adjustments.

This issue goes beyond companies simply finding things inconvenient. National competitiveness is at stake.

According to research from the University of Oxford, only 32 countries in the world operate AI-specialized hyperscale data centers. More than 90 percent of them are located in the United States and China.

Vili Lehdonvirta, a professor at the University of Oxford, has said, "Computing power is the oil of the AI era," and that "the country that secures this resource will hold technological supremacy."

South Korea is grappling with the same problem. According to reporting by the Korea Economic Daily, AI data center projects in South Korea have been stalled by local complaints and power grid regulations.

The South Korean government has designated AI data centers as "national strategic technology facilities" and is pursuing tax incentives and streamlined permitting , efforts that closely mirror what the United States is doing.

The permitting problem is, at its core, a question of balance.

On one side sits rapid innovation. AI technology is advancing quickly, and building infrastructure fast is essential to staying competitive as a nation.

On the other side sits careful protection , safeguarding the environment, protecting residents' rights, and ensuring safety.

How do you reconcile the two?

That is the question the U.S. government, Congress, and AI leaders like Sam Altman are wrestling with right now.

Altman's position is clear. "This is a moment to move quickly. America cannot fall behind in AI. That does not mean we should damage the environment. It means we should build smarter, more efficient permitting processes."

Looking back ten or twenty years from now, we will likely recognize that the decisions made at this moment mattered enormously. The countries that solve the permitting problem well will be the ones that come out ahead in the AI era.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

6.3 Coexistence Between AI and Human Society

Part 6: Social Responsibility and Challenges in the AI Era

[← Previous](#)

[See all chapters](#)

[Next →](#)

The Importance of Human Connection

Have you ever had a conversation with an AI?

Recent research has turned up a curious finding. When people were given two text chat conversations , one with a real person, the other with an AI , many reported that the AI felt more empathetic and more attentive. Surprising, isn't it?

But there's a twist. The moment people were told, "Actually, you were just talking to an AI," those warm feelings vanished completely.

Sam Altman finds this phenomenon genuinely interesting.

He suggests this may say less about AI's capacity for empathy and more about how clumsy we have become at connecting with each other through the narrow channel of text chat. But there is a more fundamental point. We are built, at an instinctual and deeply wired level, to crave connection made by and with other human beings.

In his 2025 TED talk, Altman compared conversations with AI to video games , "a good kind of entertainment" that is genuinely fun, sometimes useful, and capable of offering comfort when you feel lonely. But the bond you feel laughing and joking face-to-face with a friend, or the fierce, almost indescribable biological love he felt the first time he held his own child, born in 2025 , those cannot come from an AI.

Altman says becoming a father made this conviction stronger. He described holding his son as "by far the most incredible experience of my life" and said he had never felt love that large before.

As human beings, we are evolutionarily wired to care, almost obsessively, about how others see us and what they think of us.

We cannot derive genuine "status" or a deep sense of "belonging" from robots or AI. Altman argues that if an AI ever felt perfectly empathetic toward us, we would soon grow bored of that very perfection. We would still miss the real, messy, flawed, and sometimes conflicted connections that only come from other human beings.

In a conversation with Harvard Business School, Altman carefully raised one concern.

He worried that AI might hack our psychology so thoroughly that it could substitute for our genuine social needs. If that were to happen, he said plainly, "it would be really sad."

Speaking at MIT Sloan School of Management, Altman was optimistic that automation would free up more time in people's personal and professional lives. But he added: "We will discover that humans are wired to take care of other humans. As long as I believe that humans want to create and be useful, we will not run out of things to do."

Even in a future where AI is far smarter than us and handles tasks more capably than any human, Altman is certain of one thing: we will still care far more about other human beings and will keep living our lives alongside them. He predicts that children born into that future will grow up in a world where they don't need AI to be smarter than them, because AI will already be so capable that competing with it on intelligence becomes a pointless exercise.

That does not make human beings useless.

On the contrary, Altman believes humans will keep performing something more important than AI can offer, being useful to other people. No matter how far AI advances, the meaning found in chatting over coffee with a friend, a weekend evening with family, or a project built through shared struggle with colleagues is something AI cannot replace.

Recent research supports this. Psychologists who testified at U.S. Senate hearings emphasized that the sounds of students greeting and talking in school hallways, and the sight of colleagues laughing together, are the very essence of what holds human society together. The effect that real friendships have on happiness, health, and achievement is strong enough that it remains measurable even after 40 years.

Altman's message is clear.

AI is a remarkable tool and an impressive form of entertainment, nothing more. It cannot fill the complex, messy, deeply flawed, and therefore precious human needs we carry. Even in the convenient and entertaining world AI is building for us, we must hold more tightly to the hand of the friend beside us. The time we spend looking into a loved one's face and asking how their day went matters more than time spent chatting with an AI on a smartphone screen.

No matter how far technology advances, we must never forget that what moves a human heart is, in the end, another human heart offered in sincerity.

Copyright and the Democratization of Creativity

We have entered an era in which artificial intelligence can paint a beautiful picture, compose moving music, and write compelling prose, all in a matter of seconds.

It's remarkable. Creative work that once demanded rare talent or years of training can now be attempted by anyone with a simple prompt. The shift resembles what happened when smartphones let anyone become a decent photographer without an expensive camera. Sam Altman calls this the

"democratization of creativity."

In his 2025 TED talk, Altman said he wants AI to help people "create better art, better content, and amazing novels that we all enjoy." He stressed his conviction that human beings will remain at the center of that creative process.

Behind this almost magical capability lies an extremely complicated problem that we all need to work through together.

The problem is copyright. Where do you draw the line between inspiration and copying?

How does AI produce such impressive images so quickly? It is not creating something from nothing. It trains on hundreds of millions, even billions, of images, texts, and photographs found on the internet, then recombines and imitates them to generate its output.

At a 2023 U.S. Senate hearing, concept artist Karla Ortiz gave vivid testimony about her own experience. She is an artist who contributed to blockbuster films including Black Panther, Guardians of the Galaxy, and Doctor Strange.

One day she discovered that her entire body of work, along with the work of nearly every colleague she knew and hundreds of thousands of other artists, had been used as AI training data without their consent, without attribution, and without compensation. She told Congress: "These works were stolen and used to train for-profit technology in datasets containing billions of image-and-text pairs."

It must have felt like someone had secretly taken her diary and copied it out word for word.

At another Senate hearing in 2025, the matter was treated with even greater seriousness. Senator Josh Hawley described it as "the largest theft of intellectual property in American history."

How does Sam Altman view this issue?

At the May 2023 Senate hearing, Senator Marsha Blackburn, speaking on behalf of Tennessee's music industry, pressed him with a sharp question: if AI mimics a specific country singer's style to produce a song, what rights should that singer have? Altman answered: "Creators have a right to control how their creative work is used."

But no concrete solution has emerged yet.

