

Mayfield X IMAGINATION  
IN ACTION

# Future of AI

Agents  
Startups  
Investing  
Research  
Infrastructure & Hardware  
Energy

Consumer  
Enterprise  
Applied AI  
Health  
Physical AI  
AI Policy & Governance

FALL 2025

# Content

Future of Agents	3
Future of Startups	9
Future of AI Investing	13
Future of Research	17
Future of Infrastructure & Hardware	21
Future of Energy	25
Future of Consumer	29
Future of Enterprise	33
Future of Applied AI	38
Future of Health	43
Future of Physical AI	47
Future of Policy & Governance	51

# Building the Future

We're at the dawn of a new era—one where imagination and intelligence converge to reshape how we build, work, and live.

The *Future of AI* captures a fraction of the insights from *Imagination in Action's* summit in Silicon Valley, where visionaries, founders, scientists, and investors gathered to answer one question: **What will the future of AI look like?**

Every story here is a playbook. Each chapter explores a critical frontier where AI is being built right now: Agents, Startups, Investing, Research, Infrastructure & Hardware, Energy, Consumer, Enterprise, Applied AI, Health, Physical AI, and Policy & Governance.

The opportunity ahead is 100X: Founders can now build enduring companies faster than ever before, create entirely new industries, and solve problems that seemed impossible just years ago.

The stories ahead spotlight the founders and leaders building this future right now.

**Welcome to the Future of AI. It's bigger, faster, and closer than you think.**



**Navin Chaddha**  
Managing Partner, Mayfield

**John Werner**  
Founder & CEO, Imagination in Action

# Why Augmentation Beats Automation

Bridging the gap between intern-level agents and expert AI teammates

## Founder takeaways

- **Start simple, then scale:**  
Start agents on simple, low-stakes tasks to build trust, then scale responsibility gradually.
- **Get granular:**  
Domain-specific knowledge beats generic AI; customize agents for your specific environment and workflows.
- **Dynamic cybersecurity:**  
Security approaches must evolve for dynamic AI systems that learn and change in real-time.
- **Augmentation, not just automation:**  
Future competitive advantage comes from human-agent collaboration models, not automation alone.



Navin Chaddha, Rob Bearden, and Manish Chandra

Companies are incorporating AI into almost every work process imaginable: assembly lines, space robots, hospital systems and even mental health support. But rather than eliminating humans or their jobs, the most strategic AI companies want to work alongside them.

AI innovators need to balance the role of AI agents and humans in our work. For many, the goal is to design systems, tools and environments that will enable AI agents to complement what we do, rather than replace us. Rather than *automation*, these companies are focused on *augmentation*.

“One of the biggest opportunities — beyond the hardware — is the rise of AI teammates,” said **Navin Chaddha**, managing partner at Mayfield, at the recent Imagination in Action summit.

AI pioneers are excited about opportunities to connect more of the physical aspects of our world with these digital companions. The human-AI agent relationship could be further developed through extended and augmented reality interfaces, wearables and other embodied robotics. If done right, this points to a future where AI agents and robots are more adaptive, context-aware and seamlessly integrated into daily life and work.

### Understanding AI agents

AI agents and agentic systems may seem futuristic, but are already becoming widespread in the real world.

An autonomous vehicle, for instance, is an agentic system that integrates perception, planning and a deep understanding of

the world to operate. (See: “[Developing the World’s Most Trusted Driver](#),” page 50.)

Not all agents are created equal. Experts break down the types of AI teammates or digital companions based on their roles and expertise within a company. For instance, some can be considered comparable to interns: They need to be given clear instructions and a development path with small tasks. As agents gain experience and reliability, they may eventually be delegated more work and granted greater autonomy, like a junior employee.

Other expert agents might fill specialized roles, such as managing literature searches, generating hypotheses, and planning experiments in a virtual lab setting. In this context, there may be a designated AI “professor” agent and a group of “student” agents in the research team. That’s what **James Zou**, associate professor of biomedical data science at Stanford University, is currently building with his virtual lab replica, so that AI agents can work with each other in different domains.

Although “discovery rate is accelerated,” AI isn’t yet trained to always ask the right questions. It’s up to humans to retrain the models and change their reward objectives, Zou said.

# AI agents

refer to the combination of AI models and code that enables the models to use tools and take actions, such as updating a codebase or initiating a workflow.

# Agentic systems

are systems of AI agents that can interact with other tools and data sources to impact the real world, according to **Dan Boneh**, professor of applied cryptography group and co-director at Stanford University.

# Agents & agentic systems

differ from traditional automation because of their ability to act autonomously and adaptively.

## The key to successful enterprise agents

ChatGPT may be great at crafting talking points for your speech, but can it excel in the workplace to drive productivity and profit? It depends.

Many startup founders are embracing the idea of an “agentic-first” startup, where AI agents significantly outnumber human employees, handling many tasks previously reserved for humans, such as product design, coding, sales, fulfillment, and customer support, leading to exceptionally high revenue per employee.

Yet a recent [finding from MIT](#)’s Networked Agents and Decentralized AI initiative showed a 95% failure rate for enterprise AI solutions. Only about 5% of AI pilot programs saw rapid revenue acceleration, while the majority remain stuck with no impact on revenue.

The issue seemed to be a “learning gap” in most organizations, because they weren’t adapting their workflows to incorporate AI properly.

**Lisa Dolan**, managing director at Link Ventures, believes adoption in the enterprise has been hard because of the different verticals across a company. Dolan points to needing more training for agents in order to overcome the learning gap, building trust over time and then rolling them out into different verticals as they get better at low-level tasks.

“What I’m really interested in is the end-to-end workflow, adoption and training, similar to how you would train a junior employee,” Dolan said. She believes this iterative process could help bridge the “missing link” in current enterprise adoption, particularly if companies are hesitant to give agents autonomy when they haven’t proven themselves on simpler problems.

In another conversation with **Gamiel Gran**, chief commercial officer at Mayfield, CIOs of several organizations shared their perspectives on how AI will change enterprise technologies.

For example, **Naveen Zutshi**, CIO of Databricks, explained how the company has focused on building an AI agent’s knowledge base for handling customer support. The AI needs to be able to answer specific questions from that environment, as opposed to just generic questions. So Databricks’ goal for the software was building case routing, case forwarding, and case forecasting that utilized deep customer and product knowledge specific to the company.

### Scaling with humans at the center

If the idea is for AI to augment human capabilities, we have to view the process like a relationship. It moves beyond a simple dynamic of a user and their tools to a more integrated and collaborative partnership, which spans the entire lifecycle of AI, from initial training and design to ongoing collaboration, supervision and evaluation.

In all cases, experts agree that humans need to be at the center of these changes, including providing human input for training and refining AI models, as well as ensuring that their outputs are aligned with human intent and preferences. Companies like Anthropic have begun to train their models on user feedback to better understand how to support different customers and use cases.

In discussion with Chaddha, **Rob Bearden**, founder of Cloudera and current CEO of enterprise AI company Sema4.ai, suggested that many businesses are trying to understand the role of agentic technology and determine its use cases. The key is to understand where it fits in the stack, which components to use and then how to unlock it at scale — then repeating the application of AI across supply chain, product and customer engagement.

Bearden noted that agentic AI may convey lasting advantages to its early adopters, comparable to the way other transformative enterprise technologies have worked. “This is a force multiplier opportunity,” Bearden said. “If you do it right, there’s a huge benefit.”

**Brian Elliot**, CEO of AI development platform Blitz AI, emphasized the importance of establishing trust and regulatory compliance with autonomous decision-making agents. He advised forming a responsible AI framework and technical documentation where we can clearly trace an agent’s actions.

In short, humans-in-the-loop—keeping people involved in key decisions—is critical to mitigating risks in enterprise data and IP.

Companies also need to ensure that non-technical employees are fully engaged with the technology. The companies that succeed with agentic AI will be those that enable subject matter experts, who are often non-technical, to bring their contextual knowledge into the agent-building process.

“AI agents are a force multiplier. If you do it right, there’s a huge benefit.”

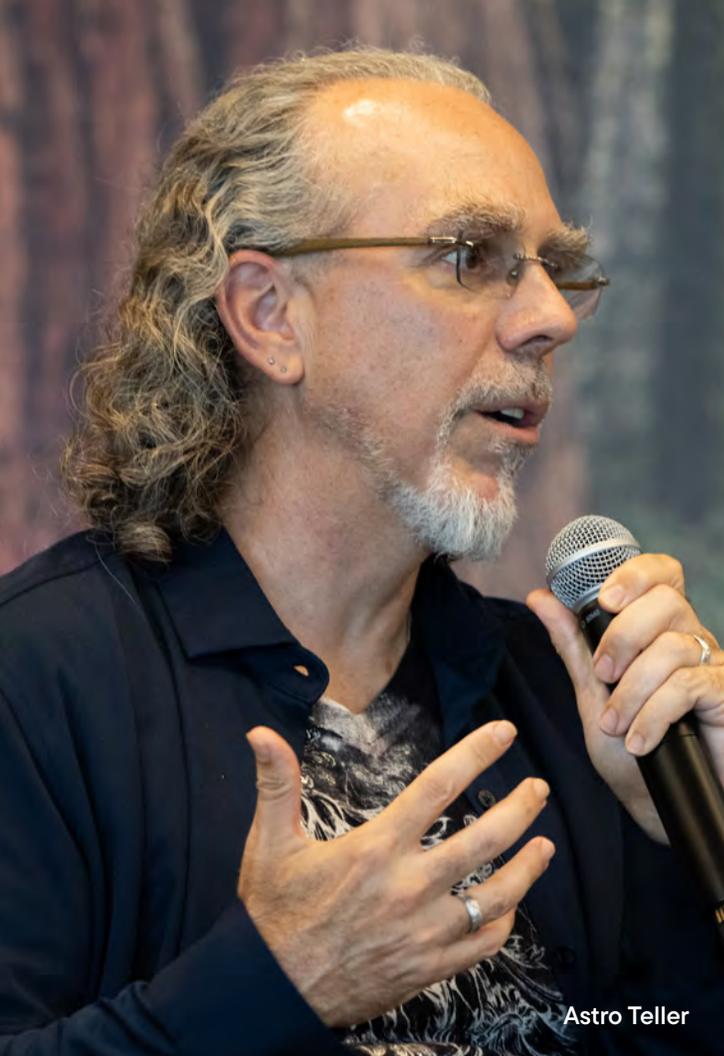
—Rob Bearden, Sema4.ai

From education to scientific research, AI agents are already being deployed as experts to leverage specialized data, models, and contextual knowledge to perform tasks that traditionally require deep human expertise. These early deployments show great promise.

Take DeepMind, the AI research and product team within Alphabet, which is using a system of agents as research assistants in discovering new targets for liver fibrosis. Another example is the AI foundation Open Athena’s software engineers and AI experts, who are working with academic labs to build scientific models for weather and plant genomics.

