LEONIS CAPITAL

The Leonis AI 100

Benchmarking the Most Important AI Startups of 2025

Nov 5, 2025

Table of Contents

Acknowledgements 4

Executive Summary 5

About This Report 7

Methodology & Dataset 7

The State of AI Startups 9

Insight 1: The Rise of the Researcher Founder and Technical CEOs 10

Insight 2: The Post-2024 Revenue Surge 14

Insight 3: Smaller and Flatter Teams, Huge Output 16

Insight 4: PLG First, Sales Later 20

Insight 5: Multiple Winners, Not Winner-Takes-All (Yet) 22

Insight 6: Winners Pivot Fast, AI Makes it Possible 24

Insight 7: Markets Break Out in Sequence 27

Who's Backing the AI 100 29

Conclusion 33

DISCLAIMERS

This document is provided for informational purposes only and should not be construed as investment advice, a recommendation, or an offer to buy or sell any security, instrument, or strategy. Nothing herein should be relied upon as financial, legal, accounting, or tax guidance.

The content reflects information available to Leonis Capital (together with its affiliates, “Leonis”) at the time of publication, including third-party sources believed to be reliable and the author’s own analysis. While efforts have been made to ensure accuracy and completeness, Leonis does not guarantee that the information is correct, complete, or current. All views and opinions are subject to change without notice, and neither Leonis nor the author(s) accepts liability for errors, omissions, or reliance placed on this material.

This document may include forward-looking statements, projections, or illustrative examples. Such statements are inherently uncertain, depend on assumptions, and are subject to factors beyond Leonis’ control. Actual results may differ materially from those anticipated. Past performance is not a reliable indicator of future returns, and investing always involves risk, including the potential loss of principal. Different types of investments involve varying degrees of risk, and there can be no assurance that the future performance of any specific investment, strategy, company, or product referenced, directly or indirectly, will be profitable, equal any indicated performance level, or be suitable for any particular portfolio.

Leonis Capital is an Exempt Reporting Adviser (“ERA”) under the Investment Advisers Act of 1940 and applicable state regulations. ERA status does not imply a particular level of skill or expertise, nor does it represent endorsement by any regulatory authority. This material is intended for a general audience and does not account for the objectives, financial situation, or constraints of any individual, nor does it establish an advisory relationship with Leonis.

Leonis, its affiliates, and its clients may hold, increase, reduce, or dispose of positions in companies or securities mentioned, independent of any views expressed herein.

AUTHORS



Jenny Xiao

Jenny Xiao is a Partner at Leonis Capital. She was previously a researcher at OpenAI and received her Ph.D. from Columbia University.



Jay Zhao

Jay Zhao is a Partner at Leonis Capital. He is a long-time VC with 3 IPOs and 8 unicorns in his portfolio.



Liang Wu

Liang Wu is a Research Affiliate at Leonis Capital. He was previously a Senior Researcher at Harvard Business School.



Oliver Wang

Oliver Wang is a Venture Fellow at Leonis Capital and a Ph.D. candidate at MIT in Computing at the Department of Aeronautics and Astronautics.

Acknowledgements

This report would not have been possible without the contributions of many individuals who generously shared their insights, data, and expertise.

We're grateful to the Leonis research fellows who contributed analysis and research, especially Armita Hosseini (Stanford University).

Thank you to our data partner, Sacra, who provided crucial data, and particularly Marcelo Ballvé, who offered editorial feedback on many drafts.

We especially thank Todor Markov (Anthropic, ex-OpenAI), Summer Yue (Meta Superintelligence), and Jason Risch (Greylock) for their in-depth feedback on early drafts. We also appreciate the many others who shared helpful comments, including Tim Lee (ex-Sequoia Capital), Aaref Hilaly (Bain Capital Ventures), David Kanter (MLCommons), Jasmine Wang (OpenAI), Allen Wu (Nvidia), Stefano Corazza (Canva), Shikhar Ghosh (Harvard Business School), and Jordan Nanos (SemiAnalysis).

Executive Summary

AI has rewritten the startup playbook. To understand how, we analyzed a dataset of over 10,000 companies from 2022 to 2025 and identified the top 100 fastest-growing AI-native startups. These startups look very different from their SaaS era counterparts.

The most striking change is in the founders and teams. Today's breakout companies are led by deeply technical CEOs: PhDs, Olympiad medalists, and alumni of top AI labs. These teams are small but unusually productive and agile, often achieving revenue per employee five to ten times higher than the SaaS benchmarks. Their experimental mindset and fluency with the underlying models allow them to pivot quickly, often within months of founding, to find product-market fit.

The business performance is equally dramatic. After years of skepticism about AI monetization, 2024 marked a turning point. Dozens of startups crossed $10M ARR in under 18 months of founding their companies, mostly through self-serve, product-led adoption. Users experience value immediately, products spread virally inside organizations, and only once adoption reaches critical mass do formal sales motions begin

The AI markets themselves are unfolding in sequence. Horizontal tools like writing and coding assistants broke out first, followed by multimodal media tools and, more recently, vertical applications in healthcare, legal, and finance. Unlike past cycles where one winner dominated, multiple companies are scaling side by side, each securing moats through workflow depth, data advantages, or vertical specialization.

Taken together, these patterns suggest the AI era is reshaping both how companies are built and where value may ultimately accrue. Team size, sales motion, and revenue velocity look very different from the SaaS playbook. But there are many open questions that remain to be answered. How durable are these early advantages once foundation models absorb more functionality? Which verticals will unlock next as reliability improves? And what forms of defensibility, whether workflow depth, proprietary data, or regulatory moats, will prove long-lasting?

The rest of this report explores each of these patterns in depth, drawing on data from the top 100 AI startups from our dataset.

How to Use This Report

If you're a **founder**, this report might serve as a roadmap for building in the AI era. Insights 6 (Pivots) and 7 (Market Sequence) highlight when to enter and how to pivot ahead of major model shifts. Insights 1 (Founders) and 3 (Teams) reveal what high-performing AI organizations look like (notably, they are small, technical, and research-led) and how that differs from the SaaS playbook. Insight 4 (PLG) outlines how to convert product-led traction into scalable enterprise sales.

If you're an **investor**, the report can help inform pattern recognition and shed light on where value is forming and how to evaluate it. Insights 5 (Multiple Winners) and 7 (Market Sequence) highlight markets that can sustain multiple winners and where defensibility still compounds. Insights 2 (Revenue), 3 (Teams), and 6 (Pivots) provide updated efficiency benchmarks (true recurring revenue, revenue per employee, and pivot speed) that better capture performance in the AI-native era.

About This Report

Leonis Capital is a research-driven venture firm focused on AI-native companies. Since 2021, we've partnered with technical founders at the earliest stages, combining technical expertise with institutional investing experience.

After three years of systematically tracking the AI startup ecosystem, we've seen patterns that we think are worth sharing with the AI community. Too much of the AI discourse is driven by hype and speculation. This report draws instead on one of the largest private datasets of AI startup metrics, built and maintained as part of our work as investors.

We hope this analysis helps the AI community reflect on what has and hasn't worked in building AI startups.

Methodology & Dataset

Much of this data set is based on Alfa, our proprietary AI research system. Alfa continuously indexes and analyzes AI startups from around the world, tracking everything from GitHub commits to hiring velocity to revenue signals. It helps us scale our analysis as an AI-native venture firm.

Between 2022 and 2025, Alfa indexed more than 10,000 AI startups globally. From this universe, we identified the fastest-growing startups based on publicly available data and signals (funding, hiring, usage, GitHub activity, Crunchbase, Product Hunt, news announcements, and ARR estimations from sources like Sacra and ARR Club) and private references from founders and investors. We've aggregated that into the Leonis AI 100 to glean further insights on the fastest-growing AI startups.

To qualify, companies must be:

- Primarily focused on AI applications, not foundation model development, since these labs operate with fundamentally different capital intensity, GTM motions, and defensibility dynamics[1]
- Founded or built their product in the post-ChatGPT era (2022 or later), we excluded legacy companies that rebranded or pivoted into generative AI
- Building AI-native products (e.g., LLMs, agents, AI infrastructure, generative media, etc.)
- Independently operated (i.e., no spin-outs from Big Tech)
- Showing clear traction (e.g., revenue, users, open-source adoption)

We tagged each company by:

- Product category (infrastructure, agents, video, devtools, etc.)
- Founder backgrounds (research, ex-big tech, co-working history, age at founding, etc)

[1] We respect the work of foundation model labs, but including them in the data would distort the insights as funding amounts, team size, and ARR growth at foundation model companies are drastically different than AI applications. Although foundation models like ChatGPT are also popular consumer apps and compete directly with application startups, we believe that model labs are in a category of their own.

- Geography and location
- Founding date
- Investors (Pre-Seed, Seed, Series A, and Series B)[2]
- Round and valuation trajectory
- Revenue growth trajectory[3]
- Go-to-market motion (product-led, marketing-led, community-led, and sales-led)[4]
- Pivots and time-to-pivot

All data is current as of September 2025. We're a small team, and this is a fast-moving space. We've likely missed things. We've open-sourced the dataset and invite founders, operators, and fellow researchers to help us improve it.

View the dataset here.

Please email research@leoniscap.com if you have any suggestions, corrections, or additions.

[2] We want to focus on tracking investors who identified the AI 100 companies early on and supported them in their initial growth. Our data is only based on publicly available information, thus, we likely have omitted some investors.

[3] This is based on a combination of news or PR announcements and revenue estimates from sources like ARR Club, Sacra, and Lakta. There will likely be inaccuracies given that many data points come from estimates.

[4] We categorized the companies' GTM motion into four buckets: 1) product-led growth, growth driven by word-of-mouth and using the product itself as the driver of customer acquisition; 2) sales-led growth, which is driven by top-down enterprise sales; 3) marketing-led growth; and 4) community-led growth.

The State of AI Startups

The last three years in AI felt like three decades.

In November 2022, ChatGPT launched and changed everything. What followed was a gold rush of demos, Twitter threads, and breathless predictions about the future of software. But beneath the hype, a more interesting story was unfolding: the emergence of a new breed of startups that is defining company building in the AI era, deviating from the SaaS playbook that defined the previous era of successful ventures.

By early 2023, thousands of AI projects had sprung up, but revenue was lagging behind usage. Many observers were skeptical that these cool demos would turn into real businesses; some dubbed it an “AI bubble” and asked where the revenue was.

The narrative changed dramatically in 2024. As model capabilities improved, customers began paying in earnest. Entire categories, from voice to video and early agent platforms, moved from demos to production. Revenue traction separated substance from hype, and the best startups proved they could scale quickly and efficiently.