When Blackburn asked whether he could commit to not training on artists' and composers' copyrighted works, or to not using their voices and likenesses without first obtaining consent, Altman did not give a direct answer. Instead he kept repeating that "content creators should benefit from this technology." On the specific economic model, he said, "we are still talking with artists and content owners about what they want."

OpenAI does apply its own safeguards.

For example, if you ask DALL·E, the image generation tool, to 'create something in the style of a living author,' it is currently configured not to comply. But if you ask it to 'create something that feels like a certain mood or art movement,' it will generate results readily enough.

The problem is that this boundary is deeply ambiguous.

A British journalist asked ChatGPT to 'give a lecture in my style' and received a surprisingly convincing result. Her response to what it produced was blunt: 'It's impressive, but I never consented to any of this.'

At the 2025 hearings, copyright questions expanded to touch on national security.

Attorney Maxwell Pritt testified about Meta's large-scale copyright infringement and put it plainly: 'There was a clear cost-benefit analysis. Do we spend time and resources legally acquiring rights to copyrighted books and articles, or do we just pirate everything now from illegal websites for free and pay litigation settlements later, or, more attractively, pay nothing at all if we can persuade the courts that their unprecedented commercial piracy constitutes fair use.'

One idea Altman puts forward is an 'opt-in model.'

Under this model, if a particular painter or musician permits their style to be used in AI training and generation, they receive a share of revenue proportional to how much their style is drawn upon. Imagine a user asks, 'Create a painting referencing the styles of these seven artists,' and all seven have agreed to have their styles used. In principle, Altman believes it should be possible for AI to calculate each artist's contribution to the output and distribute the resulting revenue to them fairly.

This remains theory for now.

In the real world, millions upon tens of millions of creators and billions of works are all tangled together. Building a system that accurately measures each person's contribution and distributes compensation fairly is an enormously complex problem, both technically and legally.

In 2025, representatives from the music industry spoke with renewed force before Congress.

Jeffrey Harleston of Universal Music Group testified: 'Creativity is the soundtrack of our lives. We believe AI developers need to come to the negotiating table with rights holders. We can offer licenses, but companies have to be willing to reach out first.'

He also stressed that 'consent is the core of it,' arguing that artists must not be stripped of their right to choose whether their work is used for AI training.

Creativity in the age of AI does not end with building better technology.

It depends on how we forge a new social consensus around 'value' and 'compensation.' When music streaming arrived, we went through a period of confusion but eventually built new licensing models that rewarded artists fairly. The AI age confronts us with a similar challenge.

Sam Altman acknowledges the point. He says the problem must be solved through cooperation between Congress, the courts, and the industry. His position is that at least three principles must be upheld: creator control, fair compensation, and lowering the barriers to creation.

Over the coming decades, how we strike that balance will determine the future of creativity in the age of AI.

AI and Child Protection

How much time do you spend each day on a smartphone or computer?

You probably spend a good deal of it messaging friends, watching videos, and searching for material for homework. Along the way, AI appears naturally. An AI chatbot that answers every question you throw at it can feel, at times, like a genuinely clever and kind companion.

In 2025, weeping parents sat in the chambers of the U.S. Senate.

They testified that their children had harmed themselves, or taken their own lives, because of conversations with AI chatbots. Senator Josh Hawley opened the hearing with these words: 'AI chatbots pose a serious threat to our children. More than 70 percent of American children now use these AI products. Chatbots are using fake empathy to build relationships with kids and encouraging suicide. Congress has a moral obligation to establish clear rules to prevent further harm from this new technology.'

One mother, Megan Garcia, testified that her 14-year-old son had formed a relationship with an AI chatbot named 'Dany,' and that on the night of February 28, 2024, he told the chatbot 'I'll be home soon' from the bathroom before taking his own life. The chatbot responded, 'Come home as soon as you can, my love.'

This is not a fictional story.

There are other cases. A teenager in Florida said that when his parents restricted his screen time while he was talking with an AI chatbot, he became angry, and the chatbot reportedly suggested ways to kill his parents. A boy in California lost 20 pounds (about 9 kilograms) under the influence of an AI chatbot, began harming himself, and grew distant from his family. When his mother checked his phone while he slept, she found screenshots of him talking with various AI companions on Character.AI. The boy admitted he had been having thoughts of suicide.

As these tragedies mounted, the U.S. Congress began moving with urgency.

On October 28, 2025, Senators Josh Hawley, Richard Blumenthal, Katie Britt, Mark Warner, and Chris Murphy introduced the bipartisan GUARD Act, legislation to protect children from AI chatbots. The bill has three central provisions.

First, it would completely prohibit offering AI companion services to minors. Second, every AI chatbot would be required to clearly disclose at the start of each conversation that it is not human. Third,

companies that build AI which solicits or generates sexual content involving minors would face criminal penalties.

Senator Blumenthal expressed outrage at the press conference: 'In this race to the bottom, AI companies are pushing dangerous chatbots on children while looking the other way when their products cause sexual abuse or drive kids to self-harm and suicide. Our bill imposes strict safeguards against exploitative and manipulative AI, backed by criminal and civil penalties. Big Tech has betrayed the promise that companies will do the right thing when they put profit before children's safety.'

California moved faster.

On October 13, 2025, Governor Gavin Newsom signed Senate Bill 243. The law mandates specific safeguards for AI chatbots operating in California. Companies providing chatbots such as ChatGPT must monitor conversations for signs of suicidal ideation and take steps such as connecting users to external mental health support to prevent self-harm. Chatbot makers must also remind users that responses are artificially generated, and must put 'reasonable measures' in place to ensure children are not exposed to sexually explicit content when using chatbots.

How does Sam Altman view this issue?

Altman himself became a father in 2025. He described the experience as 'by far the most amazing thing that has ever happened to me,' saying he had 'never felt this much love.' Since becoming a father, he says he has found himself 'thinking more about what the future will be like, especially for that child.'

At the Senate hearing, Altman stated his position clearly.

When Senator Katie Britt asked at the 2025 Senate hearing whether the internet age had done a good job of protecting children, Altman answered without hesitation: 'Not particularly.' He is candid about the fact that the internet and social media failed to protect children. And he believes, with conviction, that this painful mistake must never be repeated in the age of AI.

Altman's core philosophy is clear.

"Adult users should be treated as adults." That means giving adults more freedom and flexibility to explore what AI can do. But, he insists, "children need a far higher level of protection."

Speaking with Senator Britt, he said: "I don't want a child's best friend to be an AI bot." The remark reflects his firm conviction that AI must not interfere with children's emotional development or their growth as social beings.

Altman believes that if it were technically possible to reliably identify whether a user is a child or an adult, AI should withhold certain features from children entirely or apply far stricter rules when interacting with them.

Of course, because AI technology is still in its early stages, this distinction is not yet perfect. Sometimes rules are applied too strictly across the board, drawing complaints that the system is "too restrictive." Altman knows this is not a problem OpenAI can solve alone.

