

# AI: Charting through Uncharted Territory

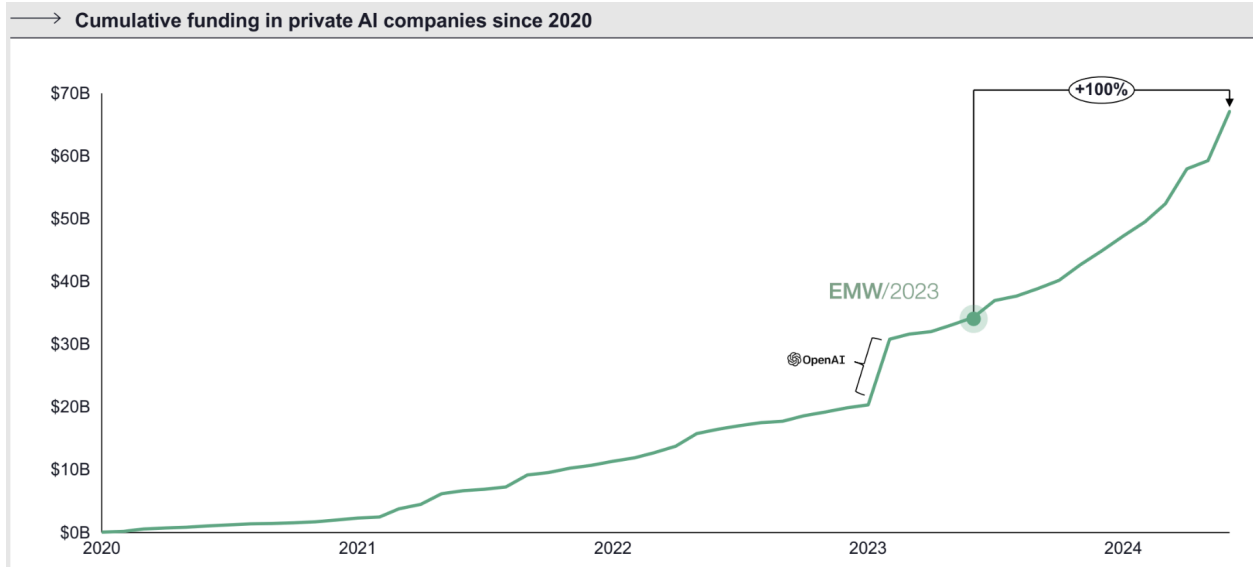
August 2024

*20-min read*

<b>Status Quo: AI everywhere</b>	<b>2</b>
<b>Generative AI: Evolution of automation</b>	<b>3</b>
<b>Market Dynamics: Imagination runs free</b>	<b>7</b>
<b>Value Capture: Battle for primacy</b>	<b>11</b>
<b>Traditional Software: Thriving of the fittest</b>	<b>17</b>
<b>Endnote: Beyond software</b>	<b>20</b>

## Status Quo: AI everywhere

A large proportion of the opportunities we evaluate have AI as a significant, if not central, component of their product roadmap. This stands in contrast to what we were seeing just 18-24 months ago in the pre-chatGPT era. Our team members have all begun to use AI-first products at least on a weekly basis for personal and/or work-related tasks.



Source: Coatue

Similarly in the public markets, the US crossover fund Coatue estimates (somewhat generously) that 2024 year-to-date, 2/3rds of the S&P Index returns and 90% of NASDAQ Index returns were driven by “AI stocks” ([report](#)). The Financial Times recently published an article entitled “Nvidia earnings now rival US jobs report for impact on markets.”

Such a boom necessarily comes with plenty of grand statements, hubris, and opportunism. The more people fan the flames of the hype on the way up, the higher chance of a subsequent downward correction. **Hype and bubbles are part of the technology landscape and aren’t something to be feared.**


This memo lays out our take on what the evolution of AI technology represents, summarizes some of the early achievements and obstacles in consumer and enterprise adoption, and considers what the implications are for value creation in both the new AI software and the traditional software space. We’re realistic about what predictions we are able to make in this fast-changing environment, so we focus on our investing principles and evaluation frameworks rather than “crystal ball readings”. Enjoy the paper and we welcome any feedback!

## Generative AI: Evolution of automation

Generative AI (gen AI) refers to a type of AI involving deep-learning models that can generate high-quality original content such as text, images, video, and others based on the data they were trained on. Gen AI tools are built on underlying AI models, such as large language models (LLMs). In other words, LLMs are the text-generating part of gen AI.