Madhavi Sewak





Astro Teller

“Jobs change. The people who invest in and participate in that change learn to harness and surf that change, and they do great.”

—Astro Teller, Alphabet X

Some agentic systems raise significant security concerns. An agent could be tricked into taking harmful actions or exposing sensitive data. Depending on the agent’s access privileges, that could be quite harmful—for example, some AI agents can read and send corporate emails autonomously.

**Moinul Khan**, CEO of AI security company Aurascape, warns of the major blind spot in how enterprises are handling the hundreds of new AI tools, including agents, being used by employees. While organizations have firewalls to keep bad actors out and web proxies to keep data safe, these traditional security stacks are often completely inadequate for the new challenges posed by AI.

Security leaders also have to understand how hackers have evolved with AI, and that many are already starting to use agents to automate the process of finding security flaws and exploiting them.

#### Job market impact — and opportunity

AI agents are poised to fundamentally reshape the labor market. Significant job disruption and displacement will take place in some areas.

“The history of automation says that jobs change. The people who invest in and participate in that change learn to harness and surf that change, and they do great,” said **Astro Teller**, co-founder and CEO of X at Alphabet. “People who resist that change tend to do less great.”

These approaches help academia and enterprise develop more AI training to advance their specific fields.

The big picture? AI benefits society when builders understand the people they’re building for.

#### Building security and responsibility

As many cybersecurity experts mentioned, a whole new level of security needs to be in place when we employ AI.

“What we’re moving towards now is where the security engineer’s job will be to train the machine,” said **Heather Adkins**, VP of security engineering at Google. “The security person becomes a guide, an orchestrator making decisions for the business with the business context.”

The most repetitive, routine tasks are likely to be replaced by agents in the near future. This will impact the ratio of humans to AI teammates at some companies, and some roles may see a reduction.

However, other roles could actually see growing employment, according to a [study by Stanford and ADP](#). The researchers found a 13% employment decline in the most AI-exposed occupations, but also found that “entry-level employment has declined in applications of AI that automate work, but not those that most augment it.”

This means many employees will need reskilling. AI literacy will become embedded into many jobs. And the nature of many jobs will change completely, as new tasks involving agents are created, opening new opportunities for employment that we haven’t yet imagined.

Another shift is that many employees will shift from being individual contributors to becoming managers of different agents. Consider a security engineer, who may go from tracking hackers and deploying patches to training and overseeing a fleet of security agents.

Looking ahead, **Madhavi Sewak**, distinguished research engineer at DeepMind, predicts the multi-agent model may be temporary: “Think about a year out, you’re probably just going to have the one agent that does all of these tasks for you.” The goal is making agents available so users can have easier lives without dealing with technical complexity.

The potential is enormous, and the possibilities are only starting to come to light.

As Mayfield’s Chaddha said, “If you look at these digital teammates, which will collaborate with humans to take us to superhuman levels, in five to seven years, this [could be] a \$3 to \$6 trillion dollar market.”

“Entry-level employment has declined in applications of AI that automate work, but not those that most augment it.”

— Stanford study

\$3-6  
trillion

Potential size of the digital  
teammate market in 5-7 years

# The New Startups: Smaller, Faster, Richer

Speed and small teams rewrite  
the startup playbook

## Founder takeaways

- **The lean unicorn:**  
Billion-dollar startups with less than 20 employees could become the norm.
- **Democratizing startup founding:**  
Non-technical domain experts will find it easier than ever to start and scale businesses.
- **100x multiplier:**  
AI can unlock market value an order of magnitude faster than previous tech revolutions like the advent of the Internet.
- **Scientific breakthroughs:**  
AI can speed the opening of new markets by powering discovery in scientific fields like drug research and materials design.



In the world of technology startups, the brass ring every founder yearns for is “unicorn” status—reaching a billion-dollar valuation. That’s long been the benchmark of success, even as founders recognized that achieving such status could take many years and many employees.

According to a [recent comment from Manlio Carrelli](#), the CEO of CB Insights, non-AI startups typically have a median of about 400 employees and take 9 years before they become unicorns. But startups in the AI field are hitting billion-dollar valuations with just over 200 employees—and in just 2 years.

That data point demonstrates that the AI economy is already driving impressive progress, but several speakers at Imagination in Action predicted an even more startling future. It won’t be long, they said, before advanced AI capabilities offer companies the ability to rapidly develop and improve products and services with streamlined internal operations and, by extension, far fewer people.

While one of the biggest societal fears about the impact of AI is that people losing their jobs, these speakers envisioned a different dynamic within the startup ecosystem: the rise of billion-dollar companies with an order of magnitude fewer workers.

“You could have a billion-dollar company with just a few people, because they’re able to scale that much faster.”

—Tarun Raison, Gruve

“You could have a billion-dollar company with just a few people,” said Gruve CEO **Tarun Raison**, “because they’re able to scale that much faster.”

**Jeremiah Owyang**, General Partner at Blitzscaling Ventures, said AI-native companies are redefining what “scale” looks like. AI startups are being structured to stay small. The average AI startup has 20 employees making over \$5 million in revenue, with revenue per employee of \$1-3 million versus the industry average of \$200K.

“It’s because they’re agentic first,” Owyang said. “A 19-year-old founder doesn’t need to manage 200 people. With 170 AI agents, they can still achieve the same operational leverage.”

One of the biggest factors driving this is the underlying power of next-gen AI systems to automate internal business



Amjad Masad

functions—IT, HR, sales, and the like—that used to require bulky teams. As Replit CEO **Amjad Masad** explained, non-technical domain experts can use AI to quickly build those in-house functions without hiring people, creating AI-driven businesses.

“If an [AI] agent can go hire a human or can go hire another agent,” Masad said, “then we’re going to see companies shrink. We’re going to see a lot more entrepreneurship, and the market is going to be more dynamic. I think more people will get rich.”

This democratization of startup building represents a fundamental shift. Previously, launching a tech company required either technical expertise or the capital to hire engineering teams. AI is eliminating that barrier, allowing domain experts to focus on solving problems in their fields while AI handles the technical implementation.

“We’re going to see companies shrink... and a lot more entrepreneurship.”

—Amjad Masad, Replit



Patrick Salyer

#### A 100x multiplier

The Internet spurred massive growth across the business world. By connecting everyone online, making information instantly accessible, and seamlessly bringing together companies with their customers, business got faster and more efficient—and the number of unicorns blossomed.

AI is changing the game yet again, especially for startups working hard to scale their operations.

“Back when I was a founder, six or seven years ago, sales development representatives (SDRs) were everything—we had a team of 100,” said **Patrick Salyer**, a partner at Mayfield. “That would be an absolute waste of money now. The rules are being rewritten, and founders need to adapt their go-to-market playbooks just as fast as their products.”



“The world’s moving much faster now,” Salyer said. That’s because AI is helping startups construct fluid systems in nearly every aspect, from how they build their products, how they sell them, how their supply chains collaborate, and how they engage with customers.

For startup founders looking to get their companies off the ground in a fiercely competitive landscape, AI is essential, helping them to establish market share and scale up faster than ever.

### A revolution in science

The technology is also a huge boon for those whose companies are built to productize and monetize science. That’s because AI is increasingly being used to advance scientific research, sometimes even resurfacing discarded discoveries for altogether new purposes.

UC Berkeley assistant professor of chemistry and biomedical engineering **Aditi Krishnapriyan** pointed out that through the transfer of data from large to small models, AI can vastly speed up advances in science such as the materials design of batteries.

Progress like this is fertile ground for startup founders aiming to become unicorns. Scientific discoveries that once took years or decades can now be accelerated dramatically, opening new markets and opportunities at unprecedented speeds. (See [“The Future of Research”](#), page 20.)

### Embracing AI

For founders, the message is clear: Building successful businesses in the future won’t require armies of employees. Instead, founders can scale with smaller teams by leveraging AI in any number of domains.

This may force society to rethink how to ensure economic opportunity, but for founders with new ideas, there has never been a better time to start a company.

The days are gone when founders would have to raise millions of dollars to go to market. Instead, it’ll be a race to see who can leverage AI and get the best ideas in front of customers first.

# 2 years

The time it’s taking AI startups to hit billion-dollar valuations—versus 9 years for non-AI startups

# Velocity Is Everything: How AI Is Rewriting the Rules

Why learning velocity matters  
more than first-mover advantage

## Founder takeaways

- **Move fast and ship:**  
Velocity of execution is more important than ever, because ideas have a short shelf life.
- **The game has changed:**  
Plummeting costs of software and models that democratize intelligence mean that execution and rapid iteration are essential.
- **Learning velocity beats first-mover advantage:**  
From zero to one, how quickly founding teams learn is the most important signal.
- **Talent is an asset:** Having the right team to execute swiftly on ideas in place is driving multi-billion-dollar acqui-hires.
- **Domain-specific wins:**  
Horizontal, broad-market models are played out. Vertical integration is the path forward.



Randall Lane, Ulrike Hoffman-Burchardi, Rudina Seseri, Dave Blundin and Navin Chaddha

When NVIDIA announced that it's putting [\\$100 billion into OpenAI](#)—which at a valuation of \$500 billion, is now the most valuable startup in history—it signaled a new stage in the race to expand the AI ecosystem. That race is white-hot now. The [global AI market](#) had an estimated value of \$233 billion in 2024, and is expected to grow to \$1.77 trillion by 2032.

AI's unprecedented pace of change is reshaping how VCs evaluate opportunities. For investors, one of the most critical investment criteria will be velocity—how fast can companies move? The dynamics that have powered previous tech innovation eras won't be the same this time around.

#### How fast can you ship?

For founders seeking funding, one question cuts to the heart of today's AI market: "Can you ship quickly?"

**"The half life of a good idea has shrunk from a few months to a few weeks."**

—James Detweiler, Felicis Ventures

Those are the words of **James Detweiler**, a partner at Felicis Ventures. The reason is stark. "The half life of a good idea," he said, "has shrunk from a few months to a few weeks."

That compression is fundamentally altering the investment landscape. As software costs plummet and AI models democratize intelligence, the value of ideas themselves is evaporating.

Companies that can't manifest their visions fast enough will watch their competitive advantages vanish in real-time. Being first to market, once the holy grail of startups, no longer guarantees success.

What matters instead? **Patrick Salyer**, Mayfield partner, puts it simply: "From zero to one, the most important signal is how quickly the founding team is learning."

#### Totality over first-mover advantage

To Mayfield managing partner **Navin Chaddha**, the demand for velocity is eroding traditional advantages like being first to market.

Rather, companies seeking investment must demonstrate the totality of their plan to develop ideas and fully flesh them out. New AI models or other fundamental advances in the technology may well be more impactful than being first to market.

“It’s much more about having the full, end-to-end workflow map and being very purpose-built,” he said. “It’s the totality of the pieces and on the go-to-market side. We really need to be delivering concrete ROI very fast.”