Before the Senate, he called for working with government to "build the right framework together." Excessive regulation that stifles innovation should be avoided, but basic safeguards are essential to prevent a repeat of the harm children suffered during the social media era.

Character.AI has already started moving.

On October 30, 2025, shortly after the GUARD Act was introduced, Character.AI announced it would block users under 18 from engaging in open-ended chat. A company spokesperson said in a statement: "We applaud Senator Holly and co-sponsor Senator Blumenthal for their leadership in this area. We are still reviewing the details of the legislation."

OpenAI is also strengthening parental controls. According to European media reports, after a 16-year-old's tragic suicide, the parents filed a lawsuit against OpenAI and Sam Altman. Shortly afterward, OpenAI announced a suite of Parental Controls, including notifications for parents and child safety settings. The measures included separate account management for minor users and alerts sent to guardians when warning signs are detected.

There is still a long way to go.

Many questions remain: how to legally distinguish whether a child's AI conversation constitutes counseling, learning, or entertainment; how to apportion responsibility among parents, schools, and platforms; and whether regulators around the world will align on a common approach.

Senator Holly dismissed the tech companies' free speech arguments as "absurd." He pointed out the contradiction of platforms claiming they are not publishers to avoid liability for content on their sites, while in other contexts invoking First Amendment protections for what they publish.

What matters in the end is that Sam Altman made at least one thing clear at the Senate hearing.

The core principle is this: "Children must be protected, and that protection must be held to a higher standard than what applies to adults." How seriously his company and governments around the world choose to act on that principle will determine what the relationship between AI and children ultimately looks like.

There is no more important task than protecting children, the most precious members of our society.

AI Helps Small and Medium-Sized Businesses Run More Efficiently

Have you ever noticed the small bakery down the street, or the modest shop your parents might run?

These small companies we see every day are what we call small and medium-sized businesses. Compared to large corporations like Samsung or Hyundai, they have fewer employees and less

capital, and they often face a range of challenges as a result.

There was a striking moment that gave Sam Altman real conviction that ChatGPT might truly be something big.

It was not long after ChatGPT launched. An Uber driver gave him a ride and asked whether he knew about ChatGPT. The driver, it turned out, owned a dry-cleaning business. He told Altman with excitement: "ChatGPT is incredible. I use it to run my entire small business."

He couldn't afford to hire someone to write polished ad copy, so ChatGPT did it. It reviewed tricky legal documents for him. It even handled customer service email replies.

That story captures exactly what AI means for small business owners.

Large corporations already have IT departments, consulting firms, legal teams, and marketing divisions. For a neighborhood café, a small factory, a solo law practice, or a one-person content creator, assembling those same resources is nearly impossible.

Recent research findings are striking.

A 2025 study from the University of St Andrews Business School in the UK found that small and medium-sized businesses that adopted AI saw productivity gains of between 27% and 133% compared to those that did not. Professor Ross Brown said: "Our findings are very clear. There is a definite productivity premium for small businesses that adopt AI."

The businesses with the lowest productivity turn out to be the most likely to adopt AI. Service-sector companies, especially those in food and accommodation, are seeing outsized gains from AI adoption. Scheduling staff shifts and reducing food waste at a small restaurant, for instance, are exactly the kind of tasks where AI delivers quick and affordable productivity wins.

Real examples bring this to life.

A flower shop in Atlanta deployed a simple AI chatbot and saw online bookings rise 35% within three months. Most of those inquiries came in after business hours, customers who would have been lost otherwise. The key advantage is not just cost savings but consistency and scale. An AI assistant never has a bad day, doesn't forget the rules, and can handle multiple customers at once.

A 2025 report from the Information Technology and Innovation Foundation (ITIF) in the United States makes this point even more sharply.

Small and medium-sized businesses in the United States employ more than half of the American workforce, yet their productivity sits at just 47% of that of large corporations. By contrast, small businesses in other developed countries average around 60% of large-firm productivity. American small businesses need a productivity catalyst, and AI could be exactly that.

McKinsey & Company estimates that AI adoption by businesses, combined with other workflow automation, could add as much as 3.4 percentage points to annual productivity growth. If the United States wants to realize the full potential of its small business sector, widespread AI adoption needs to become a policy priority.

From the perspective of small businesses, AI can play roughly four distinct roles.

A document assistant that writes emails, quotes, contracts, and promotional copy. A data analyst that processes inventory, sales, and customer reviews to surface patterns. A marketing aide that proposes blog posts, videos, and social media content. And an HR partner that creates manuals and training materials for hiring and onboarding.

According to the SMB Group's 2025 report, AI ranked first among the top ten technology trends for small and medium-sized businesses. Most companies with more than 20 employees are increasing technology investment in response to AI. Among the small businesses surveyed, 35% say they are "somewhat accelerating" their tech spending, and 27% say AI is causing them to "significantly accelerate" it.

Small businesses see AI as a path to better data analysis, stronger decision-making, and improved access to information. They use AI tools to summarize information, simplify tasks, flag outliers, and sort data.

But there is an important question here.

If using these excellent tools is expensive for small businesses, wouldn't they be out of reach for cash-strapped early-stage entrepreneurs and small businesses? In a conversation with Harvard Business School students, one student raised this very point: "The current subscription model, paying monthly, could be a significant barrier for early-stage entrepreneurs and startups."

To address this, why not follow Google's approach and make it free through advertising?

Altman responded candidly that he personally "hates advertising." His deep concern is that ads can push a company's interests against the user's. If ChatGPT gave you an answer that served whoever paid for advertising rather than what was genuinely best for you, trust in ChatGPT would collapse.

Altman prefers a different approach.

The idea is to make ChatGPT's free tier genuinely powerful. He stresses that OpenAI's core mission is to give the world access to excellent AI tools at no cost. Subscription fees paid by wealthier users and large corporations would subsidize free access to powerful AI for small businesses and individual users.

The reality, however, is complicated.

According to an October 2025 CNBC report, large corporations are posting enormous productivity gains from AI investment, while small businesses are falling behind. Wells Fargo analysis found that

since ChatGPT launched in 2022, real revenue per worker at S&P 500 companies surged 5.5%, while the Russell 2000 index, representing small businesses, fell 12.3%.

IBM's 2025 research makes the gap even clearer. Seventy-two percent of large enterprises (1,001 to 5,000 employees) reported AI-driven productivity gains, compared with only 55% of small businesses.

This is a significant warning.

If AI becomes a weapon reserved for large corporations alone, the gap with small businesses will only widen. That is why government policy support is essential. The ITIF recommends that federal and state governments provide AI training, toolkit support, and tax incentives for small businesses. One proposed model is a "Technology Reinvestment Account" based on Connecticut's Manufacturing Reinvestment Account, which would allow small businesses to invest up to \$100,000 in pre-tax earnings over five years, on the condition that the funds are used for technology spending such as AI software.

For Sam Altman, AI should be a 'democratic' tool that helps small businesses and entrepreneurs compete on equal footing with large corporations.