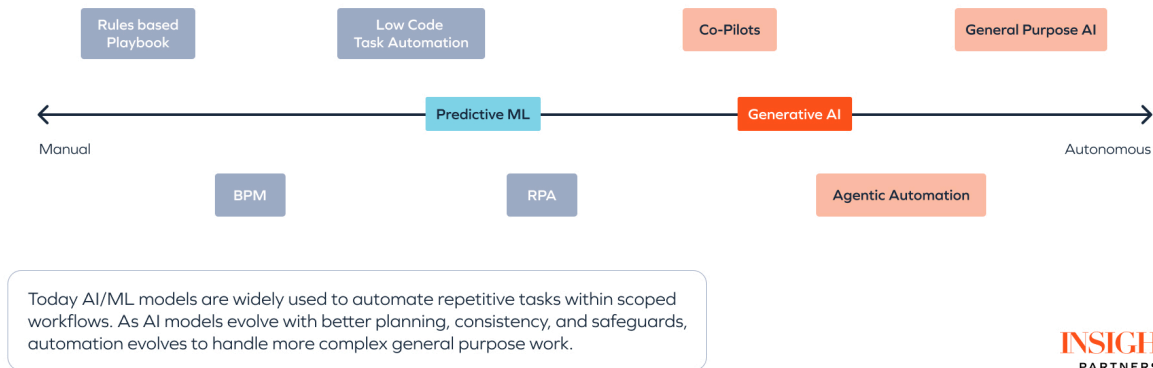
The previous iteration of gen AI had a relatively simple use case within computer vision where selfies were transformed into Renaissance-style portraits and people could see how they are likely to look when they are 80 years old, which quickly filled people's Instagram feeds. **The current version of gen AI was fueled by progress in natural language processing**, which in turn enabled computer-generated essays, poems, and even entire movies.

The underlying fundamental academic work on AI equally experienced dramatic growth.

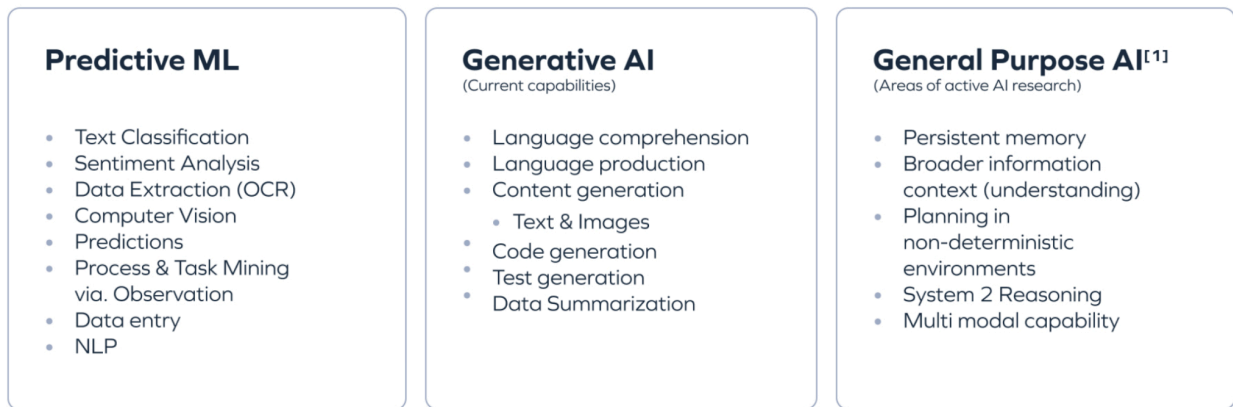


At a functional level, gen AI is simply a technological evolution of software. Insight Partners [report on AI agents](#) provides a helpful schematic on where we are currently positioned with respect to the evolution of machine learning and artificial intelligence:

## Evolution of Automation Architectures



There are important differences between various types of AI, with General Purpose AI clearly representing the proverbial “holy grail” of automation.