His advice for founders? Move away from horizontal, broad-market models. Those opportunities, which were once lucrative, are now “played out.”

Instead, pursue domain-specific and vertical models specialized for specific industries or functions, such as coding, security, IT, or chip design, to name a few. These targeted solutions can deliver the rapid ROI that velocity-obsessed investors demand. Chaddha predicted that one of the biggest economic opportunities will be AI teammates.



Ulrike Hoffman-Burchardi

“From zero to one, the most important signal is how quickly the founding team is learning.”

—Patrick Salyer, Mayfield

These agentic digital companions will soon be a \$3 trillion market, representing what Chaddha calls a “once in a century opportunity.”

#### The talent wars

In this velocity-driven market, VCs are also on the hunt for companies stockpiling a different kind of asset: proprietary expertise. Indeed, talent has become so valuable that we’re seeing numerous multi-billion-dollar acqui-hires.

This dynamic is creating unexpected opportunities. For the first time in perhaps two decades, legacy enterprise companies are vulnerable to disruption on a much faster timeline.

Link Ventures managing partner **Dave Blundin** explained he thinks startups shouldn’t be afraid of taking on incumbents. Those companies’ businesses are outdated, and some actually want to support thriving startups around them that can bolster their own efforts.

Either way, he advised startup founders that they should pursue “blue ocean opportunities where the incumbents are weak,” particularly in agentic applications, which may soon dominate monolithic workflows.

#### The self-improving revolution

What makes investors confident this isn’t another bubble? As UBS Global Wealth CIO for global equities **Ulrike Hoffman-Burchardi** put it, “AI is the first technology in human history that is self-improving. We’ve never had this before, so progress is much more rapid.”

That exponential improvement curve is already producing results that seemed impossible just years ago.

Stories of startups generating hundreds of millions of dollars in revenue in a few short years are no longer anomalies—they're the new normal. That's led many investors to believe we're just scratching the surface of what's possible.

For Chaddha, the implications are clear. AI represents the “most transformative thing in human history.”

### The new X factor

Looking ahead, the traditional costs and timelines of building tech companies are giving way to a new model—one that prioritizes speed, talent optimization, and domain expertise above all else.

The next generation of successful AI companies will look nothing like legacy SaaS businesses. Instead, expect vertically-integrated platforms delivering streamlined, end-to-end workflows.

In this landscape, “move fast and ship” isn't just advice for founders. Investors are recalibrating their entire approach to keep up with the pace, integrating AI into their investment workflows or, as Detweiler even suggested, have AI make venture investments decisions.

The next round of successful AI efforts will be much less like legacy SaaS companies and more like vertically-integrated platforms providing streamlined end-to-end workflows.

Time, not capital, is now the X factor. The AI market may be racing towards \$1.77 trillion, but only companies and investors that have internalized the velocity imperative will capture that value.



For founders who can master the art of rapid learning and execution, the returns could dwarf anything we've seen in previous technology cycles. But the window to move is measured in weeks, not months.

# \$1.77 trillion

Expected size of the  
AI market in 2032

# The University Lab Is the New Launchpad

How AI is boosting the speed and scope of scientific research

## Founder takeaways

- **The college-to-startup pipeline:**  
Lab-to-market is emerging as a major venture creation pathway alongside traditional startup models.
- **Real-world readiness:**  
Academic research must be translated into production-ready systems with long-term support to succeed commercially.
- **Turbocharging research:**  
Natural language feedback for model optimization can dramatically accelerate scientific discovery.
- **Openness pays dividends:**  
Open-source data and cross-institutional collaboration drive breakthrough discoveries and keep the U.S. competitive globally.



Navin Chaddha, Matei Zaharia and Ion Stoica

The best startup ideas are no longer being hatched in dorm rooms and garages. They're emerging from university research labs, where AI is turning academic breakthroughs into commercial goldmines.

Mayfield partner **Vijay Reddy** noted that AI for research and innovation is fast becoming a key investment area. That's particularly true for domain-specific models meant for use in research-intensive fields like life sciences and physics. He believes founders with unique problem insights and passion—rather than just technical skills—will build the most valuable companies.

The reason is simple: AI is fundamentally reshaping how science gets done. It acts as a powerful discovery engine that augments human intellect, automates tedious research processes, and uncovers insights from vast datasets that were previously inaccessible.

Within a few years, AI will manage day-to-day research processes, automating everything from literature searches and data analysis to hypothesis generation and experiment

planning, according to **Sam Rodriques**, CEO of FutureHouse, a nonprofit dedicated to building AI scientists.

This isn't just incremental improvement. Rodriques believes AI will spark a “revolution in the way that we do science.”

#### Uncovering what's hiding in plain sight

AI's ability to mine the troves of existing datasets is already helping us uncover discoveries and overcome barriers in science.

Academic research creates innovative, disciplined thinkers who push boundaries without fear—qualities valuable in corporate settings too. To maximize group potential, you must “make them curious” and help them “feel that they can push on the boundaries,” said **David Siegel**, co-founder and co-chairman of Two Sigma.

**Yossi Matias**, VP and head of research at Google, sees AI as a co-pilot in the lab. He framed AI within a “magic cycle of research,” where machine intelligence accelerates every phase of scientific discovery: identifying a problem, conducting experiments, and applying solutions at unprecedented speed.

“AI will spark a revolution in the way that we do science.”

—Sam Rodriques, FutureHouse



Peter Danenberg

For instance, Google's DeepMind "AI co-scientist" doesn't just crunch data, it reads papers, generates hypotheses, validates them, and ranks the most promising leads. It's already helped uncover new therapeutic targets for liver fibrosis, hinting at a future where AI isn't just supporting science, it's driving it.

In neuroscience, AI is mapping the human brain through connectomics, while in medicine, models like Med-PaLM, a large language model from Google Research, and even Google Gemini are passing medical licensing exams at expert levels.

In biotech, companies like Retro Biosciences are using AI to reverse aging at the cellular level, predicting protein folding and identifying therapeutic targets with impressive precision.

In other areas, AI virtual labs are designing protein binders for COVID variants faster than human teams.

### Opportunities for industry growth

But any startup in science needs to turn their ideas into go-to-market strategies by bridging the gap between academic research and commercial needs. Founders also have to balance sharing research openly versus holding it back to spin out a company.

While academic research projects often gain traction on their own, enterprises need a committed commercial entity for long-term support, bug fixes and production-readiness. This was the catalyst for the founders of Databricks, **Ion Stoica** and **Matei Zaharia**, who transitioned from academic research at U.C. Berkeley to start their own company. They wanted to ensure the continued development and support of the technology they'd created.

Stoica and Zaharia advise students and entrepreneurs to pick a big problem they are passionate about solving, then build rigorous evaluation systems to validate their progress.

"A major limitation in enterprise AI adoption is how poorly teams set up validation for their applications," Zaharia explained. For AI to be effective in science, companies need high-quality data, rigorous evaluation methods, and the engineering infrastructure to support large-scale models.

One area of research Zaharia is excited about is optimizing models using natural language feedback, not just numerical rewards. This method is much faster and could accelerate discovery in scientific fields where physical experiments are expensive.

### The new AI research landscape

Meanwhile, many companies are building large foundational models aimed at scientific research, including Google, OpenAI, Anthropic and Axiom. Each is taking a distinct approach.

Anthropic does research and develops safety protocols and enterprise applications, growing a business where customers pay for token usage. **Andi Peng**, researcher at Anthropic, mentioned the company's efforts to scale up its AI models and train their personalities and behavior to be more helpful and collaborative for people.

The company's philosophy is that building intelligence requires studying how humans use intelligence, which then informs how the models are developed, Peng explained.

Axiom is using formal mathematical proofs to train AI, aiming to build a verified "super intelligence." Founder and CEO **Carina Hong** said math is a "perfect sandbox" for developing these problem-solving approaches that can later be applied to other domains. Through self-improving systems, models can reflect and learn from their mistakes. Then multiple models can interact with each other to continuously improve each others' performance. This helps human researchers discover new questions and hypotheses worth pursuing.

By translating every problem, solution, and proof into a computer-readable format—much like computer code—Hong believes we can usher in an era of "mass intelligence" where AI can be applied at high speed to various applied sciences. She offered the example of AI potentially solving extremely difficult partial differential equations, which could lead to major breakthroughs in areas like nuclear reactor design.

### Open collaboration as competitive advantage

As research applications multiply, an interesting dynamic is emerging around openness versus consolidation.

**Tengyu Ma**, chief AI scientist at Mongo DB, assistant professor in computer science and statistics at Stanford University, and cofounder of Voyage AI, [which was acquired by Mongo DB](#), sees a trend toward more consolidation among AI models. Currently, only a handful of major models like OpenAI and Anthropic dominate the space.

However, openness may become the catalyst for breakthroughs in AI research. For humanity as a whole to stay at the cutting edge, experts argue that research should be done openly to foster more partnerships and support between institutions across academia and industry. Open-sourcing data is seen as a key solution to promoting transparency and allowing for the replication of results, a fundamental to the scientific method.

**Mike Abbott**, cofounder of Open Athena, which partners with academic labs to build large AI models, argues that conducting research "out in the open" can ensure that the U.S. stays ahead of the curve, as well.

**"We need to be working together across institutions and promote interdepartmental collaboration—that's where those big breakthroughs come."**

—Mike Abbott, Open Athena

"We need to be working together across these different institutions ... and promote interdepartmental collaboration, because it's so cross functional—and that's where those big breakthroughs come," Abbott said.

For founders emerging from university labs, the message is clear: the opportunity is massive, but success requires more than technical brilliance. It demands passion for solving real problems, strategic thinking about commercialization, a commitment to rigorous evaluation, and embracing open collaboration.



Tengyu Ma

# AI Infrastructure Is Ripe for Disruption

Solving the bottlenecks in computing, cooling, and chips

## Founder takeaways

- **Need for speed:**  
The AI industry needs a 1,000x performance improvement in the next three years to sustain current innovation rates.
- **Pain points:**  
The greatest opportunities lie in solving bottlenecks: expensive GPUs, energy-hungry data centers, networking speed, and memory capacity.
- **Thinking many moves ahead:**  
Learn from NVIDIA's playbook: they disrupted established players by envisioning what GPUs could become.

Amin Vahdat



Because of generative AI's rapid advances over just three years, it's easy to forget that the technology is built on a decades-long effort to massively increase computational power—and that AI models in development will need even more power very soon.

Large language models (LLMs) became possible thanks to a remarkable thousand-fold increase in computing performance over the past 20 years. That improvement was driven in part by Moore's Law, but also by the introduction of new tools that enabled the building of scalable, distributed systems based on commodity hardware, said Google Fellow **Amin Vahdat**.

However, performance growth has started to peak exactly at the moment when AI developers need more performance. "Just as the underlying trends slow, the demand for computing exploded," Vahdat said.

He predicted that the AI industry needs another 1000x performance improvement in just the next three years if its current pace of innovation is to continue.