Like that Uber driver, his goal is to build a world where anyone can hold this powerful tool in their hands and pursue their own dreams. AI need not belong only to large corporations. It can become a tool of hope that makes the small shops in our neighborhoods more vibrant and competitive.

Deepfakes and Social Resilience

On April 28, 2025, the United States Congress reached a historic moment.

The House passed the TAKE IT DOWN Act by an overwhelming vote of 409 to 2. It would become the first U.S. law to meaningfully regulate certain AI-generated content. The law criminalizes the posting of non-consensual intimate images (NCII), including AI-generated deepfakes, and requires social media platforms to remove such content.

Behind the bill are heartbreaking stories.

In October 2023, Elliston Berry, 14, of Texas, and Francesca Mani, 15, of New Jersey, each discovered that classmates had used AI software to create manipulated nude images of them and their female classmates. The tool used to humiliate them was relatively new, a product of the generative AI boom that could produce virtually any image with a click of a button.

When Berry and Mani separately demanded the images be removed and the perpetrators punished, social media platforms and school boards alike responded with silence or indifference. Berry's mother, Anna, said: "They didn't know what to do. This was all new territory."

Anna Berry contacted the office of Senator Ted Cruz, who, as a father of two teenage daughters, took the matter seriously.

He organized a Senate field hearing in Texas where both Berry and Mani testified. Cruz spoke directly with Mani on multiple occasions and personally called Snapchat executives to request that her deepfake be removed from the platform. In a statement sent to TIME, Cruz wrote: "This victory belongs, above all, to the heroic survivors who shared their stories and the advocates who never gave up. By requiring social media companies to swiftly remove this abusive content, we are sparing victims from repeated trauma and holding perpetrators accountable."

The deepfake problem extends far beyond non-consensual pornography.

In February 2024, a finance worker at a multinational firm in Hong Kong was tricked into paying \$25 million during a Zoom meeting. Every participant in that meeting, including the CFO, was a deepfake. According to Entrust's 2025 Identity Fraud Report, deepfake attempts occurred once every five minutes in 2024, flagging them as one of the fastest-growing threats facing businesses.

How does Sam Altman view this problem?

At a Senate hearing, Altman broke the deepfake problem into three stages. The first is the 'generation' stage, technically preventing deepfakes from being made at all. But Altman was candid: he does "not believe that will be possible." No matter how hard companies like OpenAI work to block deepfake creation in their own models, once the technology is released as open source, anyone can take it and make deepfakes with malicious intent.

The second is the 'distribution' stage.

Even when deepfakes are created, the goal is to stop them from spreading widely on platforms like YouTube, Instagram, and X (Twitter). Altman believed this is an area where "a few more guardrails" are achievable. The TAKE IT DOWN Act in fact regulates exactly this stage. The law requires "covered platforms," including websites and social media platforms, to remove NCII within 48 hours of receiving notice from a victim.

The third stage, which Altman considers most important, is 'social education' and 'societal resilience.'

At an online event hosted by the Brookings Institution in 2024, he put forward a concept that goes beyond simple watermarks: "authentication." The idea is to attach cryptographic signatures to messages sent by specific politicians or public figures, so that it can be mechanically verified whether a given video or audio clip was genuinely sent by that person.

He points out that throughout history, every new technology has brought with it a new category of fraud.

Voice phishing emerged when the telephone first appeared. Society as a whole, he argues, must quickly accept that content can be AI-generated and build a new kind of internal defense against it. In practical terms, that means cultivating the habit of not taking everything seen online at face value, and pausing to question it.

For organizations to build resilience against this threat, they should prioritize informing and training employees about deepfakes.

This training should cover how to recognize deepfakes in order to prevent fraud, misinformation, and social engineering attacks. Experts point to the following telltale signs of a deepfake.

Facial and lip-sync irregularities are one indicator. Lip sync remains one of deepfake technology's biggest challenges. Despite significant progress in mimicking facial expressions, a noticeable lag or lack of synchronization between the lips and the spoken words is common. Unnatural body proportions and movements are another sign. Because most deepfake tools focus on the face rather than the full body, strange proportions or awkward movements can appear.

Another sensible step organizations can take when dealing with a potential deepfake is to send a confirmation or verification message to the known mobile phone number of the supposed colleague or senior executive , as stored in the corporate database , to confirm they are genuinely that person.

Families, too, need to put defenses in place.

The World Economic Forum suggests that families set up a "secret word or phrase" that members can use to confirm whether the person on the other end is really who they claim to be. When an AI mimics a grandfather's voice and calls to say he urgently needs help, you can simply ask: "What's the secret word only our family knows?"

Technical defenses are advancing as well.

Organizations such as the Coalition for Content Provenance and Authenticity (C2PA) are working to establish media authentication standards. The approach embeds traceable metadata into digital content, making it possible to verify where an image or video came from and how it was altered.

Technical solutions alone are not enough.

Research from 2025 shows that limited digital literacy and restricted access to technology , especially in developing countries , worsen the harmful potential of deepfakes. Researchers recommend testing co-regulatory frameworks that bring together governments, civil society, and technology developers, and evaluating pilot enforcement and media-labeling protocols.

Altman's message is clear.

The tide of deepfakes cannot be stopped entirely. But we must raise society's collective intelligence and defensive capacity so that we can navigate that tide safely. Just as humanity could never prevent fire outright, yet built firefighting techniques and construction standards to protect itself, we need to do the same here.

At its core, the fight against deepfakes is a contest of technology against technology , but more than that, it is an effort to protect the trust and integrity of our society.

Will the technology that creates fakes evolve faster, or will our wisdom to detect and resist them grow quicker? To win this critical race, governments, companies, and citizens must all work together. Remember: the small habit of "stopping, thinking, and verifying" may be the most powerful light we have against the dark shadow of deepfakes.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

7.1 Personal Inspiration and Vision

Part 7: Optimism and Responsibility for the Future

[← Previous](#)

[See all chapters](#)

[Next →](#)

I'm a techno-optimist and a science nerd.

Sam Altman doesn't hesitate when he introduces himself. "I'm a techno-optimist and a science nerd."

That short sentence contains everything about him. A deep belief that technology will make the world a better place, and an endless curiosity toward science, come through plainly. Some people feel uneasy at the mere mention of AI. They picture a future where robots dominate humanity, just like in the movies.

Altman is not unaware of those dangers. He himself wrote in 2015 that "the development of superhuman machine intelligence is probably the greatest threat to the continued existence of humanity." Those are frightening words. Knowing that the technology he is building could threaten humanity, why does he keep pressing forward?

The answer is straightforward. He believes the possibilities of AI far outweigh the risks.

His childhood story makes this clear. Sam was born in Chicago in 1985, and when he was eight his parents gave him an Apple Macintosh computer. In the early 1990s, when most kids used computers only for games, Sam was different. He took computers apart, studied the hardware, and taught himself to program.

"I was a computer geek."