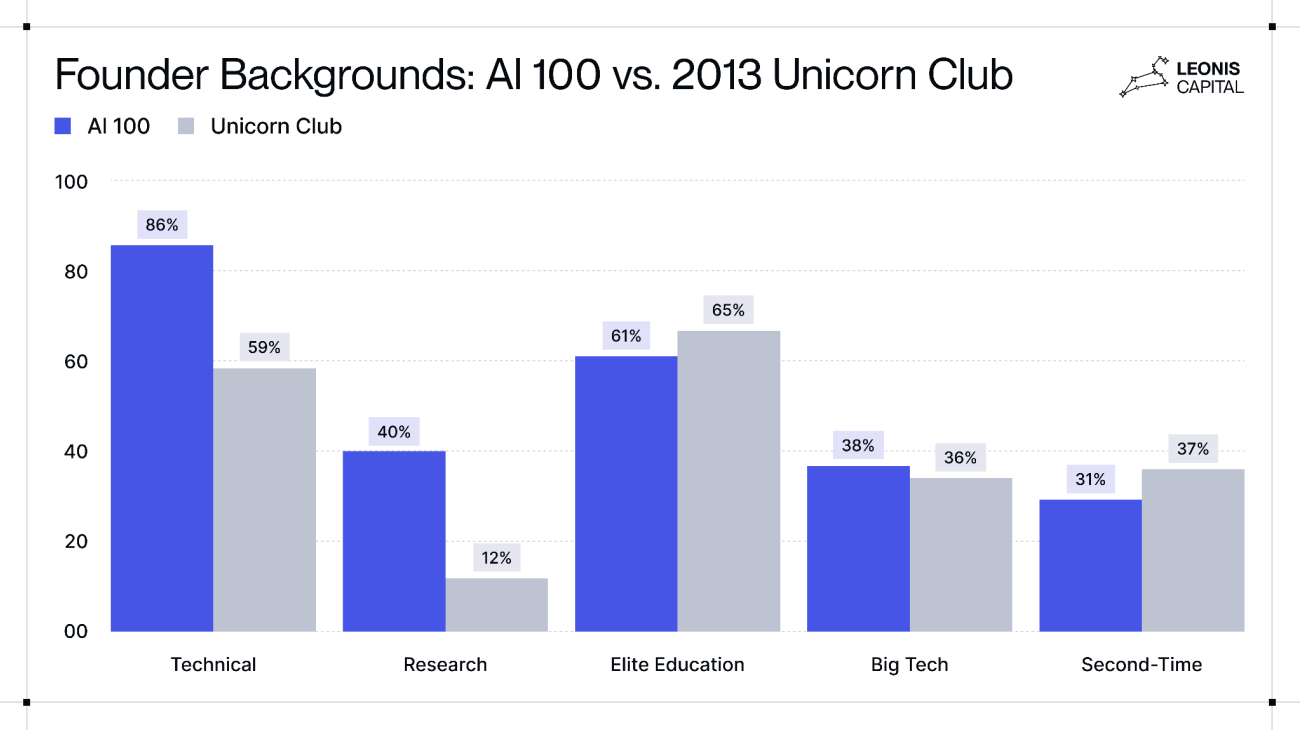
In 2025, adoption began spreading into complex verticals such as healthcare, legal, and finance. These markets demand reliability, compliance, and workflow depth, traits that raise the bar but also create stronger moats for those who succeed.

In just three years, AI startups have moved from hype to skepticism to real traction, compressing an entire market cycle into a fraction of the time that it took companies in the previous eras of technology to build out its products and businesses. From our analysis, seven patterns stand out: who the founders are, how companies scale revenue, the structure of their teams, their go-to-market motion, the shape of competition, the speed and nature of pivots, and the order in which markets unlock. The sections that follow explore each of these dynamics in depth.

Insight 1: The Rise of the Researcher Founder and Technical CEOs

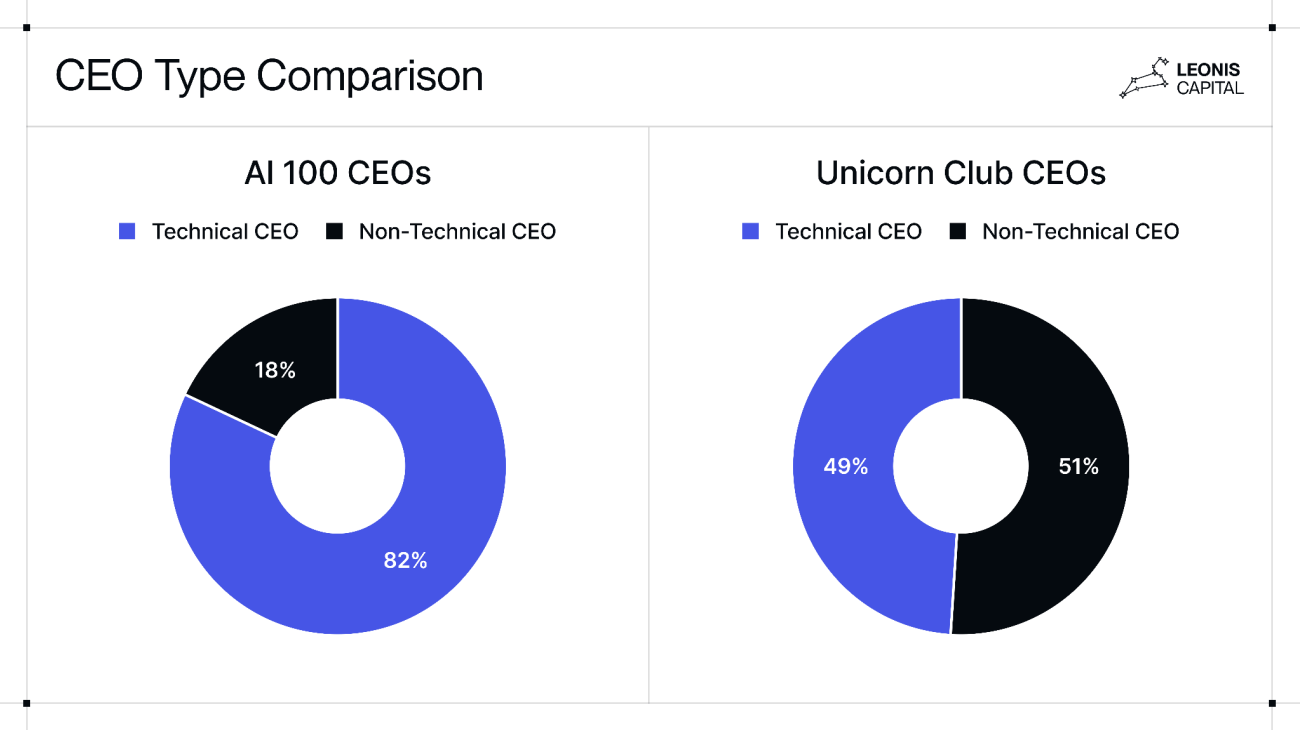
The most striking pattern in our data isn't about products or markets, it's about the people building these companies. Today's AI founders look very different from those who dominated the last decade of startups. Looking back at the 2013 Unicorn Club, when Aileen Lee coined that term, she found that outlier outcomes came from "well-educated, thirtysomething co-founders who have history together." The dominant archetypes of that era were serial entrepreneurs with prior exits, scrappy hackers shipping consumer products, product managers with customer insights, and business operators with MBAs.

This is changing in the age of AI. Today's breakout AI startups are **overwhelmingly led by founders with elite technical and research backgrounds.** Berkeley PhDs, MIT researchers, International Math Olympiad medalists, and alumni from OpenAI or DeepMind have replaced MBAs and product managers in the CEO seat. The SaaS-era model of a business-background CEO paired with a technical CTO seems to be replaced by dual-technical teams. Across the Leonis AI 100, 82 out of 100 companies are led by technical CEOs, and a striking 208 of the total 241 founders (86%) are technical. By comparison, in the Unicorn Club, only 49% of companies had technical CEOs, and only 59% of founders were technical.



This rise could reflect two related forces. First, many more technical professionals are founding AI startups than in previous waves, as the barrier to entry has dropped and venture investors actively encourage researchers and engineers to start companies. Second, among those who do, the most technically fluent founders disproportionately succeed.[5]

[5] Our dataset captures this survivorship bias: it reflects the companies that broke out, not the total pool of attempts. An interesting piece of future research should look at how the AI 100 companies compare with the total pool of AI startups.



The biggest shift shows up in research pedigree. The vast majority (58%) of the startups in the Leonis AI 100 have at least one co-founder with a research background.[6] 40% of all founders in this cohort have research experience versus just 12% in the Unicorn Club. The rise of the "researcher founder" reflects a founder archetype that barely existed in the last wave. In the AI age, founders who have seen how frontier models are trained understand where capabilities are improving, which areas remain constrained, and what directions the labs are actively investing in. That insider knowledge helps them target opportunities just ahead of the curve, knowing, for example, which capabilities are likely to improve soon and which still require fundamental research breakthroughs.

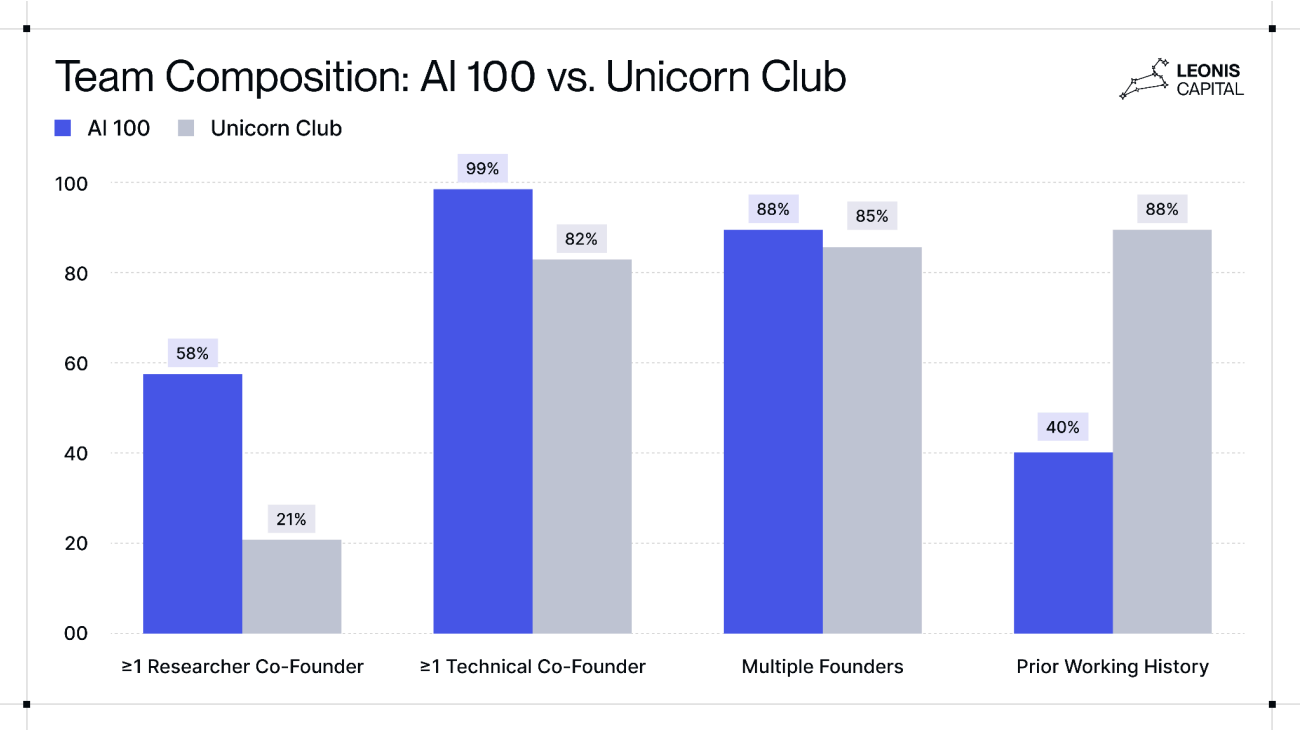
Importantly, these aren't pure researchers who've spent decades in academia either. The most successful founders combine research credentials with a builder's urgency. Many did brief stints at top labs, just long enough to understand the frontier, but not so long that they lost their builder instincts.

The rise of researcher founders originates from the fact that in AI, the tech *is* often the product. Unlike in the SaaS era, where the tech is often just a medium for something else, like marketplaces for homes (Airbnb), taxi services (Uber), or e-commerce products, the core product in AI is the tech itself. AI companies compete on first-principles choices at the model, data, and system level, decisions that need leaders with deep technical and research expertise.

This dynamic extends to every aspect of building an AI company: VCs want to back founders who've shipped production AI systems, not those who have only managed teams that built them. The best AI engineers want to work for someone who understands their work at a deep

[6] We define "research background" as having a graduate or research degree in a technical field, or meaningful research experience such as working at a research institution, publishing papers, or contributing to frontier AI research projects.

level. Even customers, especially technical buyers, expect to discuss implementation details with someone who truly understands the underlying technology.