**"Just as the underlying trends slow, the demand for computing exploded."**

—Amin Vahdat, Google

Many in the infrastructure field are excited about the challenges and opportunities in delivering those performance improvements, including new solutions for computing, networking, and power generation.

#### AI infrastructure needs new solutions

Despite massive amounts of new data center construction, there may simply not be enough computer capacity to meet expected demand, and big AI companies are snapping it up wherever they can. OpenAI recently agreed to pay Oracle \$300 billion over five years to provide compute power. McKinsey & Company estimated in a report that the amount of data center investments needed by 2030 will reach \$7 trillion, with \$5.2 trillion going to AI workloads alone.

GPUs, the chips that enable AI models to perform their massively parallel computations, are an essential part of neural networks. But GPUs are expensive, consume enormous quantities of electricity, and generate massive amounts of heat. What AI companies want are less expensive computing solutions that provide more performance while consuming less power. Other bottlenecks include networking speed and a lack of memory capacity.

The first step to creating bold new solutions is to understand the limitations of the current systems.

For example, chips used by the AI industry today were designed as long as eight years ago, and were not meant to support today's AI models, said **Faraj Aalaei**, founder and CEO of Cognichip.

"There's no way that (semiconductor companies such as NVIDIA) could have comprehended this kind of scaling this fast," Aalaei said. "The fundamental problem in our industry, the hidden problem, is that chips take too damn long to design."

### Hardware evolution ahead

**Anirudh Devgan**, CEO of Cadence Design Systems and a pioneer in electronic design and circuit simulations, is certain plenty of changes will come to semiconductors, and that they will continue to be a growth driver in AI and other sectors, including robotics.

"Chip design complexity is growing exponentially," Devgan said, adding that hardware innovations will continue to drive major changes, particularly as specialized silicon emerges for robotics and other applications.



Anirudh Devgan

**“The fundamental problem in our industry is that chips take too damn long to design.”**

—Faraj Aalaei, Cognichip

Many data centers also weren't built to support AI, says **Seshu Madhavapeddy**, co-founder and CEO of Frore Systems, a developer of thermal technologies. The people who designed most of today's existing data centers had no way to anticipate that they would need to cram so many high-performing GPUs onto their racks. Just cooling these super-powerful processors is an engineering challenge of its own.

"The reality is that all the power converts to heat, and if you don't have a very efficient means of removing heat, then you're not going to be able to actually run your data centers at a high level of performance," Madhavapeddy said.

### Golden age for infrastructure startups

"This is a great time for startups," said Frore's Madhavapeddy. "Because the data center market is growing so fast and generating so quickly, and everything is changing every year or every two years, this type of dynamic environment is where startups can really thrive."

When it comes to disrupting infrastructure, some opportunities are better than others. Aalaei recommends that startup founders take big swings and create discontinuity. Making incremental changes to existing platforms won't deliver the big rewards that VCs seek.

Still, he doesn't recommend competing directly with NVIDIA for the GPU market. They're too entrenched and too rich. Instead, he recommends that founders sizing up a market study the lessons that NVIDIA's success story teaches.



Faraj Aalaei

“When NVIDIA came up as a startup, there were already people building graphic chips,” Aalaei said. “Intel was still the king of the jungle. And NVIDIA found their way of creating discontinuity... this did not happen by accident. They had a thought process about what a GPU could do and what the future held. And they spent an enormous amount of money as a public company doing this, and they had the courage to do it and stick to it.”

**Rajiv Khemani**, co-founder and CEO of Auradine, maker of blockchain and AI applications, agreed. But Khemani cautioned that future founders should not overlook service or technology gaps in existing market segments.

“If you’re doing a new startup, you have to be very clear on what you’re going to do,” he said. “There’s a phenomenal opportunity to segment. You have to pick a segment where you don’t have a dominant player and where you can deliver a fine X value.”

The infrastructure crisis threatening AI’s future is at the same time the industry’s greatest opportunity. For founders who can solve these fundamental bottlenecks, the potential is enormous.

1000x

Increase in computing power over the past 20 years — and what will be needed in just 3 years to maintain AI’s growth rate

\$7 trillion

Data center investments needed by 2030

# Why Energy Is AI's Next Frontier

Innovation at every layer of the stack is racing to solve AI's biggest bottleneck

## Founder takeaways

- **The limits of power:**  
AI's computational demands are doubling every 100 days, creating an energy crisis that could constrain the industry's growth.
- **Opportunity in efficiency:**  
The industry is entering a "golden era of silicon hardware." Chip efficiency, intelligent cooling, and renewable generation are billion-dollar opportunities.
- **Own the stack:**  
Vertical integration from algorithms to silicon to power generation defines the new competitive advantage.



FUTURE OF ENERGY

Somewhere amid this AI explosion, there is a rather glaring problem the industry needs to solve: How to get enough power in order to actually leverage AI to its full potential.

That's the billion-dollar question for AI stakeholders right now.

AI data centers' rapid growth is creating a crisis in the American energy system. But at the same time, the availability of electric power is now one of the chief obstacles to AI growth. Addressing it will be crucial to the continued stability of both the grid and the AI industry.

Many possible solutions are on the table, including more flexible data centers, more efficient chips, using AI for grid management and alternative energy sources such as nuclear power. AI companies including Oracle, Microsoft, and Amazon have even considered powering data centers with their own small modular reactors

If companies want to scale their computation and win the AI race, they will have to figure out how to build the right energy infrastructure.

### Using AI to optimize AI

Unsurprisingly, AI itself is being used as a tool to manage energy costs. It's a bit like using "Google Maps for electrons" to utilize the grid more intelligently. In 2019, Google's DeepMind used AI to more accurately predict wind power output 36 hours in advance, [increasing the value of its wind power by over 20%](#).

**Tamanna Sait**, VP of cloud engineering at Crusoe, aims to build more flexible data centers that can minimize environmental impact and avoid consuming power. Crusoe is currently building a data center in Abilene, Texas, that uses diverse renewable energy sources, including wind, solar, geothermal and hydroelectric power.

By using AI to manage them, the data center can forecast the supply and demand, predicting its own energy needs hours in advance and balancing different power sources to meet the need.

"We are using the abundant West Texas wind to generate power, and natural gas turbines to generate the backup power," Sait said.

To understand the stakes, consider this: [AI servers in the U.S. gobbled](#) up more than 40 terawatt-hours in 2023, a huge increase from 2017. As AI scales, its power consumption will escalate, with one estimate suggesting data centers could consume up to [945 terawatt-hours globally by 2030](#). That amount is comparable to Japan's entire annual electricity output.

“We are using the abundant West Texas wind to generate power, and natural gas turbines to generate the backup power.”

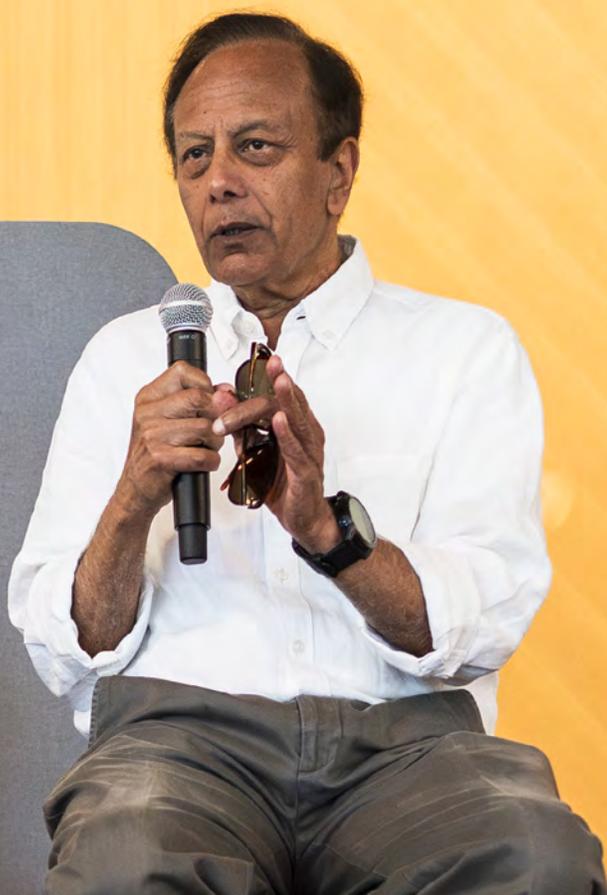
—Tamanna Sait, Crusoe

#### Revamping performance, chips, and cooling

High-performance chips generate tremendous heat, requiring extensive cooling infrastructure that adds another layer of energy consumption.

Companies are working with enterprises to use AI to optimize those cooling systems and other aspects of a data center’s operation.

Pradeep Sindhu



For instance, Bay Compute’s AI can help find where a data center is wasting energy, helping the operator to cut costs by staggering heavy usage, said **Vijay N. Gadepally**, cofounder of Bay Compute.

He said customers already see results in performance improving, comparable to the way households can save money and avoid problems by staggering their appliance usage.

“It’d be like having a single 100 amp breaker in your house, and you’re charging your car, plus turning on the dryer, plus your refrigerator, plus running the air conditioner—all at the exact same time. Of course, you’re going to run into a problem,” Gadepally said.

“But if you can stagger them and move usage around, you can do a lot more with that single breaker.”

Efficiency matters more and more. “Chip manufacturers will increasingly be measured on performance per watt rather than performance per dollar”, said **Amin Vahdat**, VP of AI and infrastructure at Google.

Google is also investing in methods like liquid cooling and analyzing performance metrics as energy initiatives grow.

Other companies are working on specialized chips that are more power-efficient than the GPUs currently used for AI workloads.

From a chip perspective, Auradine CEO **Rajiv Khemani** is taking what his company initially built for blockchain and Bitcoin to make AI energy use more efficient. The company’s technology for achieving low-power computation takes advantage of the V-squared effect, where power consumption is a function of voltage squared. Auradine runs chips at “near threshold voltage,” or below 0.3 volts, compared to typical AI chips that run at 0.75 volts, per Khemani. That greatly reduces their power consumption.

#### Costs and infrastructure continue to evolve

**Pradeep Sindhu**, founder, chairman and CTO of Juniper Networks, referred to a “golden era of silicon hardware” that will drive more vertical integration as AI grows.



FUTURE OF ENERGY

The winners will be companies where teams understand algorithms, software layers, silicon design, and networking—then integrate all of it into cohesive systems. That kind of full-stack thinking is rare but increasingly essential, Sindhu said.

The business models undergirding the AI ecosystem will transform as energy economics shift. Data centers, cloud providers, GPU manufacturers: All will need to adapt their pricing, service offerings, and strategic positioning.

Expect relentless innovation across the entire infrastructure stack as hardware and software teams work to unlock new efficiencies. The companies that crack the code on sustainable, scalable power will capture outsized value.

Balancing power supply and demand is essential for AI to scale into robotics, life sciences and enterprise applications. Without sustainable energy solutions, the promise of AI could be limited by its own hunger for power.