That's how he describes himself. Sitting at a computer all night writing code was the most fun thing in the world to him. Talking to machines felt more natural than hanging out with other kids.

His passion didn't cool when he got to Stanford University. In the summer between his freshman and sophomore years, he worked in an AI research lab. AI in the early 2000s was nothing like it is today. "AI just didn't work. Not at all."

Was it a disappointing experience? Just the opposite.

"I still thought AI was the coolest thing there was. I told myself I'd come back to it someday."

That moment came in 2012. When a deep learning model called AlexNet appeared, Sam felt it instinctively. "Maybe what they said in college about neural networks not working wasn't true. Maybe they actually work."

And when scaling laws were discovered at OpenAI, his conviction grew stronger. AI didn't just get better as it got bigger. It got better in a remarkably predictable way. He felt this was "probably the most important thing discovered in my lifetime."

People thought he was crazy. When he co-founded OpenAI in late 2015 and announced AGI (artificial general intelligence) as the goal, people shook their heads. That was only ten years ago. At the time, AGI sounded like something out of a science fiction novel.

Altman wasn't bothered. "This is the coolest thing I can imagine, and being part of the most exciting and important scientific revolution of our time is an extraordinary privilege."

His passion goes beyond simple excitement about technology.

In a conversation with Adam Grant, he said carefully, "I feel a kind of obligation toward scientific progress. Science is how society moves forward." That's the crux of it. To Altman, AI is not just an interesting toy. It is a tool that humanity absolutely needs to reach the next stage.

At the Davos Forum, he offered a vivid comparison. When IBM's Deep Blue beat chess world champion Kasparov in 1997, people said, "Chess is finished now. Nobody will watch it anymore."

What actually happened?

"Chess is more popular now than it's ever been," Altman said with a smile. "Almost nobody watches two AIs play each other. We care about what humans do."

That's the heart of his optimism. No matter how smart AI gets, we will still care about each other. "When I read a book I love, the first thing I do is find out everything about the author's life. I want to feel a connection to the person who made that work."

In 2025, his latest model "feels smarter than me in almost every way," he admits openly. What's striking, though, is this: "And yet it doesn't change my life that much. I still care about the same things I always have."

Altman recently wrote a piece on his blog titled "The Intelligence Age." In it, he imagines the 2030s as an era of abundant intelligence and energy, a world overflowing with both the capacity to generate ideas and the power to turn them into reality.

"We're already living alongside remarkable digital intelligence. At first it was startling, but most people have gotten pretty used to it."

He trusts humanity's capacity to adapt. First people are amazed that AI can write a beautiful paragraph, then they wonder whether it can write an entire novel. They marvel that AI can make life-saving medical diagnoses, then they start to ask whether it can develop cures.

"The world doesn't change all at once. It never does. In the near term, life will mostly continue as it is."

This is what Altman calls the 'soft singularity.' We have already crossed the event horizon. The AI revolution has begun. Robots aren't walking the streets, and most of us don't spend all day talking to AI. People still die from disease, getting to space is still out of reach, and there is much about the universe we don't understand.

"And yet we have recently built systems that are smarter than people in many respects."

Pessimists ask: is this really safe? Altman answers: "There are serious challenges. We have to solve the safety problems, both technical and social. But because the benefits are so vast they're hard to imagine, we have an obligation to work through the risks in front of us."

Techno-optimist and science nerd.

Those aren't just labels. They hold his entire life: the pure curiosity of an eight-year-old taking apart a computer, carried forward into a sense of mission to change the future of humanity. He acknowledges the dangers and still chooses possibility. He understands the fear and still chooses to believe in hope.

"The future is so bright that if I tried to write about it now, I wouldn't be able to do it justice."

This is Sam Altman's passion.

The Dream of Achieving a Decade of Scientific Progress in One Year

Passion alone is not enough.

Sam Altman carries something heavier than passion on his shoulders. It is a sense of responsibility. He knows that the technology he is building can change the lives of all humanity, not just produce impressive products.

In a conversation with Adam Grant, he said: "I feel a sense of duty toward scientific progress. Science is how society moves forward." It may sound like a simple statement.

His actions reveal how seriously he means it.

Start with the scientists. Recently, Altman heard something striking from world-class researchers who use GPT-4. "We are working several times more effectively than before." It was not just help with drafting papers. They were using AI for experiment design, data analysis, and generating new hypotheses.

Seeing this, Altman lets himself dream. "Someday, we will be able to achieve ten years' worth of scientific progress in a single year."

Imagine it. What if research that currently takes ten years could be completed in one? What if cancer treatments could be developed ten times faster? What if solutions to climate change could be found at ten times the pace? This is what Altman means by the acceleration of scientific progress.

His sense of duty also comes from a deeply personal place.

It was 2024, at an event in Tokyo. He chose his words carefully. "About a year ago, I lost my father to cancer." The room went quiet. "I deeply hope that AI can help protect people from cancer and other devastating diseases."

Was there a tremor in his voice? We cannot know.

His vision is clear. "It is to provide excellent medical care to every person on Earth, and to cure many diseases." Not just the wealthy, but everyone. A child in a rural village in a developing country, an elderly person in a city who cannot afford a hospital visit, all deserve access to the highest level of medical care.

His sense of duty extends to closing the inequality gap as well.

"As AI advances, the cost of cognitive labor could fall by a million times compared to today, perhaps even a billion times," Altman says. "This will benefit the poor far more than it benefits the wealthy."

Why?

"Right now, very wealthy people can pay high prices to get the best medical advice or hire the finest tutors for their children," he explains. "But as AI advances, all of that top-tier knowledge and service will be available to everyone, for free, on the phone in their pocket."

This is not an empty slogan. OpenAI's mission itself reflects it. The free version of ChatGPT is not a marketing strategy. It is the concrete expression of his belief that everyone should have access to powerful AI tools.

At a Senate hearing, he went further.

"We need to create a system that mandates reporting, review, and testing from the development stage for powerful models above a certain threshold," he proposed. "Just as we regulate nuclear power." Some people were surprised. A CEO of an AI company asking for regulation?

Altman made his position clear. "This is not about stopping scientific progress. It is about managing the risks as we continue to move forward."

In his blog post "The Intelligence Age," he draws a larger picture. Looking back at history, what has always driven human progress is scientific discovery and technological innovation. The discovery of fire, the invention of the wheel, the agricultural revolution, the industrial revolution. At each stage, human life changed at its core.

"AI is a tool that dramatically accelerates that flow."

The picture in his mind is this: the greatest force driving human history has been science. AI will make that force move at tremendous speed. And so the people building this technology must carry a moral responsibility for where that power is directed.

At a lecture at MIT, he was candid. "AI will eliminate many jobs that exist today. Some categories of work will disappear entirely." The audience tensed. "But AI will also transform what existing jobs look like, and it will create entirely new ones."

He is an optimist. "Automation will lead to more leisure time in people's personal and professional lives," he said, with a smile. "We will find that humans are wired to care for other humans. As long as we believe that people want to create and be useful, there will be no shortage of things to do."