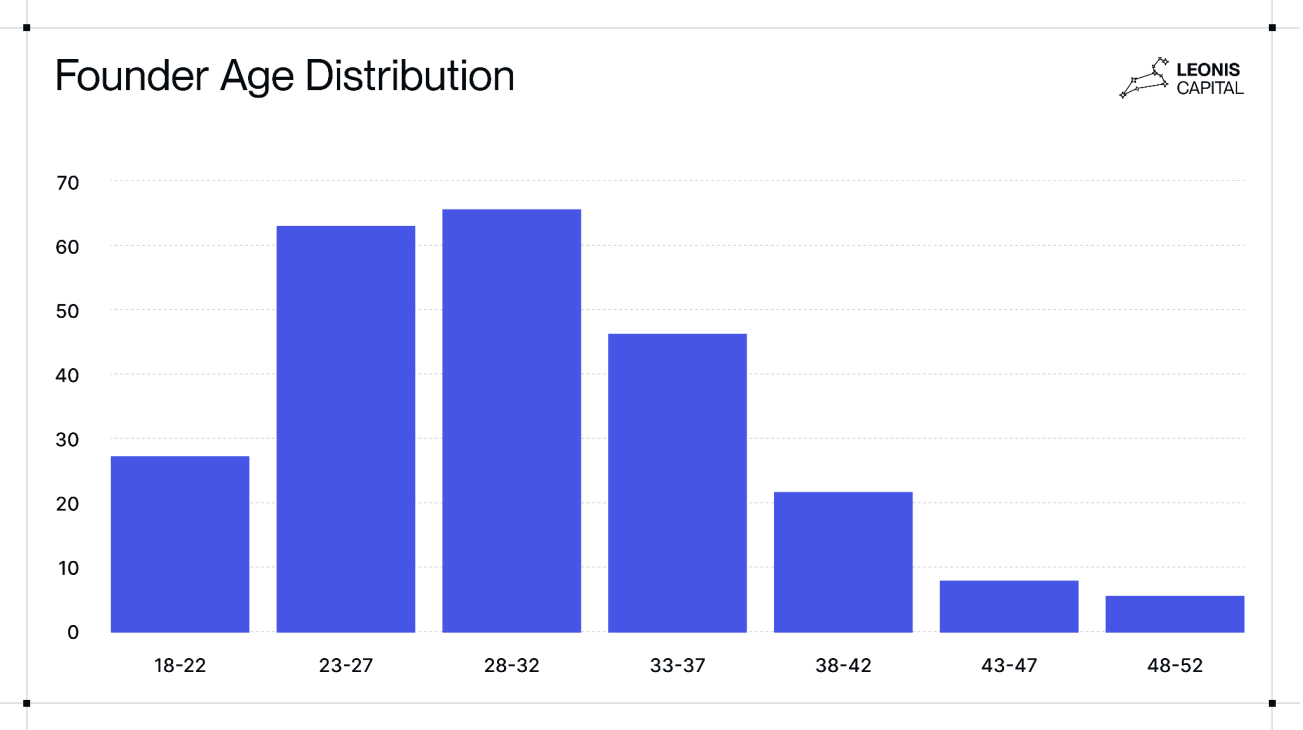


While VCs love to claim that credentials don't matter, our data says otherwise. More than 60% of the founders in our dataset have elite educational backgrounds, with degrees from MIT, Stanford, Harvard, and other top programs.[7] This continues a pattern from the last generation of unicorn founders, many of whom also came from prestigious schools. The difference in this wave is that those credentials are overwhelmingly technical: PhDs in computer science, EECS, or mathematics, and research experience at leading AI labs. Aravind Srinivas, CEO and cofounder of Perplexity, for example, completed his PhD at UC Berkeley and worked as a researcher at OpenAI before starting Perplexity. The Cursor team, though still undergraduates at MIT, had worked on AI research projects and were already paying attention to LLMs back in the GPT-2 era, spotting the AI coding opportunity before it became obvious.

Team formation also looks different. Only 12 of the AI 100 are founded by solo-founders, and while most companies have multiple co-founders (88%), just 40% had co-founders with a prior history of working together.[8] In contrast, almost all (88%) Unicorn Club companies had co-founders who had worked together before. Familiarity was the rule then; today, new teams form around technical opportunity rather than long personal history. Of course, some of this difference may reflect the stage of company building: the Unicorn Club companies were later-stage, while the AI 100 includes many earlier ventures, and the most successful among them may eventually converge to the patterns of the previous generation.

[7] 147 out of 241 founders fit into this category.

[8] Our general sense is that among AI startups, the percentage of solo-founders is much higher.



The average top AI founder is also **younger than their SaaS predecessors**, with a median age of 29 compared to 34 for SaaS. The most frequent age at founding is 26 or 27, nearly a decade younger than the prior cohort. These founders tend to jump directly from academia or research labs into startups rather than climbing the corporate ladder first. Nevertheless, extremely young college dropouts remain rare in both waves, suggesting that building successful AI products still requires experience gained from prior research or engineering roles.

The demographics of AI founders also mirror the categories that broke out first: AI coding tools, search, and generalist agents. As adoption spreads into vertical markets like healthcare and finance, we expect more domain experts who trend older to enter the founder pool. But so far, technical depth has consistently trumped industry experience. Even in vertical domains, 4 out of 13 AI 100 founders lack direct sector backgrounds, suggesting that in this early wave of AI startups, technical expertise might matter more than domain knowledge.

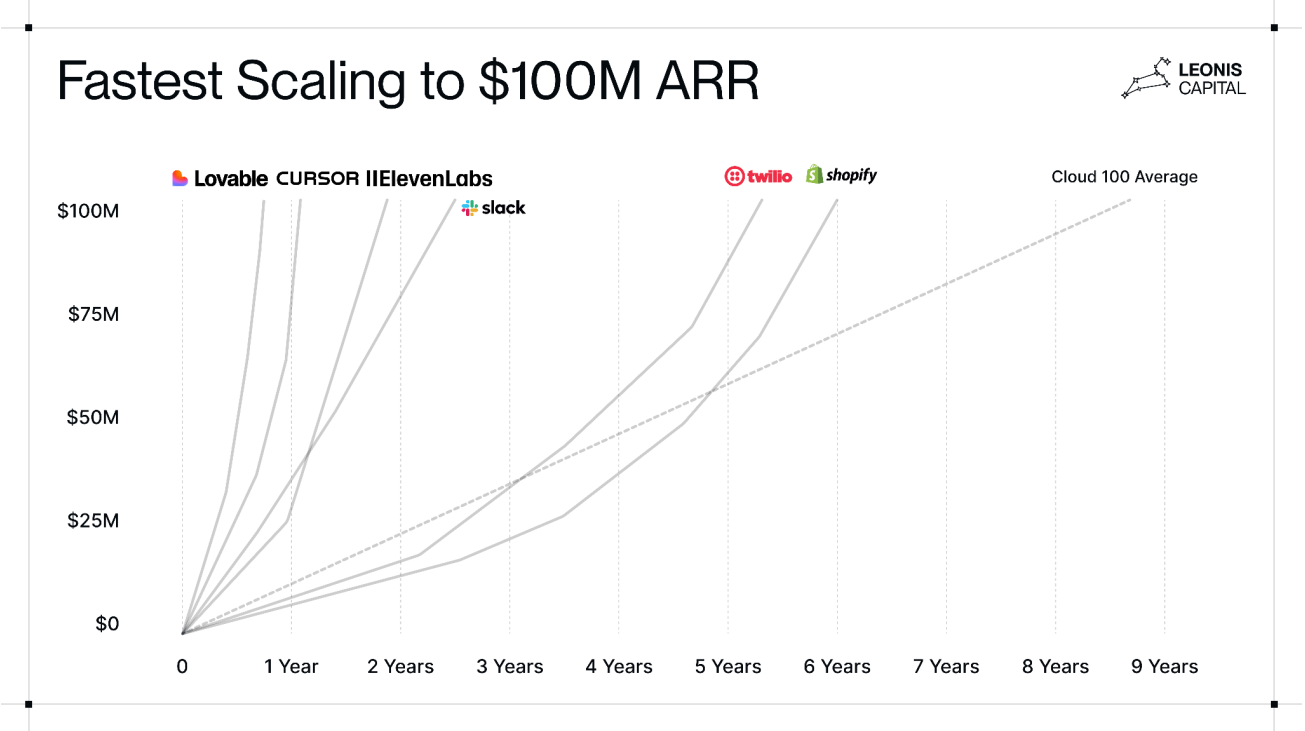
The importance of technical depth has produced a new archetype: **the "researcher founder."** Unlike the operators and domain experts who dominated the SaaS era, these founders are comfortable with technical uncertainty, quick to experiment, and able to spot inflection points before the market does. When ChatGPT's capabilities surprised even OpenAI, it was this group that immediately recognized which applications had suddenly become possible. Their rise also reshapes the timing of entrepreneurship: the best moment to start an AI company is often right after a breakthrough, when researcher founders can arbitrage their knowledge advantage before it diffuses. As the gap between research and product continues to shrink, this archetype is positioned to remain dominant and may well define the next generation of iconic companies.

Insight 2: The Post-2024 Revenue Surge

Many thought the AI funding bubble was going to pop in 2024, but that's not what happened.

Back in 2023 (and even into mid-2024), the prevailing narrative was that there was too much capital chasing too few real businesses. Sequoia Capital's David Cahn captured the mood with his famous "$200B question," warning of massive investment in AI infrastructure with minimal downstream revenue. By mid-2024, he escalated it to "AI's $600B question," painting an industry burning cash on R&D while customers remained on the sidelines. Those were all valid concerns at the time.

Then the revenue dam broke.[9]



After a year of viral demos and proof-of-concepts, late 2024 marked an inflection point. Suddenly, dozens of AI startups weren't just showing traction; they were achieving staggering revenue numbers in record time. Cursor hit $100M ARR in 12 months. ElevenLabs reached $100M in 22 months.[10] Multiple companies in our dataset went from zero to $10M+ ARR in under a year. For comparison, Slack, the poster child for hypergrowth in the SaaS era, took about 12 months to reach $10M ARR and 36 months to reach $100M.

We were less surprised than most. Even in early 2023, when skeptics dominated the discourse, we believed AI-first companies would scale revenue dramatically faster than their SaaS predecessors. Our reasoning was simple: **AI products deliver exponentially more value than static software.** A tool that writes code for you or manages an entire support queue is not a marginal feature; it replaces hours of skilled labor. **That labor-replacement logic makes**

[9] The dotted line comes from the median time to $100M based on Bessemer's Cloud 100 (2022-23 cohorts).

[10] The beta platform was launched in January 2023 and they reached $100M in October 2024.

the value proposition immediate and obvious: customers will pay quickly if the product either reduces headcount or multiplies the productivity of existing teams.

Counterintuitively, researcher-founder teams, despite their academic pedigrees, often moved faster on monetization than traditional SaaS founders. Their experimental mindset: rapid iteration, comfort with uncertainty, willingness to test hypotheses, translated directly into aggressive GTM and faster product-market fit.

But fast revenue growth doesn't always mean healthy economics. One of the biggest questions we and other investors grapple with is the gross margin profile of these companies. Despite their extraordinary growth, many still operate with poor or even negative margins. According to *The Information*, Lovable reports only 35% margins, StackBlitz 40%, and Replit just 23%. These numbers exclude nonpaying users, meaning true margins are likely lower. The reality is that most AI companies today sell products with cloud and inference costs that scale linearly with usage. Until inference efficiency improves or proprietary model access becomes cheaper, many "high-revenue" AI startups remain low-margin businesses. The tension between explosive top-line growth and fragile unit economics will define which of today's leaders can mature into sustainable, software-like businesses.

Not all reported revenue, however, is created equal. The space has developed its own creative accounting, with "ARR" used to describe everything from contracted subscriptions to one-time fees, cloud credits, and even signed LOIs. Some call this **"vibe revenue."** One example was 11x, which claimed $10M ARR based heavily on LOIs and verbal commitments. By 2025, sophisticated investors have grown far more cautious about ARR numbers and demand a clear separation between recurring revenue and one-time inflows.