“Chip manufacturers will increasingly be measured on performance per watt.”

—Amin Vahdat, Google

945 terawatt-hours

Estimated global data center consumption in 2030 (equal to the annual power output of Japan today)

# Consumer AI Is Outpacing Enterprise

From chatbots to concierges:  
Personalization at scale

## Founder takeaways

- **Massive market:**  
Consumer AI adoption is outpacing enterprise, with 500M people using AI tools daily. Low friction drives exponential growth.
- **Start focused and grow from there:**  
The consumer AI market is fragmented (130+ brands) but massive (1.7 billion global users). Find your niche and wedge.
- **Your AI concierge:**  
The next wave isn't better AI: It's context-aware AI agents that understand mood, health, and personal context. Hyper-personalization is the moat.
- **Building trust is essential:**  
The winners in consumer AI will be companies that innovate boldly while building strong privacy protections and earning user trust from day one.



Rob Bearden and Manish Chandra

In 2017, Google scientists published “[Attention is All You Need](#),” which introduced the transformer architecture. That seminal moment in AI research was the jumping-off point for the wild success of companies like OpenAI, Anthropic, and the like, and helped the AI industry explode from an estimated \$16 billion in revenue [in 2017](#) to as much as \$757 billion [this year](#).

Generative AI tools like ChatGPT offer widespread benefits for the daily lives of everyday consumers. Indeed, consumer AI benefits are currently outpacing those enjoyed within the enterprise.

A big explanation for that is how easy it is for ordinary consumers to use ChatGPT (and its rival generative AI tools). Everyday users can’t get enough of doing research, semantic search queries, financial analysis, or just entertaining themselves and their friends with AI-generated images.

One stat sums this dynamic up best: In 2024 alone, generative AI tools [accounted for \\$97 billion](#) in “consumer surplus,” the delta between what consumers will pay for something and what it actually costs.

### Opening up new experiences

But it’s not just OpenAI and rivals who are driving massive consumer AI adoption. Estimates show that, globally, as many as 1.7 billion people use AI tools, with 500 million using them daily.

And while OpenAI generates about 40% of global consumer AI spend, a recent report suggests that more than 130 brands

are providing AI tools in areas like video, business, voice and transcription, and writing.

Agents are likely one of the next big growth areas. These will offer far-reaching use cases and lead to many new consumer AI-oriented businesses that span shopping, healthcare, entertainment, and more.

For example, **Manish Chandra**, cofounder of Poshmark, envisions a personal style agent that knows who you are, what you like to wear, what appointments you have that day, and even what mood you’re in—and is able to use all that information to recommend outfits from your closet each morning.

Meta’s head of business AI, **Clara Shih**, also predicted an imminent future in which consumers employ AI assistants to help them find the specific products they want.

To do that effectively, developers of AI agents will need to figure out how to get them to respond appropriately to a variety of wildly different requests with the right “personality.”

For instance, someone asking for help on a math assignment expects a very different response from when they’re asking a model to generate some code. Yet it’s often the same model powering both responses.



Madhavi Sewak

“Personality is incredibly important in how users receive a model,” said **Madhavi Sewak**, a distinguished researcher at Google working on DeepMind. And yet it’s an unsolved problem—so far. “We don’t really yet know how exactly we can train personalities into models, but models obviously have personalities,” Sewak said.

“Personality is incredibly important in how users receive a model.”

—Madhavi Sewak, Google DeepMind

#### Opportunities in personalized medicine

Experts see potential in using AI to transform generic health data and advice into hyper-personalized care. Whether it is offering clinical support or analyzing trends for mental health professionals, AI systems can give more tailored insights—sometimes with more accuracy than professionals.

General Catalyst managing director **Quentin Clark** said he expects one solution to America’s often-dysfunctional medical system could be every consumer eventually having an AI concierge—a virtual doctor.

**Erick Tseng**, CEO of Next Chapter, is developing an AI-assisted mental health coaching service that uses a custom model trained on bespoke clinical data.

Tseng tells an anecdote in which a caregiver he knew suspected early dementia in his father. While a general practitioner dismissed the concern, an AI flagged early signs of Alzheimer’s. A specialist later confirmed the diagnosis.

Backed by Sam Altman, Retro Biosciences aims to find therapeutics that can reverse aging. The company has two programs that are particularly notable in trying to tackle the problem of aging itself by developing therapeutics for age-related diseases.

Using AI research, the company is working on treatments in autophagy, the natural recycling process in cells that breaks down old proteins. This process can become stalled in advanced age and in neurodegenerative diseases like Alzheimer’s. Another area of research is a cell-replacement therapy that swaps out old microglia (one of the four main cell types in the brain) with new, “zero age” cells created in the lab, per CEO **Joe Betts-LaCroix**.

#### Concerns about privacy and data

AI is democratizing expertise in health care, offering second opinions to those who might never afford them or who face other obstacles in getting care.

But many of these systems aren’t HIPAA-compliant, and sensitive data is often fed into general-purpose models. That’s a risk for patients and the companies alike.

The paradox becomes: The more helpful the AI, the more vulnerable the patient. This is why involving humans-in-the-loop will become more critical than ever. Specifically, that means using AI to assist and augment health professionals, rather than replace them. For instance, a therapist could enlist AI to monitor a session and surface prompts for care in the summaries. It becomes a collaborative model to build trust, where the final decisions and nuanced care still reside with a human.

Outside of healthcare, it's long been said that when we use free technologies, we are the product—and that raises additional issues about personal data.

“When Facebook and Google first started,” said **Vanessa Parli**, the director of research programs at Stanford HAI, “we didn’t even ask, ‘how are they using my data?’” But now we’re less trusting and demand more transparency and protection.

“If we can build AI and protect our privacy, there’s a big market for that.”

—Vanessa Parli, Stanford

While consumers are making it clear they’re eager for more and more things to do with AI, they’re also signaling their wariness about giving away too much of their personal data. As **Sichao Wang**, a senior manager at Cisco, put it, “we’re actually playing on...very risky ground” if there’s no clear security or privacy foundation built into today’s (and tomorrow’s) AI tools.

That idea was echoed by **Sandy Pentland**, a professor of Media Arts & Science at Stanford. There’s a growing fear, Pentland said, that although AI agents are giving consumers more power, they’re also increasing the risk of harmful trust and security breaches.

That was the motivation behind [LoyalAgents](#), a partnership between Consumer Reports and the Stanford Digital Economy Lab that aims to “make AI agents secure, loyal, and effective advocates for consumers everywhere.”

For Parli, there are lessons to be learned from industries like healthcare and finance, which have strong mandates to protect consumers’ privacy. Now, she believes, there’s a big opportunity to incorporate that mandate across consumer AI without harming corporate business interests. “If we can build AI and protect our privacy,” she said, “there’s a big market for that, and I would prefer that it’s about educating the public.”

#### Going forward boldly

It’s already clear the consumers want AI and that their enthusiasm will only generate more and more business opportunities. The companies seeking to leverage that opportunity should build boldly but responsibly.

Those that thrive in this burgeoning market will be those that both create tools people want and take their privacy and security seriously. They’ll be the ones that recognize AI’s power to hyper-personalize experiences, understanding individual mood, health, and context, while building the trust frameworks necessary to make personalization feel safe rather than invasive.

1.7 billion

Number of people worldwide using AI tools

# Enterprise AI's Real Challenge

Breaking through data silos, legacy systems,  
and corporate inertia

## Founder takeaways

- **Process and culture barriers:**  
Enterprise AI challenges stem from organizational issues—legacy systems, data silos, and resistance to change—not from AI capability.
- **So long, DIY:**  
The era of building bespoke AI models is over. Companies are now outsourcing data management and using best-in-class LLMs.
- **Transform and win:**  
Startups that help enterprises break down data silos, migrate legacy systems, and implement mission-critical AI agents represent a massive market opportunity.



Shirin Gaffary and Clara Shih

Most enterprises are struggling to integrate AI into their businesses and realize value from the technology.

According to a recent MIT study, [The State of AI in Business 2025](#), despite investments of \$30–40 billion in AI, 95% of companies have seen zero return.

This report sent shockwaves through the AI sector. What hasn't received as much attention is that [the report's authors told Fortune](#) that some large companies and startups are “really excelling with generative AI,” extracting millions of dollars in value while the rest flounder.

Companies that can close that gap with innovative AI solutions tailored to meet the needs of large and deep-pocketed enterprises stand to bank a fortune.

Enterprises need more customized and powerful AI tools, it's true. But they could also benefit from a greater willingness to embrace change, especially with respect to their workflows and processes.

That puts technical leaders in a pivotal role, said **Sharon Mandell**, CIO of Juniper Networks. “Their role is to help those who don't naturally understand the technology,” reducing friction and resistance to new technologies.

**Why AI adoption takes time in the enterprise sector**  
**Erik Brynjolfsson**, a Stanford professor and director of the university's Digital Economy Lab, agreed that there is a lot of disappointment about AI's seeming lack of value.

But, he added, that shouldn't come as a surprise at this stage.

“That's actually a very common pattern with general purpose technologies, such as the steam engine and electricity. You have these (powerful) capabilities, but getting that to turn into business value often takes much longer.”

The pattern is familiar: Transformative technology arrives, expectations soar, reality disappoints, and then—after painful adaptation—value finally materializes.

We're currently in the “painful adaptation” phase.

#### The data readiness crisis

There are other reasons for AI's weak results in the enterprise, such as a lack of data readiness. Older companies may have data in a wide variety of incompatible and hard-to-access formats. They may have enormous data lakes filled with unstructured and poorly managed data.

Or, they may have troves of valuable intellectual property accumulated over decades. Many companies, especially in the financial services and healthcare industries, need to worry about complying with regulatory requirements and privacy concerns.

“For an agent to be truly effective, it needs a lot of freedom.”

—Randall Lane, Forbes

“When organizations examined job roles or titles, most people within the organization admitted, ‘Oh, that’s not what people actually do,’” Shact said. “This (knowledge) gap makes it extremely difficult for companies to look at people’s roles, personas, and skills to determine their future needs.”

Other times, attempts to build new apps that automate or streamline workflows are stymied by the use of outmoded systems or processes. Companies may be too attached to legacy processes. Shifting to a new way of doing things may require too much change management and culture shifts.

Without understanding current processes, building AI systems to improve them becomes impossible.

### Big migrations are getting easier

Even today, many companies in regulated industries continue to operate their own data centers rather than becoming cloud-native—and that can limit AI’s effectiveness too.

In some cases that’s a calculation about minimizing perceived risk. Other times, decision makers may simply want to avoid the hassle and cost associated with migrating to new tech platforms, which often involves upgrading code.

Years ago, it was typical for banks and other regulated companies facing a large platform migration involving 20 million or 30 million lines of code to hire consultants; the projects could require millions of dollars and more than a year’s worth of labor.