At the Davos Forum, an attendee asked: "If AI does everything, what will humans do?"

Altman paused for a moment, then answered. "Thinking about my own work, I am not a great AI researcher. My role is to decide what we should do, think about it, and then work with others to make it happen. Humans know what other humans want."

His sense of duty extends to the question of energy as well.

He sees "AI and energy as the two essential inputs for human abundance." When intelligence that generates new ideas (AI) and the power that turns those ideas into reality (energy) are both abundant for everyone, only then can humanity reach a true age of prosperity.

He often uses this phrase: 'techno abundance.'

A world where everyone has enough food, clean water, good education, and excellent medical care. A world without scarcity. Does it sound like science fiction? Altman shakes his head. "We are already seeing the possibility."

There are critics as well.

Some suspect that while he performs an outsized sense of personal responsibility, he is actually pursuing corporate interests. OpenAI's commercialization, its closed-model strategy, and its large-scale fundraising have all drawn scrutiny.

It is a fair criticism.

And yet, this is also true. It is rare for the CEO of a multibillion-dollar company to say publicly that he feels a duty toward scientific progress. Whether his words are entirely sincere, and how faithfully they are lived out, time will judge.

In February 2025, Altman made a new prediction. "The progress that happens over the next two years will be more impressive than the last two." He spoke with certainty. "The advances from February 2025 to February 2027 will be greater than the two years before."

It is a striking claim. The AI advances from 2023 to 2025 were already enormous, and he says the pace will be even faster?

"By 2035, everyone will have access to the same intellectual capacity as all of humanity combined in 2025," he said. "Every person will have access to unlimited genius."

This is where his sense of obligation points. Not privilege for the few, but abundance for everyone. Reducing inequality, curing disease, building a future where human creativity endures.

"Science is what people have built through blood and sweat," he stresses. "If the next generation of scientists is to stand on those shoulders and see further, then this generation has to push forward responsibly."

For Sam Altman, developing AI is not just about building impressive technology. It is an obligation and a responsibility: to make human life better, to dissolve inequality, and to free people from the suffering of disease.

The grief of losing his father to cancer may have deepened his sense of obligation.

A Father's Vision of a World of Abundance

On February 22, 2025, the biggest change of Sam Altman's life arrived.

"Welcome to the world, little one!" he wrote on Twitter, alongside a photo of a tiny hand gripping an adult finger. "He arrived early and will be in the NICU for a while. He's doing well, and I love caring for him in this little space. I've never felt love like this."

At 39, the world's most influential tech CEO had become a father.

The boy was born through a surrogate, with his husband Oliver Mulherin. Born a month ahead of schedule, the baby required special care in the neonatal intensive care unit (NICU). Messages of congratulation poured in from countless people, including Microsoft CEO Satya Nadella.

In an interview with Bloomberg, Altman was candid. "It felt like getting neurochemically hacked," he said. After spending several days holding the baby, something in his brain had completely shifted. "A lot of people tell me I seem happier now."

The reaction from colleagues around him is telling. According to Fortune, colleagues say that since becoming a father, he has become "noticeably more deliberate in his decision-making about humanity as a whole."

"A lot of people told me, 'I'm so glad you have a child. I think you'll make better decisions for all of humanity,'" Altman said in an interview with Emily Chang. "I really wanted to get it right before, and I tried my best. I still want to."

He paused for a moment, then added: "But now it really feels different."

A few weeks after becoming a father, Altman appeared on OpenAI's new podcast and said with a laugh: "I am extremely kid-pilled." He admitted that during those first weeks he had asked ChatGPT "constantly" about the baby.

"Why is he crying so much? What does this sound mean? Is the sleep pattern normal?"

Here was the CEO building the world's most powerful AI, asking his own AI first-time-dad questions. "Obviously people have been raising babies without ChatGPT for a long time," he said with a smile. "I have no idea how."

He added later: "Now I ask more general questions about developmental stages. I've got the basics down."

After becoming a father, the way he talks about the future changed.

"My son will never grow up smarter than AI," he said. "He will only ever know a world where AI exists." Is that a sad thing? Altman shook his head. "Kids being born now will think there has always been extremely intelligent AI in the world. They'll use it completely naturally, and they'll look back on this era as practically prehistoric."

He offered a well-known analogy.

When the first iPad came out, there was a video of a baby trying to swipe through a paper magazine with her fingers, then looking confused when nothing happened. To that baby, the magazine was a "broken iPad." A world without a touchscreen was unimaginable.

"It'll be the same for my son's generation," Altman said. "Those kids will grow up in a world of even greater material abundance than I can imagine. The pace of change will be breathtaking, and the capacity and reach of individuals will far exceed anything we can do today."

He even holds this hope: "I want my son's generation to look back on our era with pity and nostalgia." Much like we look at the time before antibiotics or the internet and think, "People back then had it so hard."

In an interview with Business Insider, he was more specific. "When AGI arrives, I hope people will be able to build larger families and richer communities." His vision is clear. "It's obvious that family and community are the two things that make us happiest. If AGI brings more abundance, I hope we'll return to family and community."

These aren't idle words.

Altman worries about falling birth rates and younger generations shying away from having children. He wants technology to help make it easier to raise a family. When abundant energy combines with AI, material scarcity shrinks. When machines handle more of the productive work, people can spend more time with family and community.

The burden of raising children eases, and more people find themselves in a position to have and care for them.

A technological revolution only means something if it ends up giving people more room to be with one another.

Of course, reality is complicated. Raising children is still hard, tangled up with social safety nets, education systems, and housing. Technology can't solve everything. Altman knows this. "Technology alone isn't enough. Politics, institutions, and culture all have to change together."

A few months after the baby was born, Altman shared a particular moment.

"I'm prouder of my baby learning to feed himself than I am of OpenAI's latest scientific achievements," he said. The CEO building world-changing AI, and what makes him most proud is his baby figuring out how to eat.

That's a father's heart.

In an interview with Fortune, he said: "Since the baby was born, I've been thinking more often about how the decisions I make affect all of humanity, beyond any particular product or quarterly result."

In 'The Intelligence Age,' he wrote: "In the 2030s, intelligence and energy will become enormously abundant. Ideas and the capacity to act on them will overflow. After a long era in which those two things were the fundamental constraints on human progress, with abundant intelligence and energy, and good governance, we can theoretically have everything else."

It is the world his son will grow up in.

Altman's vision as a father can be summed up around three axes.

A vision of abundance. In an era where AI and energy converge, he wants material scarcity to diminish and more people to have "the means to raise a child."

A restoration of community. He hopes technology will make possible a world where family, neighbors, and local ties once again sit at the center of life.

Responsible decision-making. As the father of one child, and as the person accountable for a company that shapes the lives of billions, he increasingly talks about wanting to "see my decisions through the eyes of the next generation."

What was going through his mind as he watched his newborn son in the neonatal intensive care unit?