The revenue surge, however, may partly reflect financial acceleration and customer experimentation rather than purely organic growth. Many AI startups raised unusually large rounds at premium valuations, giving them cash to scale usage and revenue faster than typical SaaS counterparts. Some of this capital is used to **subsidize model costs**, allowing companies to price below true compute and model expenses to accelerate adoption. At the same time, customers today are actively looking for AI solutions, often **experimenting with multiple vendors**. If we controlled for cash on hand and experimentation effects, the gap between AI and SaaS revenue velocity would likely narrow but not disappear. The fundamental driver still appears to be product value: when a product automates or replaces skilled labor directly, conversion happens almost immediately.

Even accounting for exaggerations, the AI revenue numbers are extraordinary. The real test comes after generating revenue: retention. Explosive growth is exciting, but net revenue retention (NRR) and logo churn will determine which of these companies build enduring businesses. Many early AI products solve an initial problem brilliantly but fail to embed themselves in daily workflows, while others face the constant threat of being absorbed by foundation models. Today's $10M ARR AI writing tool could be tomorrow's free ChatGPT feature. Still, the revenue surge of 2024-2025 definitively answered one question: customers will pay for AI, and they'll pay quickly and handsomely.

Insight 3: Smaller and Flatter Teams, Huge Output

One of the most striking features of this wave of AI companies is how much they achieve with such lean teams. Where the last generation of SaaS unicorns often scaled headcount in parallel with revenue, today's AI startups are rewriting the relationship between team size and output. Their leverage doesn't just come from capital efficiency; it comes from using AI internally to build, sell, and support their products.



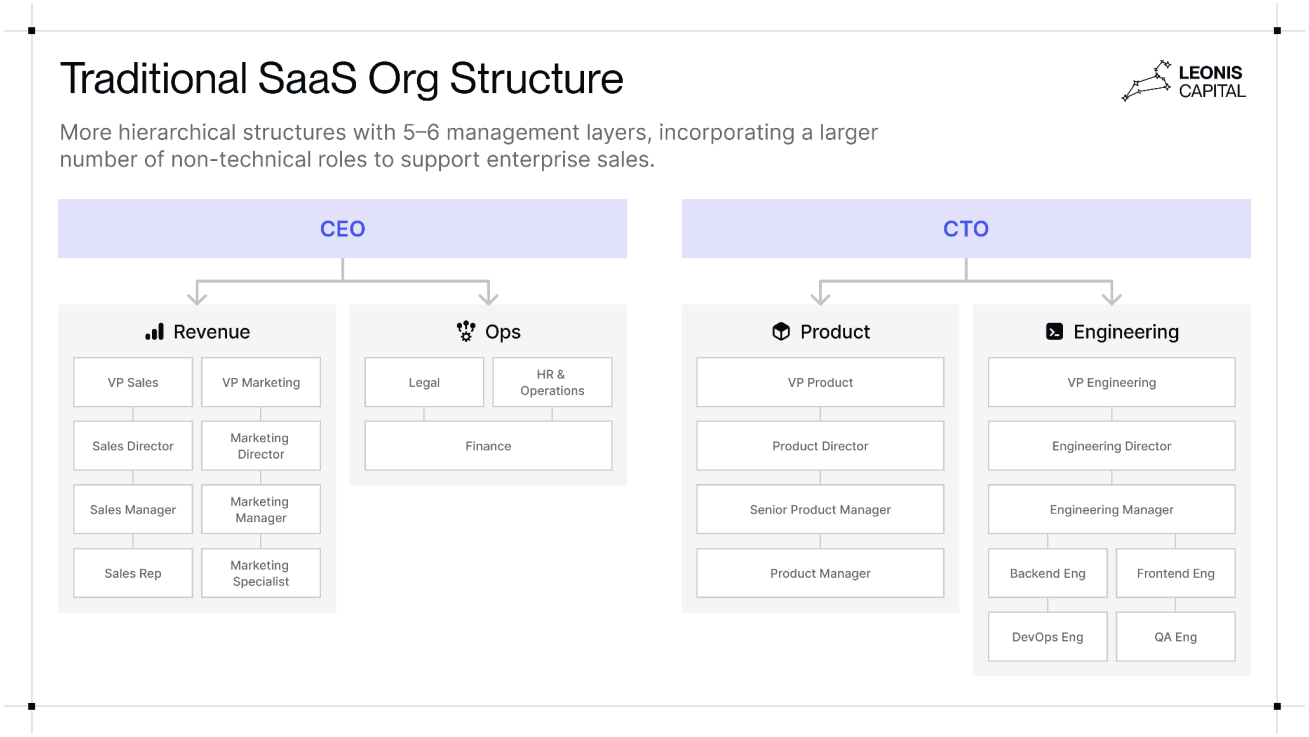
The numbers are staggering. In our dataset, several companies crossed $1M ARR with fewer than five employees. One company reached $5M ARR with a team of five; another scaled to $50M with a team of just 28. Midjourney reportedly reached ~$200M ARR in 2023 with only 40 people ($5M per employee), and Lovable scaled to about $100M ARR with 45 people ($2.2M per employee).[11] These figures dwarf even top-tier, pre-IPO SaaS companies, which averaged ~$300K per employee, showing AI companies can achieve 3-10x more revenue per employee.[12]

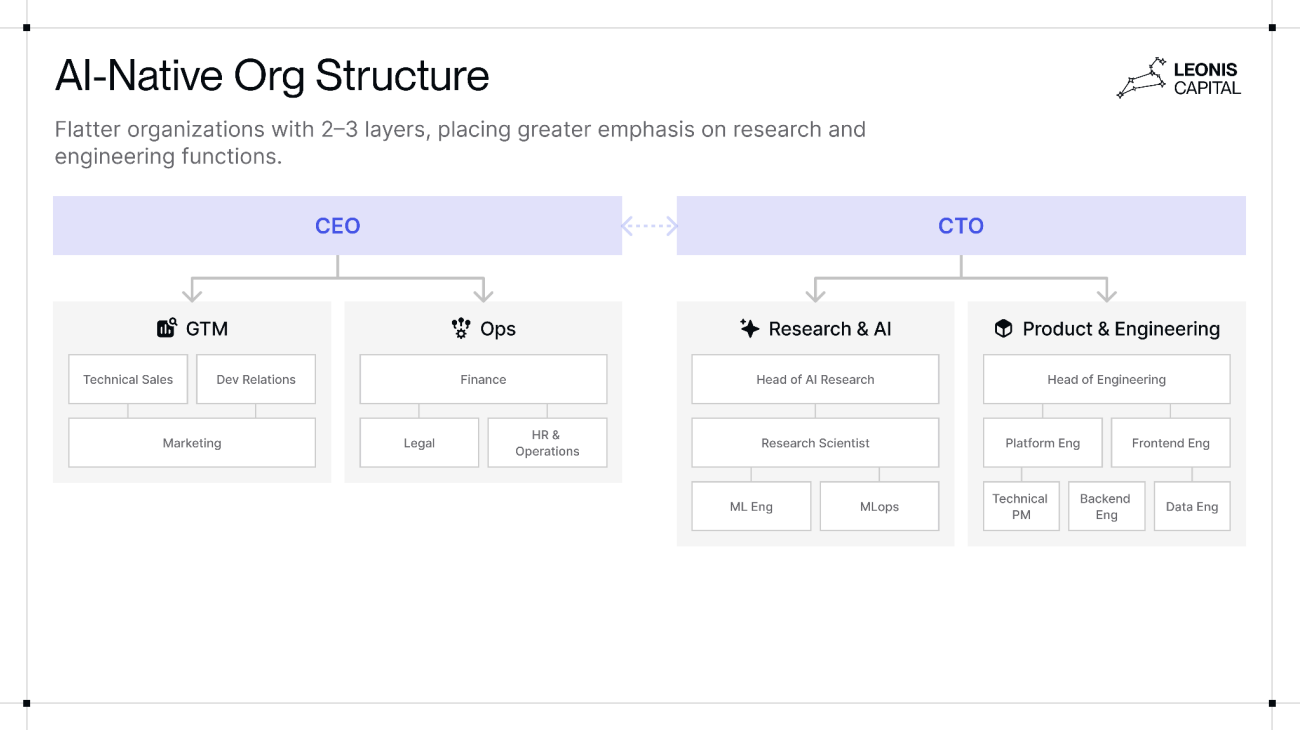
[11] Data for the revenue per employee illustration: *TechCrunch* (Midjourney), *TechCrunch* (Lovable Series A), *TechCrunch* (Lovable ARR Milestone), *TechCrunch* (Replit), *Bloomberg* (Perplexity), PitchBook profile, RIAA / MBW article. For headcount data, we used Built In NYC for Runway and LinkedIn Sales Navigator for the others. We discounted LinkedIn headcount due to contractors, investors and ex-employees often being counted in erroneously. Other revenue data from Sacra's estimates.

[12] There is a caveat to the revenue per employee metric though. Many AI-native companies spend more on specialized agencies, fractional contributors, and software but none of these expenses show up in the revenue per employee's denominator. This could lead to an illusion of higher capital efficiency. However, since traditional SaaS companies also use services that are not reflected in the payroll, we believe that the effect of these services is relatively small.

This isn't just efficiency; it's **a structural shift of how startups are built**. Traditional SaaS organizations often developed several layers of specialization as they scaled: product managers mediating between engineers and customers, SDRs driving top-of-funnel, and customer success teams managing adoption. AI companies are collapsing many of these layers. Engineers build systems that build features. AI agents handle most support queries. The same person who architects the system often demos it to customers and ships code to production. GTM happens through developer relations and technical sales instead of large SDR teams. Even business staff are more technical, leaning on AI tools to bridge skill gaps. The result is a noticeably flatter, more research-heavy structure, typically with 2–3 layers at the early stages, where every function orbits the technical core.

This collapse of boundaries extends to the executive level as well. In many AI-native startups, **the traditional CEO-CTO divide is blurring**. Because so many CEOs are deeply technical, they often lead product development, write or review code, and shape engineering direction directly. The CTO role, in turn, leans more toward research and infrastructure leadership rather than being the sole technical authority. The result is a founder dynamic that feels closer to a two-engineer team scaling an idea than a classic business–technical pairing.





The lean headcount of AI companies does not mean they are capital-light. The difference is where the money goes. In the SaaS era, new funding rounds translated directly into headcount: armies of sales reps, engineers, and product managers. In AI, capital is deployed less toward people and more toward infrastructure: GPUs, inference costs, and data licensing. These companies are labor-light but compute and data-heavy. In other words, where SaaS companies turns dollars into people, AI companies turns dollars into compute and data.

This headcount efficiency is also aided by **product standardization**. Many AI startups deliver largely uniform products across customers: an AI coding agent or image generator typically looks the same for every user, perhaps linking to a different codebase or dataset but requiring little customization. By contrast, traditional enterprise SaaS products often demand heavy integration with each client's existing infrastructure and data systems. The relative uniformity of AI-native products reduces implementation complexity and customer-specific engineering effort, enabling high revenue per employee.

It's worth noting that **not every company in the AI 100 achieves this level of revenue per employee**. The outliers like Midjourney, Lovable, and Cursor share common traits: horizontal products with broad appeal, delivered as pure software with minimal services, and monetized through self-serve or viral adoption. Their business models allow infrastructure spend to scale directly into revenue without large customer-facing teams. By contrast, vertical AI companies, particularly in healthcare and law, often require professional services, integration, and compliance support.[13] This tends to push revenue per employee closer to SaaS norms, but it is balanced by the benefits of enterprise offerings: higher margins, longer contract durations, and more defensible customer relationships.

[13] For example, as of 2025, Harvey has an estimated revenue of $75M and 340 employees, which translates into about $220K in revenue per employee.