[1]: Yann LeCun, Objective-Driven AI – Ding-Shum Lecture, Harvard University and public domain sources

At some point, General Purpose AI becomes AGI (Artificial General Intelligence). The boldest predictions around AI come from people around OpenAI, such as the alum Leo Aschenbrenner’s [160-page paper](#). **AGI itself is a vague and ambiguous concept** so we won’t debate it much further, though there are attempts at creating consensus around what AGI means, for example the [ARC-AGI project](#). **Researchers seem to agree that ChatGPT4 passed the Turing test**, one of the holy grails of AI – making machine behaviour sufficiently indistinguishable from human behavior in the tested aspects. **The big picture remains, however, that LLMs don’t have the capacity (or only very basic one) to understand the physical world, the ability to remember and retrieve things, persistent memory, the ability to reason, and the ability to plan, which are all features of organic intelligence.** Curiously, LLMs perform better when emotional manipulation is used ([paper](#)).

The evolution of AI is summarized as follows by Open AI:

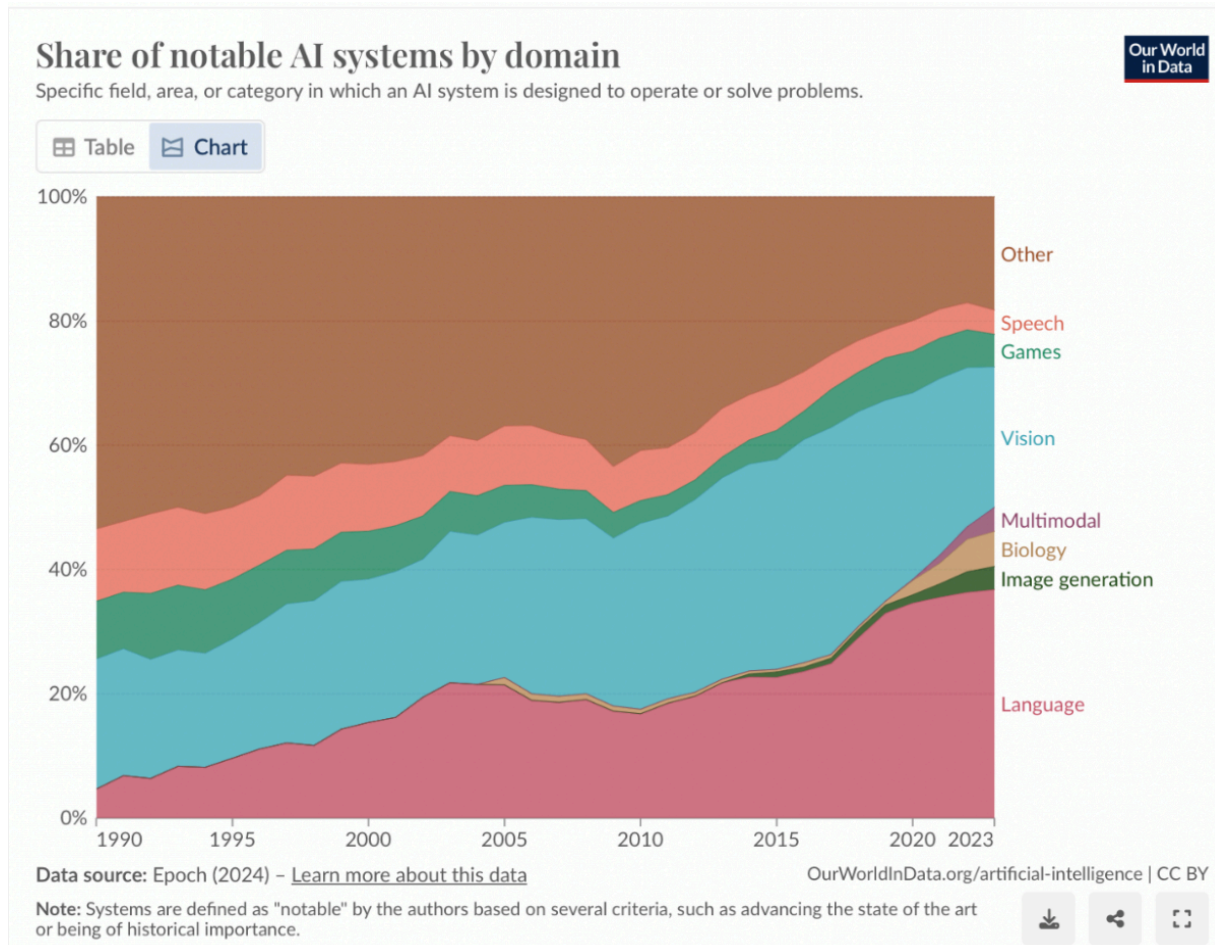
**Stages of Artificial Intelligence**

Level 1	Chatbots, AI with conversational language
Level 2	Reasoners, human-level problem solving
Level 3	Agents, systems that can take actions
Level 4	Innovators, AI that can aid in invention
Level 5	Organizations, AI that can do the work of an organization