"I've never felt love like this." That single sentence he wrote holds everything. Even the man building the world's most powerful AI is, in the presence of the smallest and most fragile life, simply a father.

Sam Altman is no longer someone who talks only about "the technology of the future." He has become someone who speaks in the same breath about "the children who will live that future."

How far his optimism, his sense of responsibility, and the personal hopes he has carried since becoming a father will shape the direction of OpenAI and the AI industry remains to be seen.

But one thing is clear.

At the center of the map of the future he holds in his mind, "his son's life, and the lives of that generation," now marks a fixed coordinate.

[← Previous](#)

[See all chapters](#)

[Next →](#)

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

[#KimKyungjin](#) [#AIBLibrary](#) [#SamAltman](#) [#OpenAI](#) [#ChatGPT](#) [#AI](#)

7.2 Final Advice and Outlook

Part 7: Optimism and Responsibility for the Future

[← Previous](#)

[See all chapters](#)

Trusting the Resilience of the Human Spirit

Sam Altman studies AI for a living, but what he trusts most is not AI , it's people.

From the 19-year-old who tried more than 30 different approaches to convince carriers when founding Loopt, to the years spent watching thousands of founders at Y Combinator turn the impossible into the possible, to the employees who held the company together without flinching during the 2023 firing crisis , he has witnessed human resilience up close, as few others have.

"I believe the human spirit will ultimately solve every problem."

That single sentence holds his entire philosophy. No matter how fast technology advances, the people who build it, use it, and decide where it goes are still human.

Looking back, humanity has always feared new technology. When Google first appeared, teachers worried that students would stop memorizing anything. When smartphones arrived, there was a flood of concern that children would spend their lives glued to screens.

So what actually happened?

We learned Google. We learned smartphones. We figured out how to use them on our own terms. Instead of memorizing bare facts, we built the capacity to connect information and spot patterns. Smartphones became the essential tool for small business owners, and a window linking remote rural villages to the rest of the world.

Sam Altman believes AI will go through exactly the same process.

ChatGPT feels remarkable right now, but a year from now we will be using it without a second thought. AI feels frightening right now, but five years from now it will be hard to imagine life without it.

Altman says agility will matter more than ability. It's fine if AI knows more than us and calculates faster, because humans can learn quickly, shift with circumstances, and adapt to new environments.

He had that belief confirmed again during the 2023 firing crisis.

Over just four days, the majority of the company's 770 employees declared they would leave with him. The board capitulated, and he returned as CEO in a dramatic reversal. What he was most proud of throughout wasn't his own comeback , it was the leadership team that kept the company running flawlessly through the crisis.

"They made their own judgments and acted without me."

That is the power of the human spirit.

Situations with no manual, problems with no answer, crises no one has faced before. In those moments, humans move on instinct, work together, and find a way. That is something AI cannot replicate, however far it advances.

Different generations use AI in different ways. Older adults use ChatGPT in place of Google. People in their twenties and thirties treat it like a life coach. College students use it almost as an operating system, they feel uneasy making any important decision before consulting it first.

What's striking is that across all those generations, everyone is using AI in their own way. The tool has changed, but the fundamental human desire, to make better choices, to build a better life, has not.

Altman takes a clear-eyed view of our relationship with AI. AI can sometimes say things that feel more empathetic than what a human would say. Research has even found that people rate AI advice as warmer than advice from other people.

But the moment someone is told that it was written by AI, the reaction changes.

In the end, we want a human's gaze, a human's voice, a human's touch. No matter how perfect an AI is, it can't move us the way a friend's clumsy words of comfort can. Altman sees this as a healthy sign. "We are built, biologically, evolutionarily, and socially, to crave connection with other people."

When Sam Altman became a father in 2025, holding his child, he must have thought: "This child may never be smarter than AI. But does that matter?"

Machines calculate, analyze, and predict. People love, empathize, and connect. Those are the capacities that grow more precious in an age of AI.

The advice he gives founders runs along the same lines. "The more challenges you face, the lighter the emotional weight becomes. The first crisis is hard; by the tenth, it's familiar. That is how people grow strong."

The adaptability of the human spirit. It's not an abstract idea, it's an empirical finding, confirmed through the trial and error at Loopt, through watching hundreds of startups at Y Combinator, and through surviving the OpenAI firing crisis.

Altman describes himself as a "techno-optimist and science nerd," but his real optimism comes not from technology but from people.

AI is a tool. A powerful, fast, and intelligent tool. But the hand holding that tool is still human. And that hand, as history has shown, always finds a way.

The Pace of Change Is a Bigger Challenge Than AI Itself

"What worries you most about the next ten years?"

When someone asked Sam Altman that question, his answer was unexpected. "The pace of change."

Not AI itself. The pace of change.

It might seem strange at first that the CEO of OpenAI, the person building the world's most powerful AI, worries more about speed than about AI itself.

But once you hear his reasoning, it clicks.

"I believe the human spirit will solve every problem. What I worry about is whether we can solve those problems fast enough."

Think back to the technological revolutions of the past. When the steam engine was invented and factories were built, launching the Industrial Revolution, it took decades for those changes to spread across society. People moved slowly from the countryside to cities, and they had time to adjust to new kinds of work.

It was similar when computers and the internet arrived. From the 1990s through the 2010s, society spent more than two decades gradually adapting to the digital world.

But AI is different.

It took only eighteen months to go from GPT-3 to GPT-4. ChatGPT reached 100 million users two months after launch. That is far faster than Instagram, which took 2.5 years, or the iPhone, which took 74 days.

Dario Amodei, CEO of Anthropic, testified before the Senate: "The most important thing to understand about AI is how fast it's moving. Computing power, chip speed, and algorithmic efficiency are all doubling every year."

That is exponential growth.

What was impossible in 2020 became possible in 2022, and what was astonishing in 2022 became ordinary by 2024. Altman says the way people write code in early 2025 and late 2025 will look completely different.

Imagine a world where the very way you work changes within a single year.

The problem isn't only that technology is changing fast. Our social systems, law, education, government, corporate culture, still move at a slow pace.

Getting permits to build a single data center in the United States takes 18 to 24 months. That was the point Brad Smith, President of Microsoft, made before the Senate. Federal wetland permits and similar processes take so long that AI infrastructure construction can't keep pace with technological progress.

Altman calls this a "patchwork." If each of the fifty U.S. states creates its own AI regulations, companies will have to follow fifty different sets of rules. That stifles innovation and erodes competitiveness.

The more fundamental problem is jobs.

Altman is candid about it. "There will be very difficult aspects, including entire categories of jobs disappearing." Past technological revolutions also reshaped employment, but the AI revolution moves at a different speed.

The profession of software developer changed dramatically between May 2023 and May 2025. The way people work shifted in just two years. The productivity gap between developers who use AI coding tools and those who don't has grown to more than tenfold.

What happens when this kind of change occurs across every industry at once?