Early data suggests the influence is also running the other way, and the lean model pioneered by AI-native companies is spreading across SaaS. ICONIQ's *State of Software 2025* finds ARR per FTE up 11% year-over-year, growing more than 2x faster than OpEx per FTE, evidence of rising headcount leverage as AI tooling scales across teams.[14] In short, SaaS and AI may be converging, with the leaner AI-native mode becoming an operating standard.

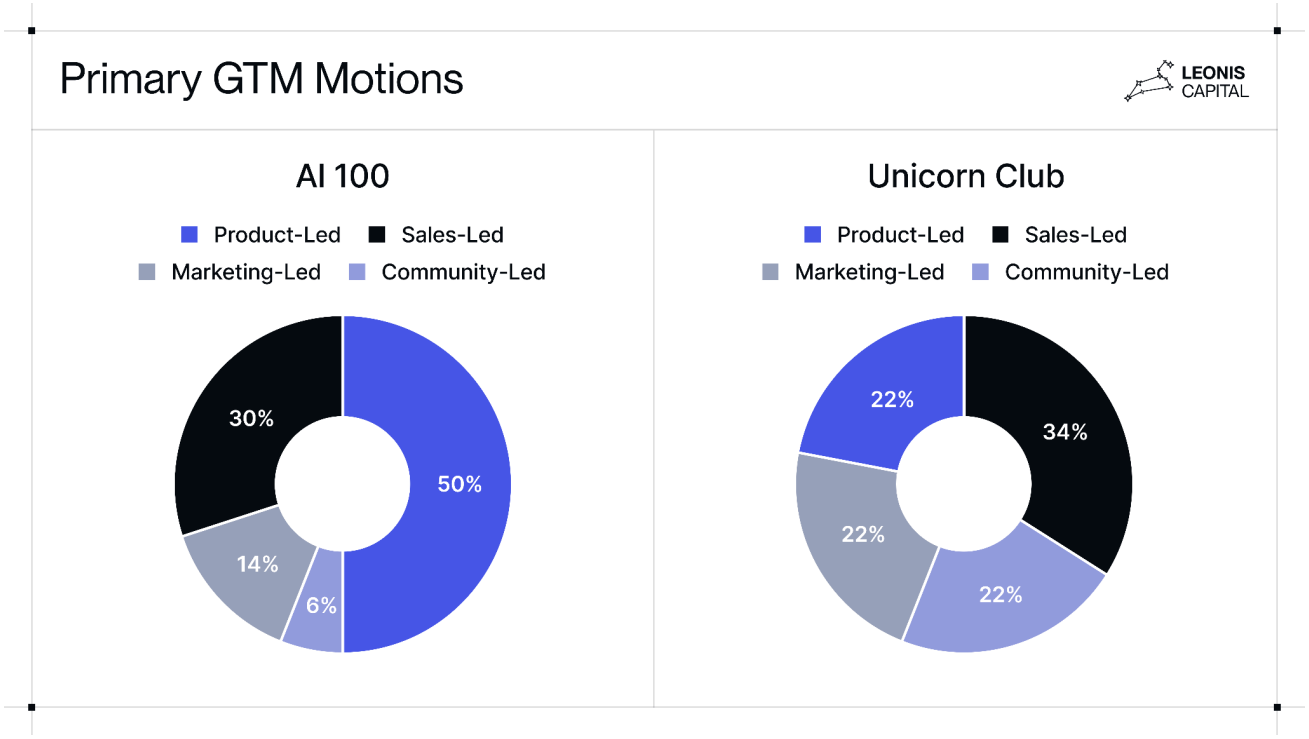
An open question is **whether this model can scale**. Today's AI companies achieve extraordinary leverage by automating work internally and compressing roles, but that same simplicity may limit how far they can grow. As customers demand deeper integrations, regulatory compliance, and enterprise support, even the leanest teams may need to add more structure and hire more across these teams. The tension ahead is clear: can AI-native companies preserve their ~$1M-per-employee efficiency as they expand, or will they gradually adopt the hierarchies of the SaaS era they replaced?

[14] The benchmark is derived from a moderate-sized cohort, 125+ enterprise software companies, with multi-quarter financial & operating data, and ARR > $10M. In addition to AI tooling, offshoring, and org changes are cited as possible factors. But it should be noted that layoffs have anecdotally been enabled by AI.

Insight 4: PLG First, Sales Later

The death of enterprise sales has been predicted before, and in the age of AI, many believe it's finally here.[15] After all, products like Cursor and Granola seem to sell themselves.

Across our top 100 AI startups, product-led growth (PLG) dominates early traction. Over 80% launched with self-serve onboarding, showing how much reach a small team can achieve through pure product appeal. According to our data, enterprise sales hasn't disappeared, it simply moved downstream. The sales motion now happens after adoption, formalizing demand that already exists rather than creating it from scratch.



Sales-led motions still matter, especially in vertical and complex domains. Companies like Harvey in legal tech and Abridge in healthcare rely on enterprise sales from day one because their customers demand integrations, data security, and compliance.

For horizontal tools, the pattern is different. Developers adopt a coding assistant like Cursor individually; within weeks, it spreads across the team; months later, sales steps in to manage procurement and pricing. Many founders didn't hire any sales team until well after reaching significant ARR. The sales motion follows usage, not the other way around. This inverted funnel has become the standard AI GTM sequence: **PLG drives awareness and proof of value, then sales consolidates and scales it**.

In the SaaS era that preceded today's AI cycle, sales-led growth still dominated across companies, despite breakout PLG stories like Slack. Among B2B SaaS companies with more

[15] Analysts made similar predictions during earlier software shifts: In the SaaS wave in the early 2000s, tools like Salesforce and Zendesk were said to make enterprise sales obsolete by lowering the friction for adoption. In the mid-2010s, the PLG wave saw the rise of Slack, Dropbox, and Atlassian, showing that products can spread virally without sales teams. In the developer-led adoption wave of 2019-2022, the rapid growth of APIs from Stripe, Twilio, and Datadog were cited as evidence that bottom-up demand would replace traditional sales motions.

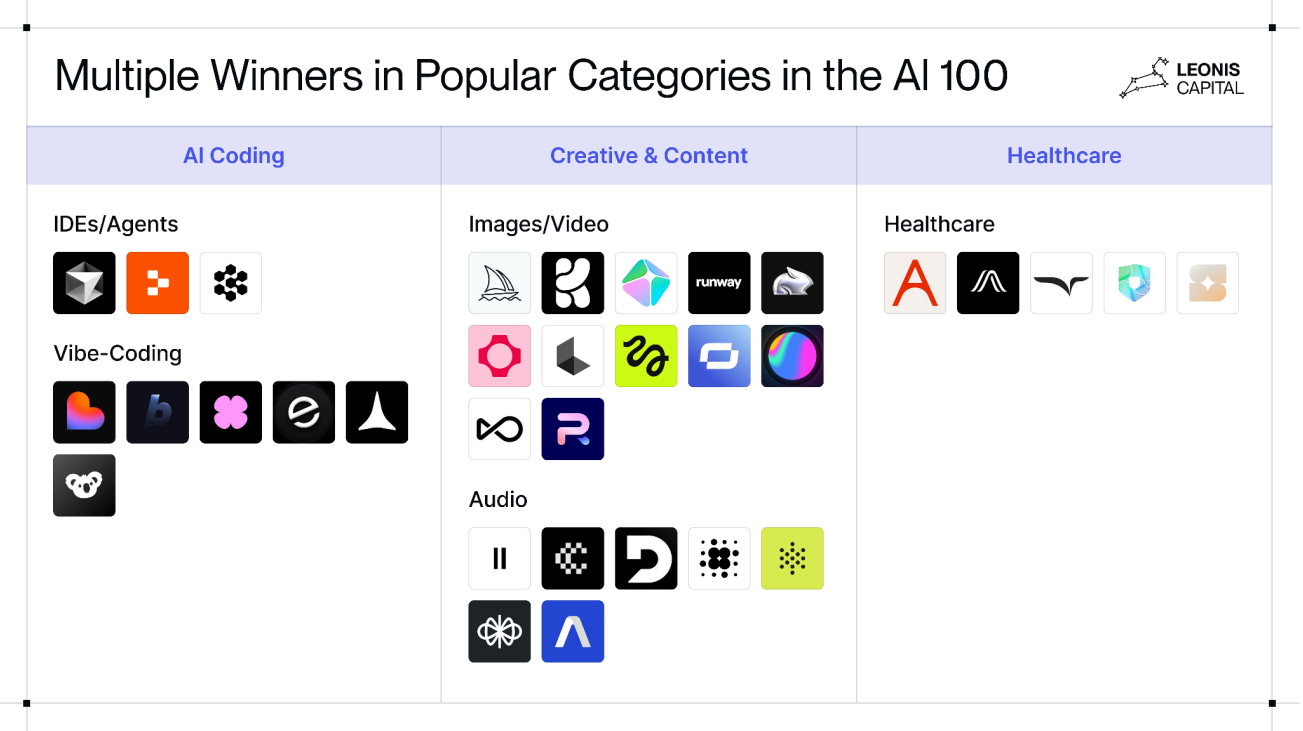
than five employees in 2022, 71% were sales-led and just 29% product-led, according to a Notion Capital dataset of 30,000 companies. The 2013 Unicorn Club reflects this dynamic: more than half of the top startups leaned heavily on sales or marketing-led motions, with product-led growth accounting for only about one in five.

Today, that PLG playbook born in the SaaS era applies more universally. Slack, Atlassian, and Dropbox showed that virality and usability can replace cold outreach, but AI compresses the cycle. Time-to-value has dropped to seconds, APIs enable instant integration, and open-source diffusion creates demand faster than any outbound campaign.

In short, enterprise sales isn't dying, it's being redeployed. In the new order of AI distribution, the product gets in the door, and sales builds the contract around it.

Insight 5: Multiple Winners, Not Winner-Takes-All (Yet)

In the current AI boom, many sectors are seeing multiple successful companies flourish side by side. This runs counter to the pattern of past tech waves like search and social media, where a single platform often dominated. It seems that **AI has become a rising tide that lifts all boats.** A huge surge in demand and investment is enabling numerous AI startups in the *same* segments to prosper simultaneously, at least in the beginning.



In coding, products like Replit's Ghostwriter, Cursor's IDE, and Cognition's Devin have each built loyal followings through unique features: Replit for cloud-based collaboration, Cursor for local speed and control, and Devin for full agent autonomy and end-to-end workflows. Developers adopt different tools depending on their environment and needs, and often use several simultaneously.

Furthermore, the AI coding market is splitting into two distinct categories: products for **engineers** who want IDE integration, version control, and code reliability, and products for **vibe-coders** who want low barriers and rapid prototyping. The fact that Lovable and v0 can coexist with Cursor and Replit illustrates the fact that the AI coding market is both massive and diverse. The sheer scale means there's room for multiple players to grow without cannibalizing each other's success.

Similarly, the creative and content domains, spanning image, video, and audio generation, have also seen a Cambrian explosion of tools, with many early winners across different creative tasks.

The demand for AI-generated content is so widespread that no single service can satisfy all use cases or audiences. Consider AI image generation. Even though Stability AI and

Midjourney laid the groundwork, newer platforms still emerged and thrived. Many creators use Krea or OpenArt for generating and editing images. The open-source nature of image models like Stable Diffusion has spawned dozens of specialized tools, from AI design assistants to image-upscaling services.

In AI video generation, the market is so "crowded" that multiple competing companies buy ads against "Synthesia" as a search term. HeyGen, for example, offers an alternative avatar video platform that's also popular in the same enterprise training market. But both companies have achieved significant traction, with Synthesia reaching $100M ARR in March 2025 and 60,000 business customers, while HeyGen generated more than $100M in ARR.

AI voice is another example of this multi-winner dynamic. There are three voice AI providers on our AI 100 list: ElevenLabs achieved a $6.6B valuation, focusing on high-quality voice synthesis, while Cartesia stands out for ultra-low latency real-time conversations. Meanwhile, Deepgram built a $21.8M revenue business specializing in enterprise speech-to-text. The lock-in effects are minimal, and customers routinely combine multiple providers to optimize different requirements.