But **Brian Elliott**, CEO of Blitz AI, a maker of an autonomous software-development platform, said the tools now exist to enable worry-free technology migrations.



Randall Lane and Brian Elliott

These companies have all sorts of data controls, and are typically wary about anything that might affect their data, said **Randall Lane**, chief content officer at Forbes.

“For an agent to be truly effective,” Randall said, “it needs a lot of freedom (to manage data).” That’s going to make more conservative companies reluctant to roll out agentic or other AI technologies very fast.

### The knowledge gap

Sometimes the problem isn’t bad data—it’s no data at all.

**Laura Shact**, a principal in the human capital practice at Deloitte Consulting, said the COVID-19 pandemic revealed to many executives that they had no idea what work employees were actually doing.

Blitzly is built around a large-scale platform that deploys multi-agent systems to complete development tasks such as large-scale migrations that usually take six months to a year or more.

“Our developers were able to do that in two and a half months,” Elliott said, also noting customers get the added benefit of not having to pay a consultant’s markup fee.

The technology exists. The question is whether enterprises will embrace it.

### Breaking down the walls

A growing number of C-suite executives are recognizing that data silos stand in the way of AI success. Execs at these companies welcome services and tools that can tear down the barriers separating data from the people who can use it.

In the AI era, this is what’s required to unlock value from a company’s data: Toppling silos and providing AI with access to as much data as possible.

Even when attitudes shift, plenty of potential pitfalls remain. Individual business units may adopt a range of unintegrated AI tools or isolated AI agents to suit a narrow set of needs. A lot of organizations are still without a single source of truth and data sources remain scattered. In cases where companies intend to enable AI to make decisions based on the totality of their data, they can inadvertently wall off information to “the detriment of the larger organization,” said **Philip Rathle**, CTO at Neo4j, a database company.

It helps to take an expansive view of data. **Abhishek Mehta**, CEO of Tresata, a data and predictive analytics company, says that the least successful companies he sees are those that think “the only data that is valuable to you is within your four walls.” In contrast, Tresata’s most successful clients combine their company’s data with information from other sources. He added that this strategy generates “supernatural value.”

These experts all agree that unleashing AI’s potential means combining isolated stores of data. That is what **Lisa Dolan**, managing director Link Ventures, calls the “big unlock.”

### The shift away from DIY AI

Some of the earliest enterprise AI efforts centered on creating bespoke AI models. Companies hired data scientists and engineering teams and set out to build their own LLMs and custom data integrations. Some of those companies quickly became disillusioned with the expense and complexity of this approach.

That’s all changed. Many companies have now concluded that it no longer makes sense to build their own AI and data management infrastructure. They don’t want to pay their employees to organize, store and clean data, or to create models from scratch—not when it’s easier and more effective to outsource those chores or use best-in-class LLMs from the leading AI vendors.



“Now buyers want specific AI teammates and want to buy a complete service, not just the technology.”

—Patrick Salyer, Mayfield

“DIY didn’t work well,” said **Patrick Salyer** of Mayfield. “Now buyers want specific AI teammates and want to buy a complete service, not just the technology.”

The shift has opened opportunities for AI companies upgrading the enterprise stack. Those include the agentic systems that enable goal-directed execution rather than task-oriented execution. Instead of asking AI agents to perform individual tasks, agents increasingly have the ability to figure out what tasks need to be performed to achieve a stated goal, then perform those tasks autonomously, said **Rao Surapaneni**, vice president of Google’s CloudAI Business Application Platform.

“Agents are now able to collect, collate, and synthesize lots of disparate information; summarize it; and reduce cognitive load.”

—Rao Surapaneni, Google

Another way customers are deploying agents is to perform deep research, which Surapaneni says he himself uses every day.

“With current reasoning capability,” he said, “agents are now able to collect, collate, and synthesize lots of disparate information; summarize it; and reduce cognitive load.”

Agents also are evolving into the primary interface for interacting with technology, according to **Nancy Xu**, vice president of Agentforce for Salesforce. “I think the

conversational interface that many of us are familiar with in terms of how we interact with agents is becoming a new user interface,” she said.

#### The path forward for enterprise AI

Companies targeting the enterprise market should focus on agents that perform mission-critical tasks in observability, governance, security, and performance measurement.

Long term success requires better pricing models. Charging by the number of tokens used, as many AI companies currently do, will not be a viable long-term strategy for the enterprise market.

Vendors must also help enterprises understand that their own data can create competitive moats. But to accomplish that, enterprises need to embrace new ways of leveraging data, breaking down internal silos and developing new workflows and processes.

Will enterprises be ready to take advantage of those capabilities? That depends entirely on their ability to overcome their own inertia. For the companies that can help them close that gap, the opportunity is extraordinary.



# Vertical Is the New Horizontal

Where cost savings end and new business models begin

## Founder takeaways

- **Think beyond efficiency:**  
AI augmentation doesn't just make existing operations more efficient, it opens doors to adjacent markets and new revenue streams.
- **Automating repetitive knowledge work:**  
The biggest ROI comes from industries with high labor costs and iterative workflows—healthcare, finance, legal, customer service—where AI can automate knowledge work while maintaining quality.
- **Vertical is the new horizontal:**  
Industry-specific vertical AI is becoming the new horizontal—transforming “sleepy SaaS” into one of the most exciting categories for building companies.
- **AI with personality:**  
Researchers and scientists are pursuing an even more ambitious goal: building human qualities into AI models.



Sri Pangulur, Bratin Saha, and Lisa Dolan

The rise of applied AI is fundamentally changing how companies approach innovation and competition within vertical industries. It's reshaping business operations and workflows, with companies layering AI capabilities on top of existing systems.

"This is the first time that we really have the technology that can automate knowledge work by taking in data and insights and giving recommendations," said **Bratin Saha**, chief product and technology officer at DigitalOcean.

#### Where AI delivers immediate ROI

Industries enjoying the most success with current-generation AI models share a common characteristic: They have iterative workflows where built-in error resiliency allows for rapid improvement.

**"This is the first time that we have technology that can automate knowledge work by taking in data and insights and giving recommendations."**

—**Bratin Saha**, DigitalOcean

Think about drafting legal filings, processing customer service requests, analyzing financial data, ensuring sales tax compliance, creating content, or developing software. These workflows can absorb small imperfections because humans are still in the loop, but they benefit enormously from AI's speed and pattern recognition.

Fields like healthcare, finance, or legal have high labor costs, "so if you can automate some of it, the ROI is there for the taking," said Saha.

Others pointed to Klarna's bold [pre-IPO decision](#) to sever its relationships with leading enterprise software vendors and replace hundreds of SaaS apps with an AI-powered knowledge graph. It's an extreme example, but it illustrates the magnitude of change possible.

According to Mayfield partner **Sri Pangulur**, the current sentiment is that "industry-specific vertical AI is the new horizontal."

Applied AI represents unprecedented potential change for vertical software, transforming it from what might be called "sleepy SaaS" to one of the most exciting categories for building companies.

**Anirudh Devgan**, CEO of Cadence, agreed. Devgan emphasized that AI's biggest impact will come from domain-specific applications, particularly because the "horizontal parts of AI will get commoditized."

### The augmentation advantage

Cost savings are one thing, but the bigger benefit lies in how AI augmentation enables entirely new business models.

Sri Pangalur points to three key benefits: boosting revenue per customer, growing gross margins, and expanding total addressable markets (TAM). These, in turn, are spurring innovation in how companies price and package their offerings.

DigitalOcean's experience with its WordPress-based Web hosting platform Cloudways provides a concrete example. The platform typically fielded roughly 600,000 annual support calls due to outages, Saha said.

By deploying AI-powered agents that quickly spot the causes of the outages and recommend fixes, Cloudways reduced support costs and expanded its TAM into proactive site management and troubleshooting. What was once a massive operational burden and cost center has become a value driver.

Even in entertainment, AI has a powerful role. **Michael Wise**, the former CTO of Universal Pictures, sees AI as either a positive or negative force.

**“Where we’re really seeing AI being helpful is as a creative accelerant.”**

—**Michael Wise, ex-Universal Pictures**

“In my world, we’re in a moment in time where AI could be a hero or a villain,” Wise said. “A more villainous flavor is if AI is purely a cost savings mechanism,” automating visual effects and replacing writers and actors.

But, he added, “Where we’re really seeing AI being helpful is as a creative accelerant.” That includes artists using AI to create more incredible visual effects or entirely new experiences. In other words, AI also has the power to augment the power of Hollywood’s best storytellers.

### Breaking down legacy barriers

Perhaps the most significant opportunity and challenge lies in integrating AI with the siloed legacy systems that power most enterprises.

**Lisa Dolan**, the managing director of Link Ventures, calls this the “holy grail” in the business of medicine: building AI integrations with siloed systems like CRMs, ERPs, or electronic health care records platforms like Epic.

“With AI,” Dolan says, “the big unlock is to combine these siloed pieces of data.” When health care providers can leverage consolidated information from many sources, they can enable health care providers to deliver care to patients faster than ever before.



The AI doesn't just speed up existing processes. Bringing LLMs to hospitals will enable entirely new workflows and a more comprehensive, personalized view of patients.

The opportunity extends far beyond healthcare. Financial services organizations are integrating data from disparate sources to improve decisionmaking and serve clients more proactively. Legal teams are connecting case management systems with research databases and document archives. Manufacturers are linking supply chain data with production systems and quality control.

### AI that can “read the room”

Even as applied AI reshapes business operations and workflows, founders, researchers, and scientists are pursuing an even more ambitious goal of building human qualities into AI models. The aim is to make AI more understanding, more empathetic, more capable of navigating the complex dynamics of human relationships. If successful, this could transform everything from mental healthcare to materials science.

Next Chapter CEO **Erick Tseng** is working on an AI system that can truly “read the room.”

His mental well-being company is developing methods to “embody AI with relational intelligence.” Based on the concept of emotional intelligence, this is an effort to detect and understand interpersonal dynamics that goes beyond analyzing words.

Consider an AI couples' therapist. Today's AI can process what people say and perhaps detect sentiment in their language. But Tseng's vision is to use computer vision that could detect not just sharp language between a couple, but nonverbal cues like body language, posture, eye contact—or the lack thereof—and even the sudden flush of skin.

“All these different inputs get composited into a math model,” Tseng said, “and what the AI [could] do, which no AI can right now...is say, ‘wait, timeout, I'm going to interrupt the conversation,’” allowing the AI therapist to monitor and intervene before the session goes off the rails.



Erick Tseng



Zhenan Bao

### Teaching AI to feel

While one application of AI is understanding human emotions, researchers are also exploring ways to give AI physical sensations similar to human touch.

**Zhenan Bao**, a professor of chemical engineering at Stanford, is leading work on developing new materials and sensors that can mimic the sense of touch or measure gut health in the intestine. Her lab sits at the intersection of materials science, sensing technology, and AI.