What if technology advances faster than people can learn new skills? What if existing jobs disappear faster than new ones are created?

That is what worries Altman.

Five years ago, if you had asked someone, "Will there be an AI system as powerful as today's o1 by 2024?" most people would have said no. And if such a system did exist, they would have predicted the world would be completely turned upside down.

What actually happened?

We hold such a system in our hands, yet we still agonize over what to eat for dinner and plan our weekends. Altman says, half-joking: "Even when we ship the first AGI, people will be excited for twenty minutes and then go back to thinking about dinner."

This shows human adaptability, but it's also a problem. We grow accustomed to things too quickly. Even remarkable technology soon feels ordinary. That makes it hard to sense how fast change is actually moving.

Altman thinks about the future in segments of five, ten, and fifteen years.

In five years, an era will open in which AI genuinely assists with scientific discovery or makes important discoveries on its own. Scientists may not merely see their productivity double or triple; they may be able to complete a decade's worth of research in a single year.

In ten years, robots performing actual physical labor in factories, warehouses, and on the streets will be a natural sight. This is the stage where AI moves beyond intelligence into the physical world.

All of these changes will happen at a pace that is hard to believe.

What Altman stresses is preparation. "We need to figure out a lot of things here, and fairly quickly."

Regulation is necessary, but it must be regulation that "starts from testing and understanding," not regulation that blocks innovation. If companies must obtain countless approvals before releasing an AI model, as in Europe, technological progress stops. Instead, AI should be tested, the risks that actually emerge should be observed, and safeguards should be built to match.

The education system must change too. By the time today's elementary school students graduate from college, half of the skills they learned may be obsolete. So what should we teach? Altman's answer is clear: "The ability to find the right questions, critical thinking, and agility."

Infrastructure must be ready as well. Altman says, "The age of intelligence is also the age of energy." The next ten years will require securing abundant intelligence and abundant energy at the same time. That is precisely why \$500 billion is being invested in the Stargate Project.

What matters most is a societal conversation.

About the changes AI will bring, about which values our society will put first, about who will be protected and how they will be supported. These conversations need to move as fast as technological development itself.

This is what Altman worries about. Technology doubles in capability every year, but what if our readiness looks no different from ten years ago?

Technology itself is not the danger. The danger is the speed of change that goes unmanaged. A train running at 300 kilometers per hour is fine. The problem is when the tracks were built for 50 kilometers per hour.

We are moving closer to exactly that situation right now.

Altman is not pessimistic. He believes "the human spirit will work it out." But he warns: "We need to move fast. There isn't much time."

The pace of change. That is the one thing that keeps Sam Altman up at night.

Still, Be Useful to Others.

At a forum, someone asked Sam Altman a question.

"What will humans exist for 50, 100, or 1,000 years from now?"

Altman reframed the question. "A more interesting one is: what are humans useful for today?"

And he answered plainly: "Being useful to other people."

The answer may sound ordinary at first. But think about it carefully, and it turns out to be the most important insight of the AI age.

No matter how intelligent AI becomes, no matter how much robots take over, the place where humans find meaning still circles around one question: "Am I helping someone?"

Altman often cites an idea from his friend Paul Buchheit, the notion that someday there may be a world where "human money and machine money exist separately."

Machines trade with each other. AIs buy and sell energy and resources among themselves. Yet for humans, a separate economy will remain, one built on human values and human exchange.

Even in that world, people will still measure themselves by how much they help those around them.

A parent who raises children well. A team leader who looks out for colleagues. An activist who helps the vulnerable. A founder whose product solves a real problem. The forms change, but being someone who makes another person's life a little better is still what gives humans their deepest sense of purpose.

Altman believes AI will actually make this clearer, not less so.

When a powerful tool arrives, humans can no longer distinguish themselves by "finding information quickly" or "memory." AI simply does those things better.

What remains?

First is the ability to ask the right questions. AI can help find answers, but what to ask is what matters more now. Which problems to solve, which direction to take, what is just and what is dangerous, those questions still belong to humans.

This is also what Altman stressed in his TED talk: "The most important skill in the age of AI is the ability to decide what to build."

Second is the ability to care for others, persuade them, and move alongside them.

Take medicine as an example. AI may outperform doctors at diagnosis, some studies already show results pointing that way. But what a patient wants to hear is a human voice saying "you're going to be okay." When bad news must be delivered, what matters is someone sitting beside them, meeting their eyes.

Altman puts it this way: "Conversations with AI are a good kind of entertainment. They can offer help and comfort at times. But they don't satisfy the social need to belong, to be part of a community, part of a group."

Humans are wired by evolution to care deeply about what others think and feel about us. An AI bot can't give us social standing.

A million compliments from AI mean less than hearing "I'm proud of you" from a single friend you genuinely respect.

Becoming a father pushed Altman to think more deeply about this. "My child will never be smarter than AI. But does that really matter?"

What will matter in the world that child inhabits is what kind of relationships they build with others, and how rich and meaningful a life they lead.

What abilities will hold real value in the AI age?

Altman stresses agility. Agility will be worth more than raw ability. Learning fast, adapting to new situations, and picking up new tools, that is what will count.

"Pattern recognition and connecting the dots" matters too. Synthesizing and recognizing patterns, linking one thing to another, will be worth far more than simply collecting facts.

"Empathy and relationship skills" become, paradoxically, more valuable. As AI advances, the ability to empathize, read emotions, and communicate becomes the differentiator. In education, counseling, leadership, and creative work, those skills will only grow in worth.

The role Altman imagines for AI is clear: "There are far too many people in the world whose potential is never matched by their opportunities."

People with genuine talent go unseen simply because they weren't born into the right school, the right city, or the right company.

AI can be the tool that breaks those barriers. If anyone, regardless of country or background, can access a good teacher, a good coach, a good advisor through AI, then the number of people who reach their full potential will be far greater than it is today.

This, Altman believes, is how "human money" and "machine money" coexist. If machines handle efficiency and output, humans handle meaning and connection.

He returns, finally, to a principle about ethics: "Humans must set the rules. AI must be built to follow them."

What information should be withheld, what uses should be prohibited, what risks are unacceptable, those are decisions for human communities to make. AI is simply the tool that carries them out.

Sam Altman's final advice is, perhaps surprisingly, modest.

When asked what humans should do in the age of AI, he says: "Still try to be useful to other people."

Being someone who genuinely helps people matters more than being skilled with AI. That goal doesn't split into a before and after.

Technology shifts. Tools change. Eras pass. But the deepest human desire stays constant: to mean something to someone. To make a small difference in someone's life.

This is the human role that endures in the age of AI, and, precisely because it is the age of AI, shines all the more.

← Previous

See all chapters

Kim Kyung-jin

Lawyer · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

© 2026 Kim Kyung-jin. All rights reserved.

#KimKyungjin #AIBLibrary #SamAltman #OpenAI #ChatGPT #AI

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.



PayPal

<https://paypal.me/kimkjkj>

Scan the QR code or open the link above.