Healthcare is another domain where we're seeing multiple concurrent winners, largely because healthcare itself encompasses varied, complex needs. Abridge and Freed AI both target the real-time AI scribe market, but they've each charted distinct paths to scale. Abridge reached a $5.3B valuation by deploying across 150+ enterprise health systems, while Freed AI adopted a lean direct-to-clinician SaaS model, converting 20,000 clinicians and achieving $19M in ARR in two years. Their coexistence underscores how even narrowly defined workflows can fragment into multiple viable sub-markets, depending on user type, integration depth, and GTM motion.

Several structural factors drive this multi-winner landscape. AI's wide range of use cases means companies can carve out distinct niches, each requiring specialized expertise and domain-specific data that few competitors can easily replicate. Unprecedented levels of funding ($131.5B in AI startup investment in 2024 alone) combined with off-the-shelf models have lowered barriers to entry while enabling rapid development and experimentation. The result is an ecosystem where companies complement rather than cannibalize each other's success.

However, it's worth considering whether this represents a fundamental shift in how AI markets operate or **simply reflects the industry's early stage before inevitable consolidation**. As market dynamics mature and capital becomes more selective, the better-funded and strategically positioned companies will likely outcompete and absorb their competitors. Arguably, this is already happening: Cursor is reportedly generating more than three times Cognition AI's revenue, and Abridge has established a substantial lead over Freed AI in enterprise adoption. Also, companies whose main product is the agent itself, rather than the surrounding workflow or vertical layer, face a harder long-term position, as they are the most direct targets for foundation models to absorb. For now, though, AI appears to be creating expanding opportunities that reward innovation and specialization across multiple parallel winners.

Insight 6: Winners Pivot Fast, AI Makes it Possible

Many successful AI companies in our dataset pivoted dramatically in its first year. Not minor adjustments or feature tweaks, but complete rebuilds of their core product.

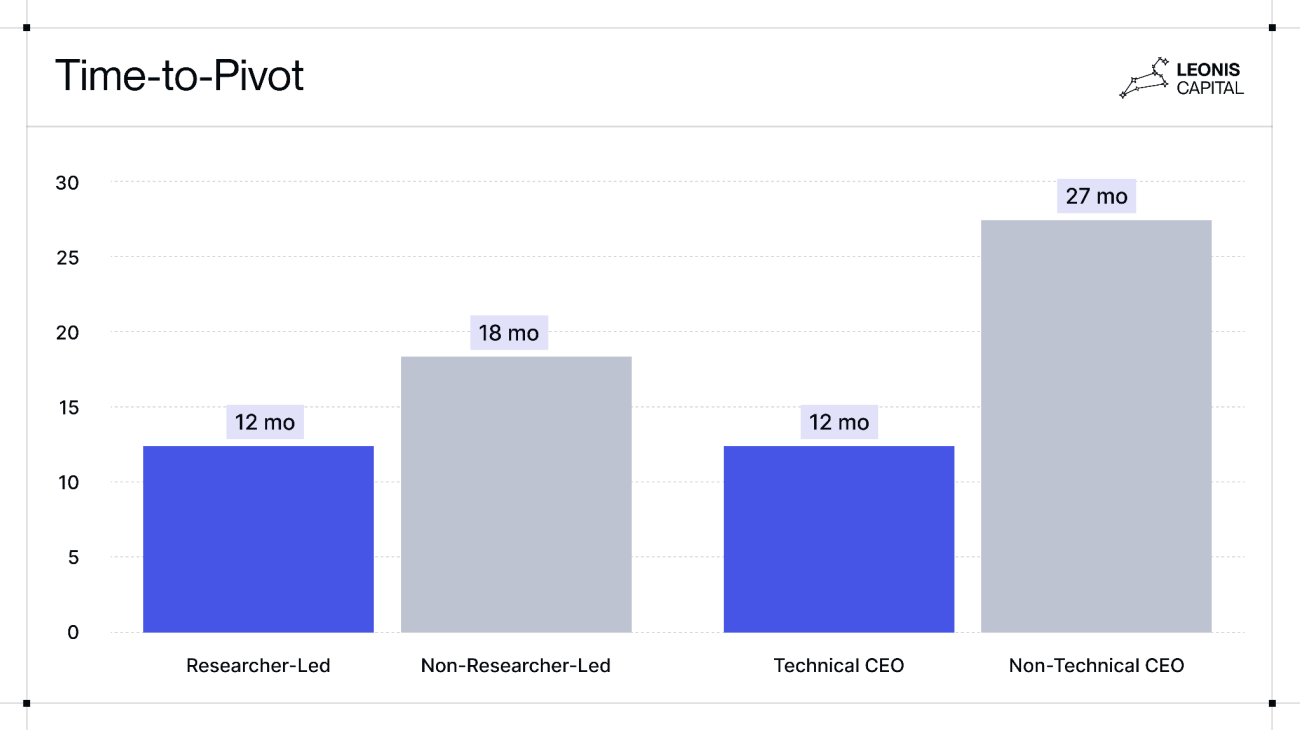
In the pre-AI era, pivots looked very different. Even successful ones unfolded over years, not months. Twitter, which started as a podcasting app named Odeo in 2005, took roughly 12 months to pivot, which was considered fast at the time. Coinbase spent two years as a Bitcoin wallet before becoming an exchange. Slack took four years to pivot from gaming company Tiny Speck to enterprise messaging. Likewise, Lyft, Notion, and Twitch all took four or more years to abandon their original ideas. These timelines made sense in a world where product direction was constrained by customer feedback cycles, infrastructure build-out, and slower iteration loops.

Fast Pivots Behind the AI 100

LEONIS CAPITAL

Company	Original Idea	Final Product	Time to Pivot
CURSOR	AI-powered CAD design for mechanical engineers	AI coding assistant	~9 months
Lovable	Open-source prompt-to-full-stack-app generation tool (GPT Engineer)	AI vibe coding platform	~3 months
Windsurf	Infra for managing large-scale GPU operations	AI coding assistant	~10 months
manus	AI browser plugin	General-purpose AI agent	~12 months
Retell AI	Infra for testing voice AI models	AI agents for call centers	~3 months

Conventional wisdom suggests that researchers and technical founders would be wedded to their original vision. Our data shows the opposite: **Two-thirds of AI 100 startups (66%) pivoted at least once**, compared to only half (54%) in the Unicorn Club. Researcher-led teams pivoted within a median of **12 months**, versus 18 months for non-researcher teams. **Technical CEOs tend to lead fast pivots** (12 months), while companies with non-technical CEOs took more than twice as long (27 months).



Why? Because they live at the model frontier. They read release reports, benchmarks, and arXiv papers, so that when model capability jumps appear, they see them first. Their advantage isn't just market intuition but technical foresight.

This sensitivity to new model developments often triggers the earliest and sharpest pivots. Take Manus as an example. The company started as an AI browser plugin called Monica in 2022. But after Chief Scientist Peak Ji joined a year later, they quickly recognized that frontier models could reason, plan, and execute across tasks. Within months, the team shifted focus to building a platform to orchestrate these emerging capabilities and eventually a generalist AI agent that surprised the tech world.

The Cursor story is similar. The founders began by building AI-powered CAD software for mechanical engineers, a technically ambitious project requiring spatial reasoning and 3D model generation. But they ran into fundamental limits: insufficient training data, models that couldn't reason about physical constraints, and a market that wasn't ready. When the team gained early access to GPT-4, within minutes of testing it, they realized it was shockingly good at coding. They made the call to abandon months of work on their CAD product and rebuild from scratch as a coding assistant.

Another interesting trend within the AI 100 is that **many teams started with infrastructure plays but later pivoted into applications**. It made sense since technical founders start out building what they know and usually for users like themselves. But as foundation models advanced, they saw two things converging: New capabilities making many infrastructure layers redundant and value starting to accrue to applications that controlled the data, distribution, and the user experience.

So they moved up the stack. Windsurf's journey captures this perfectly. The team started by building Codium, an infrastructure layer for managing large-scale GPU operations. But as AI infrastructure matured, they realized that GPU management was becoming a commodity

problem. Meanwhile, a much bigger opportunity was emerging as AI coding advanced. The team quickly pivoted up the stack, rolling out an AI coding assistant that went on to reach more than 1 million users.

The stories of Manus, Cursor, and Windsurf are typical of this generation of AI startups, where **new model capabilities trigger technical pivots**. In the SaaS era, pivots are driven by market learning and customer demand. You build an MVP, talk to users, iterate on features until you find product-market fit. Slack, for instance, started as a gaming company's internal communication tool before they realized the chat product was more valuable than the game itself. But the underlying technology largely stayed the same. AI pivots are different. They're often driven not just by customer feedback but by model capabilities as well. New model releases can suddenly make or break an entire product category.

AI also makes pivots unusually fast and cheap. Because most applications are built on the same foundation models, reinventing a product rarely requires rebuilding the stack. The same APIs and orchestration frameworks can be re-combined into an entirely new product in days. AI agents across different domains share the same underlying components: evaluation, logging, routing, and orchestration infrastructure. Once that layer is built, switching use cases becomes relatively easy, since the difference lies more in prompting and data context than in core system design.

This fluidity extends to teams. Technical talent in AI is far more portable than business talent in traditional SaaS, where domain specialization locks people into verticals. For example, in e-commerce, success depends on deep domain context and understanding of supply chains and logistics, so skills don't transfer easily. In contrast, in AI, a technical team fluent in agent architecture, retrieval, and fine-tuning can jump from voice AI to code generation without hiring new expertise.

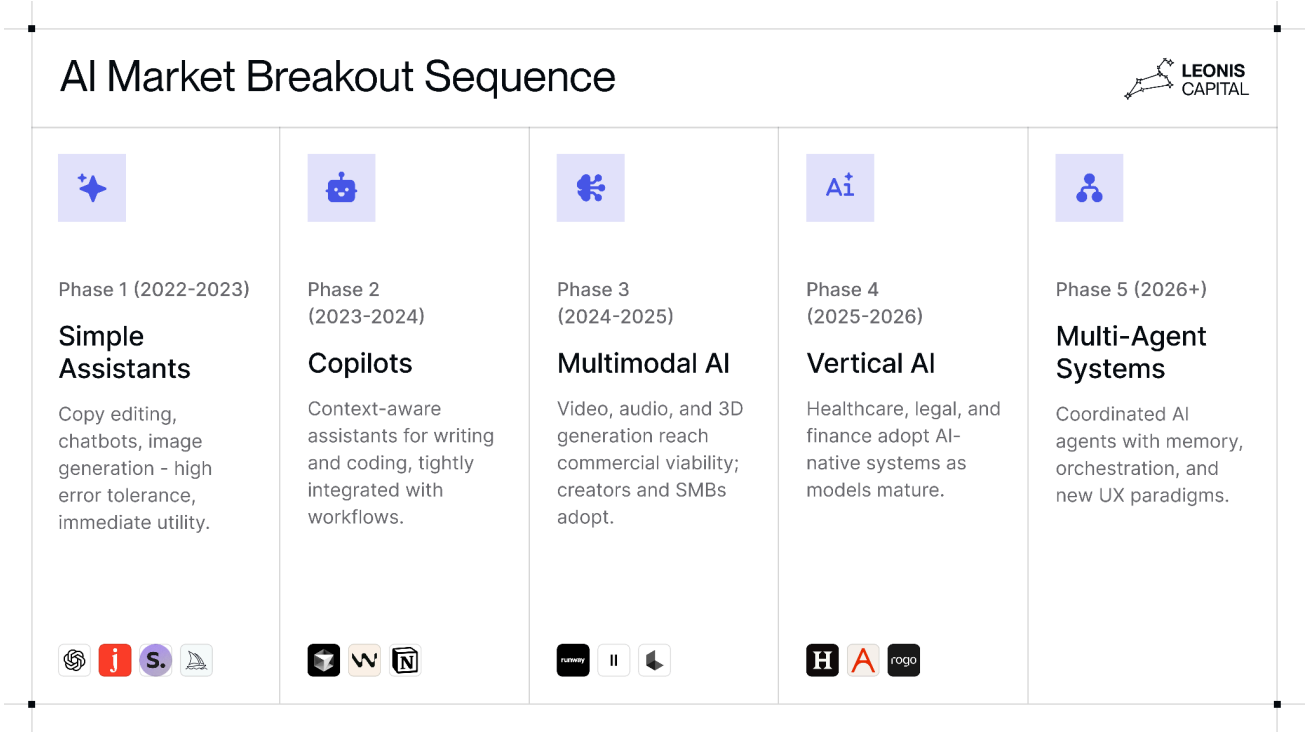
But proximity to frontier models cuts both ways. The same dependency that enables speed also threatens survival. As foundation models absorb functionalities like coding and agentic workflows, entire product categories collapse into base-model features. Many companies pivot not by choice but out of necessity, trying to escape commoditization before it reaches them.