Bao's team is using AI to analyze data from these sensors and design the next generation of sensing materials themselves, creating a powerful feedback loop. Even better, she said, they've found ways to get AI to predict the next set of materials they need for their work. "Our devices are really sensing systems," she said, that "can create a lot of data for training algorithms, and in turn the AI can be used to help us understand what these data really mean."

The upshot? Now the lab can begin to make systems that intuit things like: What kind of sensing information is generated when a person touches an object? How many sensors are needed to emulate human touch? What materials would make those sensors more effective?

By answering these questions, AI is accelerating the pace of materials science research—which in turn creates better sensors, which generate better data, which improves the AI. It's an exponential cycle of improvement.

### The transformation is already underway

The adoption of AI to transform industries is no longer theory, and companies across the globe are increasingly using the technology to change the way they operate, even as AI is also helping develop physical systems that can make people's lives better.

Though it's still early, what was clear is that the potential benefits are enormous. Mental health services could become more accessible and effective. Materials science could advance at unprecedented speed. Human-computer interaction could become more natural and productive.

The momentum behind these efforts suggests that we're approaching a future where AI doesn't just process information, but it understands context, recognizes emotion or touch, and responds with appropriate sensitivity.



# Personalized Medicine, Proven Outcomes

From predicting delivery dates  
to detecting dementia early

## Founder takeaways

- **The virtual provider:**  
Personalized medicine and AI agents are converging to create the next wave of healthcare innovation.
- **Privacy is foundational, not optional:**  
Build HIPAA-compliant, privacy-first solutions from day one, with ethical guardrails baked into clinical workflows.
- **Clinical value drives adoption:**  
Physicians will use tools that demonstrably improve patient outcomes. Founders must prove it through institutional partnerships and validation studies.
- **Integrate or die:**  
Seamless integration into workflows and systems are essential. If it adds friction, it won't get used—no matter how sophisticated.



Erick Tseng and Mudit Garg

Artificial intelligence, long imagined as a potential tool for diagnostics or administrative efficiency, is now reshaping the very architecture of health care. From personalized mental health support to predictive ultrasounds and drug discovery, AI is not just augmenting medicine but transforming it.

With the growth of AI, health care may start to buck its reputation as a slow adopter of new technologies. The use cases are rapidly spreading, with models capable of interpreting medical images and CT scans. For example, AI has already been used to [successfully diagnose diabetic retinopathy](#) from retinal images with greater sensitivity than human ophthalmologists.

The problems AI can solve are no longer theoretical. Some of these AI apps are already being used in hospitals and clinics. As AI models' speed and precision increases, health care stakeholders are no longer looking at merely incremental gains and process enhancements, but paradigm shifts in how they practice medicine.

### Developing trust with providers

The winners in health care transformation won't be those with the most code or capital. They'll be the ones who earn trust.

When **Robert Bunn**, founder of Ultrasound AI, wanted to transition from the oil industry to women's health care, his challenge was to get doctors to trust an AI tool that could predict what they couldn't: the actual delivery date of a baby.

Ultrasound AI uses ultrasound images of the fetus as early as six weeks into the pregnancy and can deliver remarkably accurate results. This enables providers to know if a premature birth or miscarriage is coming, and intervene earlier than ever.

"I got into this for very personal reasons," Bunn said. "My wife had numerous miscarriages, and I decided I wanted to apply my skills to solving that problem."

Doctors initially told him it couldn't be done—especially not by someone without a medical degree. But after a year of work and clinical validation, his tool has now been in use in multiple countries across South America for the past year.

### Using enterprise AI, from automation to doctors' chatbots

**Mudit Garg** sees a similar potential to improve automation across health care. He's the CEO of Qventus, a startup building AI tools for more than 200 hospitals across the country. It is already being used for prepping patients for surgery, scheduling, and releasing patients after care.

For patients, accessing health care can often feel like "a rock being pushed up the hill," Garg said. Tackling the data problem helps them manage the patient journey with greater ease.

“For patients, accessing health care can often feel like a rock being pushed up the hill.”

—Mudit Garg, Qventus

In health systems, the degree of unstructured data and interactions (everything from visits, claims, faxes) far outpaces anything we capture in a structured standpoint, Garg explained. Structured data is organized in a predefined format, making it easy to process and analyze; unstructured data has typically been much harder to understand.



Additionally, a host of logistical inefficiencies—missed appointments, redundant paperwork, fragmented records—have plagued health systems.

AI agents can now streamline these processes, parsing unstructured data like faxes and handwritten notes, and generate actionable insights. What once required a team of nurses and administrators can be handled by intelligent systems in seconds, for instance to manage surveys and intake forms. It’s also possible to use AI to give providers a more complete picture of patients, Garg said.

Garg sees greater adoption and hunger for these kinds of innovations, especially given the rising costs of health care for both systems and patients. Manually doing everything isn’t sustainable, so these hospitals have been rapidly adopting the platform and its AI features.

At Stanford Health Care, a new AI agent for clinicians helps expedite daily tasks, including chart reviews and asking specific questions regarding patients. ChatEHR, artificial intelligence software developed at Stanford Medicine, allows the staff to interact with patients and data in plain language, much like in ChatGPT. For example, they can ask: “Has this patient ever received therapy?” or “How many times has this patient received a blood thinner?”

As ChatGPT became popular, it became obvious to the Stanford team that it would make sense to create an app for direct care, said **Topher Sharp**, clinical professor in medicine, care and population health at Stanford Health Care. “For physicians, this would be an ability to augment their search,” Sharp said.

Sharp knew that faculty at Stanford would use ChatGPT if they didn’t have a better alternative. So ChatEHR draws from a variety of data sources, including Epic Systems, a widely-used Electronic Health Record (EHR) software platform, to access patient data. But it does so in a secure environment, ensuring the data isn’t being used to train the model.



Jarrod Helzer, Anurang Revri, and Topher Sharp

“Our core business isn’t AI, but we do want to use it,” added **Nigam Shah**, professor of medicine at Stanford and chief data scientist at Stanford Health Care.

As with many AI integrations, the end goal is to reduce the cost of health care by reducing and streamlining repetitive tasks. But the broader vision extends far beyond cost savings to improving—and even transforming—the way patients receive care.

“Our core business  
isn’t AI, but we do  
want to use it.”

—Nigam Shah, Stanford

# The Race to Build Intelligent Machines

From warehouse floors to hotel lobbies:  
Robots meet reality

## Founder takeaways

- **Brave new world:**  
Physical AI systems face an impossible challenge: preparing for scenarios that haven't been imagined yet, let alone tested.
- **It's a wild world:**  
Real-world deployment is the only way to discover edge cases and failure modes—laboratory testing cannot replicate the complexity of uncontrolled environments.
- **Safety is crucial:**  
Founders must balance ambitious vision with pragmatic execution and safety, recognizing that different applications have vastly different risk tolerances.

Ashish Kumar, Rodney Brooks, and Nico Enriquez

The rise of AI has reinvigorated the robotics industry. Robots were once simply pre-programmed machines that were effective only within highly controlled environments. But neural networks and LLMs have breathed new life into robotics, enabling companies to take the first steps towards building machines that perceive their surroundings and make real-time decisions about how to interact with the physical world.

The term “physical AI,” also known as embodied AI, describes systems that understand and adapt to the physical world with the help of software and sensors and the continuous processing of sensory data. In a very brief time, physical AI has begun transforming key sectors beyond robotics: self-driving vehicles, healthcare, military weapons, manufacturing, and security.

“I’m excited about the capabilities that we’re going to have in manufacturing, and really take that to a new level,” said **Steve Cousins**, executive director at the Stanford Robotics Center. “I’m also excited about the possibility of taking robots out of constrained environments into real human environments, and being able to operate there successfully.”

But if the industry is to one day develop machines with the agency and utility of C3PO and R2-D2, it must first overcome many technological challenges. For ambitious technologists, the sector offers the chance to solve some intriguing and daunting problems.

“I’m excited about the possibility of taking robots out of constrained environments into real human environments.”

—Steve Cousins, Stanford

#### Imperfect humanoid robots

Not everyone agrees on the best direction forward. **Rodney Brooks**, the pioneering roboticist who co-founded Roomba maker iRobot and developed foundational robotic architecture, offered a less-than-dazzling view of the near-term potential for humanoid robotics, suggesting that robots would continue “doing real work” primarily in warehouses, plants, and factories—sectors where they’re already deployed.

His most pointed critique? “We have zero grasping capability,” he said, highlighting how humanoid robots still struggle with taking basic actions.

For **Ashish Kumar**, a Tesla researcher who worked on that company’s Optimus AI and Robotics program, perfection isn’t a prerequisite for commercialization or revenue generation. He pointed to AI chatbots, AI video production, and self-driving cars as examples of imperfect technologies already generating substantial value.



The issue with humanoid robots, he argued, is simply scale: Once production volumes increase, costs will plummet, like they do with other technologies.

### Engineering challenges remain

Those seeking to build better robots face an array of technical hurdles. Some of the challenges include building smaller and more perceptive sensors, more stability in two-legged robots, smaller sensors that enable robots to better perceive diverse environments and improved capabilities for robot actuators—the components that drive robotic limbs.

AI has enhanced robots' real-time decision-making, but providing them with absolute autonomy in any environment is still a ways off. The problem is that there are too many variables in the physical world, which requires robots to process too much data.

These aren't abstract problems. **Samir Menon**, CEO of Dexterity, an AI-powered robotics provider, shared a story about a robotic arm his company built that excelled at moving mail packages—until it malfunctioned one day. His staff discovered the arm clutching a perforated box leaking liquid that damaged its components.

The culprit? The box was designed to let earthworms inside breathe. No amount of pre-deployment testing had anticipated earthworms by mail.

### When robots meet the real world

Along the same lines, Stanford's Cousins shared an anecdote about putting a robot into a hotel.

"When you're in factories, you're behind safety cages or you're very careful to control the environment," Cousins said. "You can make everybody wear steel toe boots, and you can keep people behind yellow lines.

"But when we put our robot into a hotel for the first time, the very first time we looked at what the robot's cameras were seeing, we saw a three-year-old's bare toes. A child had obviously hugged the robot. We're like, whoa! We weren't thinking about this when we designed the safety system."

That incident prompted a complete redesign of the safety system, Cousins said.

### Building on progress

The many challenges notwithstanding, newcomers to physical AI will find a host of innovative tools and pioneering technology to build on. Embodied intelligence has made massive strides in touch-and-feel perception. Robot locomotion has matured, resulting in improved motors. Quadrupeds, the robots that get around on four articulated limbs, are more agile.

The manufacturing and warehouse sectors are seeing real commercial traction. These environments, while less constrained than traditional factory settings, still offer enough structure to make current physical AI capabilities viable.

For all the legitimate technical concerns, many in the field remain bullish on physical AI's trajectory.

That tension—between measured caution and audacious ambition—will likely define physical AI's evolution over the coming decade. The winners won't be those who choose one approach over the other, but rather those who know when each is appropriate.