The best founders anticipate this cycle and pivot ahead of the curve. They predict where capabilities will mature next and build for what the next generation of models will unlock. They position themselves in a way where new model releases expand their product's power rather than replace it. That's how model volatility becomes leverage instead of risk.

Insight 7: Markets Break Out in Sequence

AI markets don't break out randomly. They unlock when model performance crosses a hard capability threshold. Understanding where those thresholds sit and when they're about to shift is now one of the most important skills for founders and VCs. Enter too early, and the technology can't deliver; enter too late, and the advantage disappears.

Here's how the dominoes have fallen:



Early breakouts like **writing and coding** were not accidents. These were domains most optimized by foundation-model labs from the start because they had massive structured training data, clear task definitions, and a large user base of early adopters. Individual use cases generally accelerate early adoption and developers are among the fastest to experiment with new tools.

As models matured, **verticals such as healthcare, legal, and finance** began to unlock. These sectors benefited because foundation models improved reasoning, retrieval, and long-context understanding capabilities. Some were even fine-tuned on vertical data and tasks. With these improvements, vertical AI agents finally became viable, leading to the rise of startups like Rogo in finance and Abridge in healthcare. This shift mirrors the SaaS evolution, where horizontal tools emerged first but software later specialized by industry.

The critical pattern throughout the LLM era is that **once a performance threshold is crossed, entire new applications appear**. When Claude 3.5 dramatically improved coding reliability (handling multi-step reasoning, debugging, and even end-to-end app generation), an entire wave of "vibe-coding" startups appeared within weeks. Earlier, Stable Diffusion catalyzed a number of image-generation companies because it made high-quality synthesis accessible.

And now with OpenAI's latest Sora release, we may be seeing an equivalent inflection point in video-first startups.

For founders, the critical insight is that **timing matters as much as execution**. Entering too early, before the tech is mature, erases even the best execution. We saw this firsthand: many AI-coding startups before 2021 struggled because the underlying models just weren't good enough. The opposite mistake is entering after the inflection point, when the window has already closed. By mid-2024, markets like AI note-taking and chat-based customer support had become saturated; dozens of nearly identical copilots were competing on UX polish instead of underlying differentiation. These founders weren't wrong about demand, but they were just late to the capability curve. The best timing sits in between, close enough to the inflection point to deliver real value, but early enough to compound before the market crowds in.

The best founders often appear slightly ahead of their time. Harvey started working on their product in 2022 when they saw the potential for LLMs to handle legal tasks. While early versions didn't perform well because of model limitations, the company gained credibility and mindshare among law firms. When model capabilities caught up, usage began to scale rapidly. Similarly, Freed AI began working on AI medical note-taking in early 2023, even before the release of GPT-4. Its usage accelerated after model quality improved, validating the team's early conviction. The top founders today are probably already building for Phase 5 markets, betting that they can benefit from improving model capabilities and multi-agent systems. They're trading current difficulty for future dominance.

As the AI market matures, a clear pattern emerges: **moats deepen as products move from horizontal to vertical**. Horizontal tools compete directly with foundation models and face constant feature absorption. Vertical AI applications, by contrast, gain durability from depth by embedding domain expertise, regulatory compliance, and proprietary data into their workflows in ways general models can't easily copy.

Each leap in model capability doesn't just unlock new categories, it redefines where advantage comes from. In the early phases, success came from speed and distribution, whereas in later phases, it comes from precision, workflow depth, and timing. The next decade of AI will likely be defined by founders who build just ahead of model capability shifts, in verticals where reliability and data depth truly matter.

Who's Backing the AI 100

The AI gold rush has pulled in capital from every corner of venture capital: tier-one firms writing large checks, solo GPs running micro-funds, strategic investors, accelerators, and a new wave of technical operator-investors.

At the Pre-Seed and Seed stage, Y Combinator leads by a wide margin, backing 21 companies, more than 20% of the entire AI 100.[16] YC remains the leading accelerator for founders building in the AI space.

Platform funds are also moving aggressively upstream, with a16z and Sequoia backing several companies at Seed or Pre-Seed. This marks a departure from their traditional focus on Series A and beyond. In AI, winners often reveal themselves earlier, forcing even multi-billion-dollar funds to compete for Seed round allocations.[17]

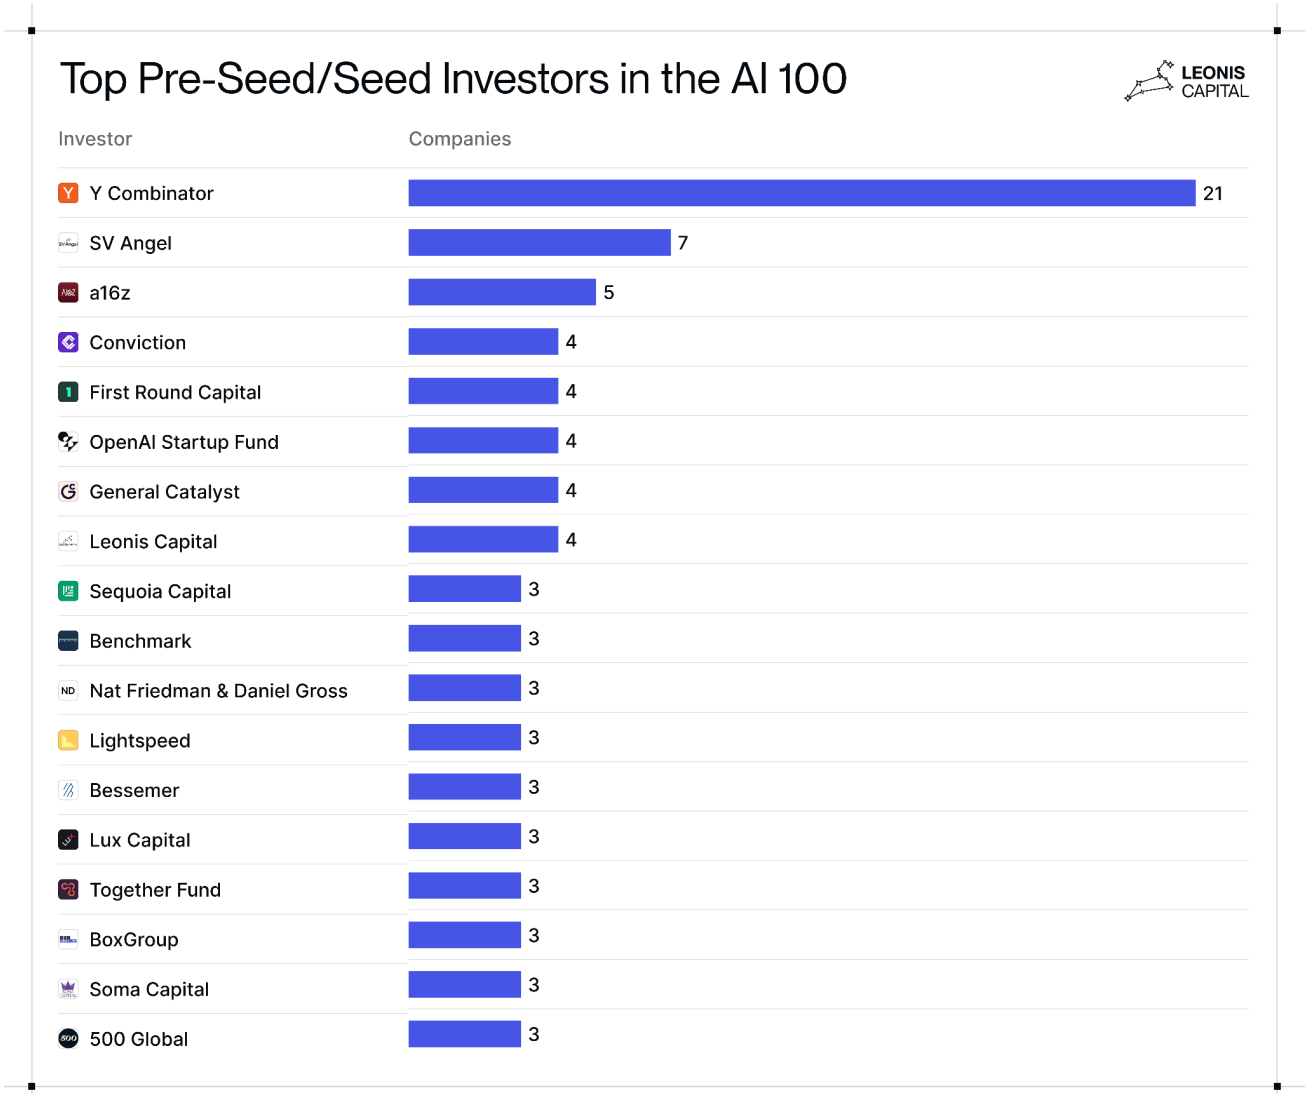
Angel networks like SV Angel were also fairly active (7), while First Round Capital (5) shows that established Seed funds still get into some of the best deals. Meanwhile, a new class of AI-native investors, such as Conviction (4) and the duo of Nat Friedman and Daniel Gross, has quickly become founders' preferred partners for their brand, technical judgment, and decision speed.

The OpenAI Startup Fund deserves special mention. Although strategic investors typically enter at later stages, they have quietly become one of the most credible early-stage investors in the AI ecosystem, leading the Pre-Seed or Seed rounds of four heavy-hitters: Cursor, Harvey, Cognition Labs, and UnifyGTM. OpenAI's deep technical expertise has become a core part of the fund's value proposition, turning its venture arm into a magnet for top technical talent.

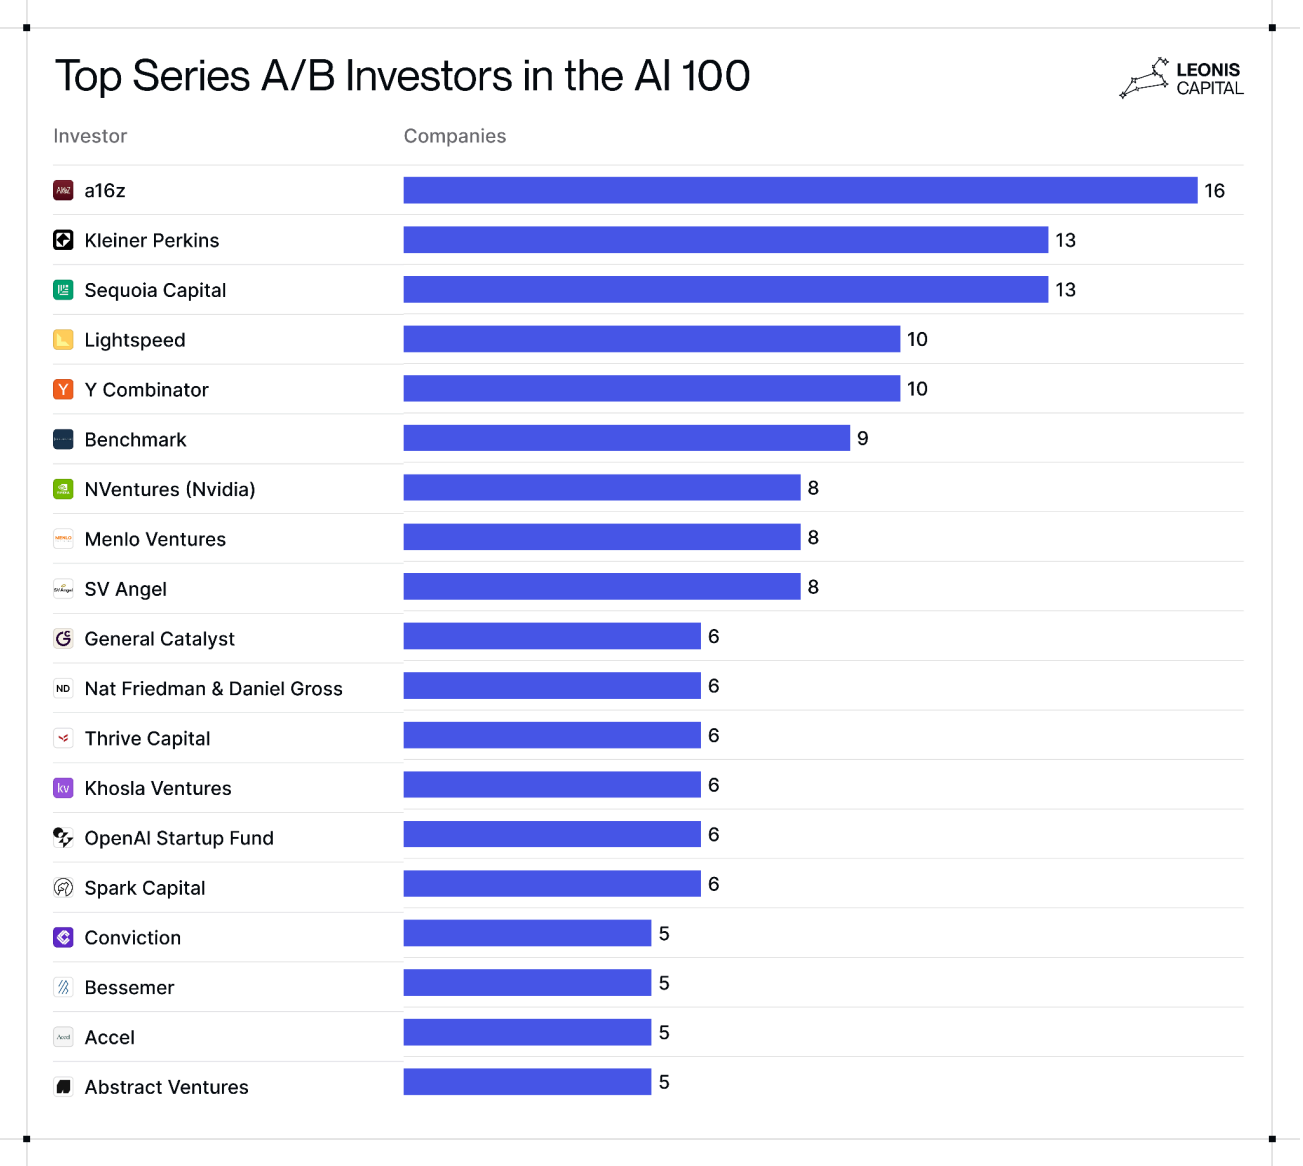
Beyond the major names on the list, there is a long tail of 177 Seed and Pre-Seed funds that backed one or two companies on the AI 100 list. This high degree of fragmentation suggests that despite the entrance of platform funds, there's still ample opportunity for emerging managers who can provide differentiated value.

[16] Investment counting methodology: Our analysis applies within-stage deduplication and cross-stage double-counting rules. If an investor participates in both Pre-Seed and Seed rounds for the same company, they are counted only once in the Pre-Seed/Seed category; similarly, participation in both Series A and Series B rounds counts as one investment in the Series A/B category. However, if an investor participates in both early-stage (Pre-Seed/Seed) and later-stage (Series A/B) rounds for the same company, they are counted in both stage categories. Therefore, the 195 unique Pre-Seed/Seed investors and 280 unique Series A/B investors represent overlapping rather than mutually exclusive sets, with the actual total unique investor count being significantly less than 351 due to follow-on investment behavior.

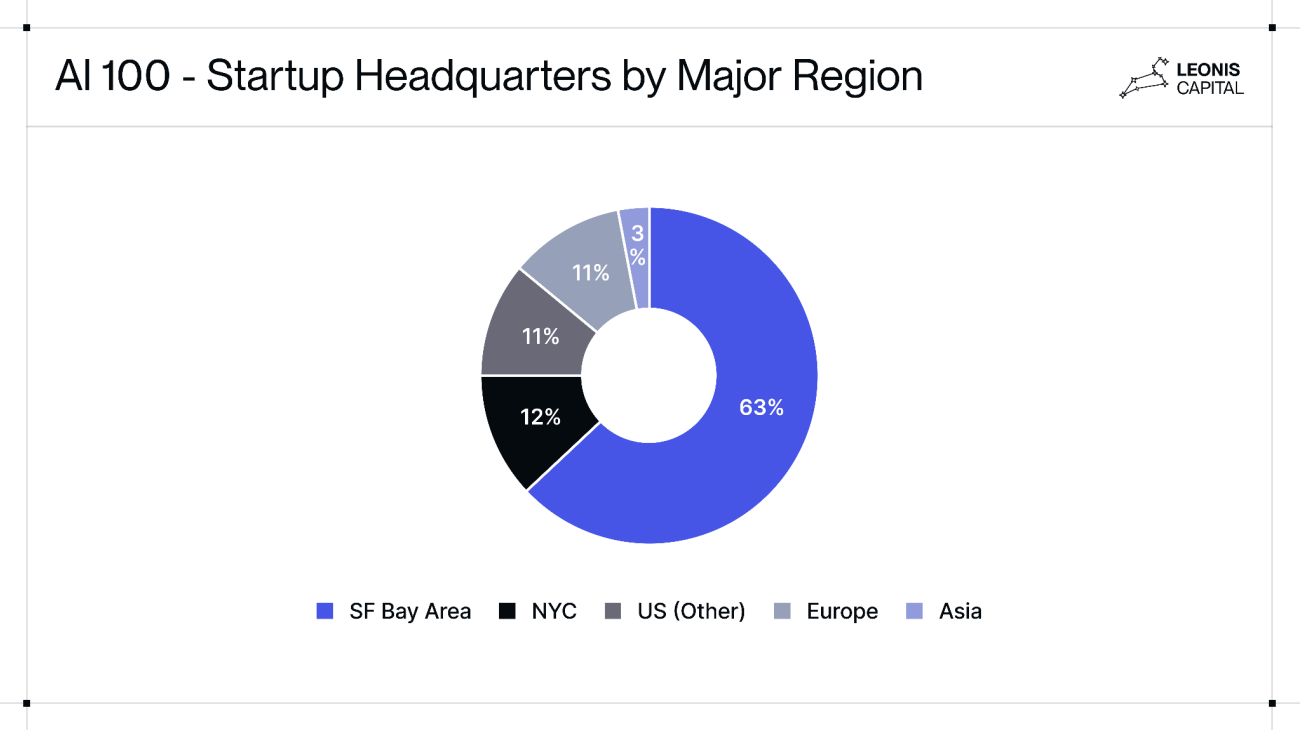
[17] Another reason for platform funds to invest earlier is that as they grow larger, they face increasing deployment pressure and have begun treating Seed investments as strategic "loss leaders" to secure access to later rounds. Bill Gurley is particularly critical of this dynamic.



As companies mature to Series A and B, the investor landscape consolidates quite a bit, with the top funds dominating the list: a16z leads with 16 investments, followed by Kleiner Perkins (13), Sequoia (13), Lightspeed (10), Benchmark (9), and Menlo Ventures (8). These firms together account for a large share of the Series A/B rounds, especially in lead or co-lead positions. At these later stages, brand, platform strength, and capital scale become a decisive advantage.



Strategic investors are also well represented in the Series A/B stages, but are playing a more targeted game. NVentures, Nvidia's venture arm, has backed 8 companies in the AI 100, concentrating on infrastructure plays such as Together AI and Fireworks AI, areas where GPU efficiency directly ties to Nvidia's business. At the Series A and B, the OpenAI Startup Fund continues to double down on its Seed winners while also leading rounds in application-layer companies like Speak AI (language learning) and Ambience Healthcare (clinical documentation). Its emphasis on applied use cases reflects OpenAI's broader objective: expanding a developer and product ecosystem built on its foundation models.



Geographically, Silicon Valley's pull on AI has never been stronger. Talent may be globally distributed, with vibrant AI hubs in London, Paris, New York, and Beijing, but both capital and headquarters continue to cluster in the Bay Area. The vast majority (63%) of AI 100 companies and almost all of the Series A and B funds are headquartered in the region. European stars like ElevenLabs and Lovable often raise their earliest rounds locally, yet by Series A, they turn to Silicon Valley VCs. Asian players like Manus and HeyGen follow the same trajectory.

The investor landscape for AI startups is both crowded and open. More than 350 funds have invested in at least one of the AI 100 companies. As capital continues to pour into AI, founders have unprecedented choice. The real challenge isn't access to capital, but picking the right kind of partner: YC's network, a16z's platform, or the OpenAI Startup Fund's technical depth. In a market this fast-moving, who you raise from can matter more than how much you raise.

Conclusion

The past three years have compressed what would normally be a decade of startup evolution. Researcher founders built companies that reached $100M ARR faster than SaaS unicorns reached $10M. Five-person teams generated more revenue than 50-person organizations. Pivots unfolded in weeks after model releases, and markets broke out in sequence as AI capabilities crossed critical thresholds.

But the patterns that defined 2022 to 2025 might not continue into the next decade as is. Some reflect structural change, others, early-cycle noise. **The structural forces are clear**: the sequential market breakout pattern tied to capability thresholds, the rise of researcher-founders fluent in model dynamics, and the shift toward headcount-lean but compute-heavy organizations. These are durable because they stem from the physics of the technology itself: how quickly models improve, how AI replaces human coordination, and how capital converts to compute instead of headcount.

The other patterns might not endure: simultaneous multi-winner markets, compressed revenue timelines, and hypes in specific domains. These are byproducts of an early, fast-expanding ecosystem and will likely normalize as capital tightens, models absorb features, and distribution advantages consolidate.

One problem that all AI startups face is **the enterprise adoption lag**. According to an MIT study, 95% of enterprise AI pilot programs deliver little to no measurable business impact. The problem isn't the technology, it's organizational integration: entrenched processes, data readiness, and change management. This gap suggests that the next phase of growth may not just come from top-down initiatives but from bottom-up adoption that gradually forces organizational change from within.

For investors, this shifting terrain is even trickier to navigate. Traditional heuristics, like headcount growth, organizational maturity, and staged GTM sequencing, are losing predictive power. Value no longer compounds through scale but through **foresight**: anticipating where functionality will commoditize and build moats where models can't easily follow. In this fast-moving technical landscape, investors must also learn to **treat pivots as a feature, not a flaw,** representing a structural response to a stack that keeps redrawing where value lies.

One question everyone's asking is **whether we're in an AI bubble**. The math is uncomfortable: there was $162.8B flowing into AI startups in the first half of 2025 alone. By historical standards, this looks frothy. But bubbles aren't defined by high prices alone, they're defined by the gap between price and fundamental value. **The real question is whether AI applications are building real defensibility vis-à-vis foundation models.** The answer will determine not just which startups survive, but whether the application layer captures lasting value or becomes a footnote in the foundation model era.

What we do know is that the companies that win the next phase won't simply execute today's playbook faster. They'll recognize new capability thresholds, verticals ready for adoption, and moats that compound rather than erode. The playbook for AI applications is still being written, and the most important chapters are probably ahead of us.