The robots are learning to touch our world. Now we need to figure out if we're ready to touch back.

### Developing the World's Most Trusted Driver

Autonomy is a major challenge for physical AI, but Waymo already employs AI to ensure its cars are safe on the road, even in the face of the most unexpected external actions.

“Modern AI enabled us to bring together all the components that you need to build an autonomous agent that operates in the real world,” Waymo distinguished engineer **Vincent Vanhoucke** said. “You can interpret images and sensor inputs, you can reason, you can plan—and all of that is in a common framework.”

As Vanhoucke put it, following the rules of the road is the easy part. The much harder problem to solve is how to engineer the company's vehicles to respond properly when people, animals, and other vehicles “do crazy stuff.”

The Alphabet-owned company's cars feature more than 20 high-resolution cameras, meaning that at any given moment, they're ingesting a huge amount of data about what's going on around them. “We have to assume the worst in many situations,” he said, and “at scale, operating the long tail of difficult scenarios dominates the technology.”

Previously, Waymo leaned on the data it derived from millions of miles of autonomous driving. Now it's able to leverage vision language models that can understand and parse the world outside the vehicles at a deep level, in real time.

While the company's cars have already been on the streets of San Francisco, Los Angeles, Phoenix, Atlanta, and Austin, it plans to grow to many other cities over time, so it needs an AI that's capable of adapting quickly to new locations.

To that end, Waymo built its own model based on the types of driving data its cars collect—images, context, maneuvers, and so on—that it uses to predict how its vehicles should navigate city streets, even streets that are new to it.

The resulting performance was on par with the previous on-board-only systems yet is far more scalable and simple. And now, it's working to refine the models even further.

In that effort, the conversational nature of modern AI is a plus.

“The nice thing about driving is that it's like a conversation,” Vanhoucke said. “You're stopping at a stop sign and there's another car there, and you're going to nudge forward, nudge forward, maybe they go and then you stop. It's a visual conversation, and that's not just a cute metaphor. It's literally the problem that we have to solve.”

“We have to assume the worst in many situations, and the long tail of difficult scenarios dominates the technology.”

—Vincent Vanhoucke, Waymo



# Winning Over Lawmakers and Public Means Earning Trust

If they're to influence AI policy decisions, AI leaders must become educators about the benefits of the technology

## Founder takeaways

- **Many different regimes:**  
Global AI regulation is fragmenting rapidly across the U.S., Europe, China, and Asia. Understanding these differences isn't optional—it's survival.
- **It's not just about compliance:**  
Earning trust from policymakers and consumers is the foundation for influencing AI policy. Safety and transparency are business imperatives, not just checkboxes.
- **Learn to tell your story well:**  
AI leaders must become educators, not just innovators. The ability to communicate benefits clearly can prevent legislative backlash.



At home and abroad, companies working with AI should keep close tabs on how their AI products or services are perceived—by employees, the public, and especially lawmakers.

AI remains a widely misunderstood technology and startups attempting to expand their businesses into different markets must navigate rapidly changing legal and regulatory environments.

Meanwhile, breathless stories in the media continue to stoke fears about AI replacing workers, criminals misusing AI to commit fraud, and students relying on AI to cheat on tests.

Setting the record straight is critical. If AI leaders can't clearly communicate the advantages of AI, legislative and consumer backlash are real possibilities.

### Navigating complex regulatory environments

Companies operating globally already must adhere to [varying rules in Europe](#), the United States, China, and other countries—some more complex and burdensome than others.

According to **Jeff Hancock**, the director of the Social Media Lab at Stanford, Europe emphasizes precautionary oversight, with lots of regulation but spotty enforcement. The U.S. leans toward market-driven innovation and more lax regulation, but with excellent enforcement for the regulations that do exist. China prioritizes state control and social stability with strict regulation and enforcement. And other Asian markets, like Japan and Korea, have very little regulation and enforcement, because they want to encourage as much innovation as possible.

Regulation impacts many aspects of the technological supply chain, creating uncertainty for companies unprepared to deal with the varied legal landscape. For example, in recent years, the U.S. and China have each attempted to prevent [NVIDIA and AMD](#) from selling high-performance [GPUs in China](#).

“We have shifted our entire supply chain to make sure that any final aspect of the supply chain is not done in China,” said **Rajiv Khemani**, co-founder and CEO of Auradine. “There is some uncertainty with regard to what could happen.”

Even within the U.S., state-level regulations are proliferating. Illinois recently passed the [Wellness and Oversight for Psychological Resources Act](#), which banned the use of chatbots to provide therapy without medical supervision. The new law also prohibited other AI uses in mental health care, even by licensed professionals. Critics argue the law unfairly prevents legitimate AI-powered applications from helping people who could benefit from them. It's a cautionary tale about what happens when innovation outpaces public understanding.

### Sovereign AI and military defense

Adding to the confusion is the rise of sovereign AI, or AI models that are exclusively controlled by a country in order to adhere to its specific national interests.



Howard Wright

**Howard Wright**, head of the startup ecosystem for NVIDIA, talked about what an enormous opportunity sovereign AI represented.

As they compete to become the world's dominant AI power, China and the United States have begun adopting different sovereign AI strategies. Other countries are also pursuing this, building their own LLMs running in their own, secure datacenters in order to protect state secrets.

AI also has a role to play in warfighting. The war in Ukraine shows that the era of autonomous weapons is here. In mid-2025, Ukraine launched an unprecedented and orchestrated attack on five Russian airbases [by AI-enhanced and semi-autonomous drones](#). The attack damaged or destroyed 41 aircraft.

Military experts now agree that AI may become the catalyst for a new autonomous weapons age.

"We're going to see more AI tooling used to control and orchestrate drone attacks," said U.S. Air Force

Col. **Jason Hansberger**. "But if we build this complexity and this level of mass (attacks), then we're going to need a lot of AI tools to help in decision making."

Alongside the growth of sovereign AI is a fear that international competition might stifle the overall growth of the AI industry by limiting global information sharing.

International rivalries might be a speed bump, but won't fundamentally change the advantage of open source, said **Jose Plehn**, CEO of BrightQuery.

"I strongly believe that the solution to open source, AI, and any potential competition that may arise out of that, is open data," Plehn said. "Data and knowledge have no borders. Once knowledge is out there, it's out there. It's very hard to contain."

**"The solution to open source, AI, and any potential competition that may arise out of that, is open data."**

—Jose Plehn, BrightQuery

### Start by building trust

When it comes to influencing policy and winning over lawmakers and the public, the best thing startups can do is earn trust. Achieving that starts with proving that their AI works the way it's supposed to and isn't a threat, **Anna Makanju**, vice president of Global Affairs for OpenAI, said during one panel discussion.

"The thing that gives me the most optimism (about AI) is that there's a business imperative for these things to be safe," said Makanju, a former White House foreign affairs official. "Ultimately, consumers and governments don't want to integrate something that's going to be unsafe."

“There’s a business imperative for these things to be safe.”

—Anna Makanju, OpenAI

Another way to build trust, Makanju added, is to design AI systems that enhance core human needs, specifically autonomy and competence. When serving regulated sectors, such as financial services and healthcare, companies must incorporate accountability, traceability and human decision into their systems.

One way to understand a market’s regulatory environment is to hire experts with a global political perspective. **Quentin Clark**, managing director at venture firm General Catalyst, said during a panel discussion that his company works with the [Tony Blair Institute for Global Change](#) to help think through the impacts of AI applications.

#### Real risks and anxieties remain

Criminals continue to think up new ways to employ AI to commit fraud on businesses and individuals. Parents and educators are alarmed about how some

school children use AI to cheat on tests. Workers around the world still fear AI will cast them into unemployment.

Anxiety over potential job loss exists even in the tech sector.

“When you look at AI and its capabilities without proper context, you feel like AI might be replacing software engineers,” said **Ursheet Parikh**, a partner at venture firm Mayfield. “AI agents are writing code and fixing bugs... but if AI can do 50% of your work, you now have 50% of your other time to review the outcome of the AI and train the AI to do better next time.”

Parikh said that everyone knows AI-generated code is not perfect and requires human review.

Others at the conference predicted AI automation will mean job augmentation rather than job loss. Several speakers said they see a day coming when every worker will supervise and manage their own groups of AI-agent underlings. But that optimistic future requires clear communication now.

#### The cybersecurity opportunity

Defending against AI-powered crime represents an opportunity for AI startups. **Eric Schmidt**, the former Google CEO, said the world faces three key potential threats: 1) from some kind of new bio-weapon, 2) misinformation aided by new AI-powered deep fakes, which can undermine democracies, and 3) more sophisticated cyber attacks. “If you can write code, you can also write cyber attacks,” Schmidt said. “AI models program better than I ever did.”

Anna Makanju



# “AI models program better than I ever did.”

—Eric Schmidt, former Google CEO

AI security companies need to become skilled AI ambassadors to help important constituencies understand the nature of the problems and how AI can help.

“Everybody has a firewall to make sure the bad guys aren’t coming in, and everybody is using some type of web proxies or secret service edge to keep the data safe,” said **Moinul Khan**, cofounder and CEO of Aurascape, an AI security company. “AI (security) tools are different. They’re not static software applications. AI tools are really dynamic. They’re learning. They’re adapting and evolving in real time.”

To deal with dynamic, AI-powered attackers, enterprises need a dynamic, AI-powered defense, Khan added. Traditional cybersecurity defenses won’t cut it.

“Your existing enterprise network, the inline stack that is supposed to protect you? 90% of the time they’re completely lost.”

## Charting a path forward

The companies that successfully navigate AI’s regulatory future will share common characteristics: They’ll prioritize safety and transparency, invest in policy expertise early, communicate benefits clearly to multiple audiences, and build compliance into their products from day one.

Most importantly, they’ll recognize that regulation isn’t just a constraint—it can be a competitive moat. The founders who master regulatory complexity will be positioned to scale globally while others remain stuck in single markets. They’ll be the ones who earned trust, educated stakeholders, and turned policy challenges into strategic advantages.

Moinul Khan



FUTURE OF AI POLICY AND GOVERNANCE



Erik Brynjolfsson and Anna Makanju

# What's next?

Join the conversation. Connect with fellow builders. Shape what's next.

- For future Imagination in Action events:  
[www.imaginationinaction.co](http://www.imaginationinaction.co)
- To learn more about Mayfield:  
[mayfield.com](http://mayfield.com)
- Subscribe to Founder Insights:  
[Founder Insights Newsletter](#)



## The Future OF AI

Fall 2025

### Publication Credits

Karen Lee/Mayfield  
Dylan Tweney/Tweney Media  
Terhi Ignatius/60Nrth  
Imani Sires  
John Werner

Editor-in-chief  
Managing editor  
Designer  
Photographer  
Photographer

Antoinette Siu  
Daniel Terdiman  
Greg Sandoval

Writer  
Writer  
Writer