Today, gen AI is good at essentially three important functions (Level 1)

- 1) **Creativity.** Image, video, music, text creation
- 2) **Companionship.** Chatting in various modes
- 3) **Basic language reasoning.** Summarization of content, Q&A

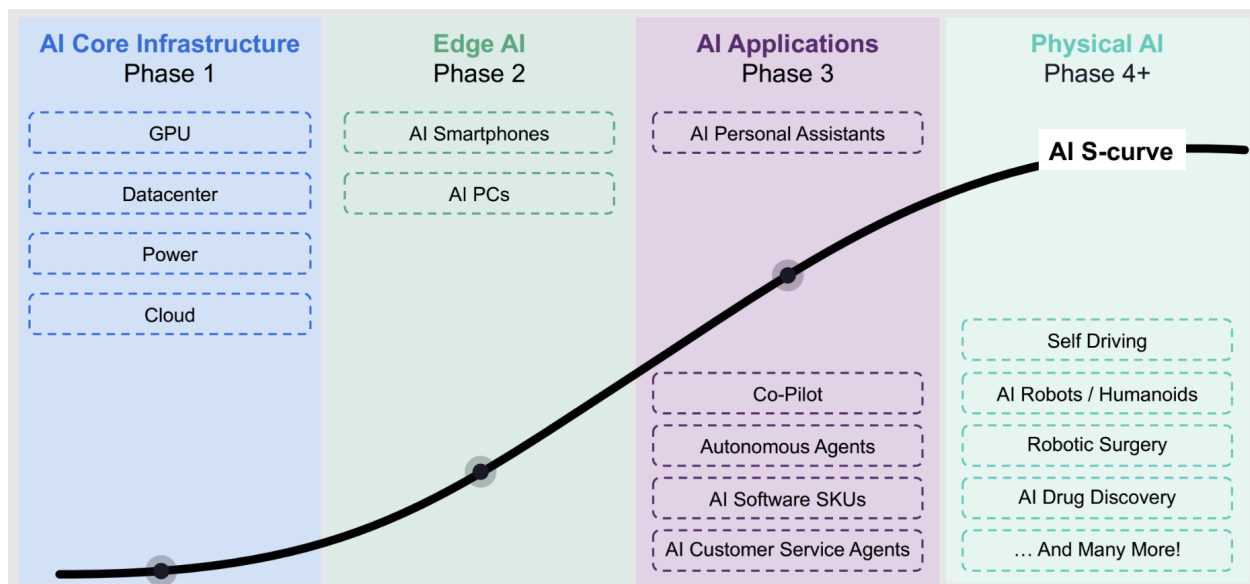
There is progress being made in Level 2 and 3 capabilities. A substantial part of the effort is being deployed on textual analysis and the promising ‘multi-modal’ analysis is in its early days.



This is important because language is only a very approximate representation of percepts and/or mental models, so if we can capture and manipulate richer data beyond text, it can lead to better outcomes. **Regardless of future improvement, the current state of AI models is already quite advanced and even without further progress has the potential to transform industries.**

*So when do the AI robots show up?* The pace of intelligence growth for smart robots will be complex and gradual, but robot hardware has been developing across various form factors (e.g., drones, autonomous vehicles, humanoids, etc.) and industries (e.g., industrial, commercial, home, and defense). The transition from single-purpose robots pre-programmed for specific tasks to general-purpose robots capable of adapting and learning for various tasks and situations is accelerated through spatial AI models.

Even projecting an improving cost profile and data collection, complex AI robots may be far away. To illustrate, while we had the first autonomous vehicle in 2005 ([Stanley](#)), robotaxis don't yet operate anywhere near the economic efficiency of Ubers today, as it's extremely resource intensive to navigate a complex 3D space. **It's hard to gain conviction as to the precise timelines on which the complex gen AI use cases will emerge and be deployed on an economic and regulation-compliant basis in the physical world.**




Source: Coatue. S-curve describes the typical behaviour of a new technology. It starts out expensive, bulky, not very widely adopted; improvement is slow as the fundamental concepts are being figured out. A period of rapid innovation and massive adoption follows, up to an eventual slowdown in meaningful improvement

## Market Dynamics: Imagination runs free

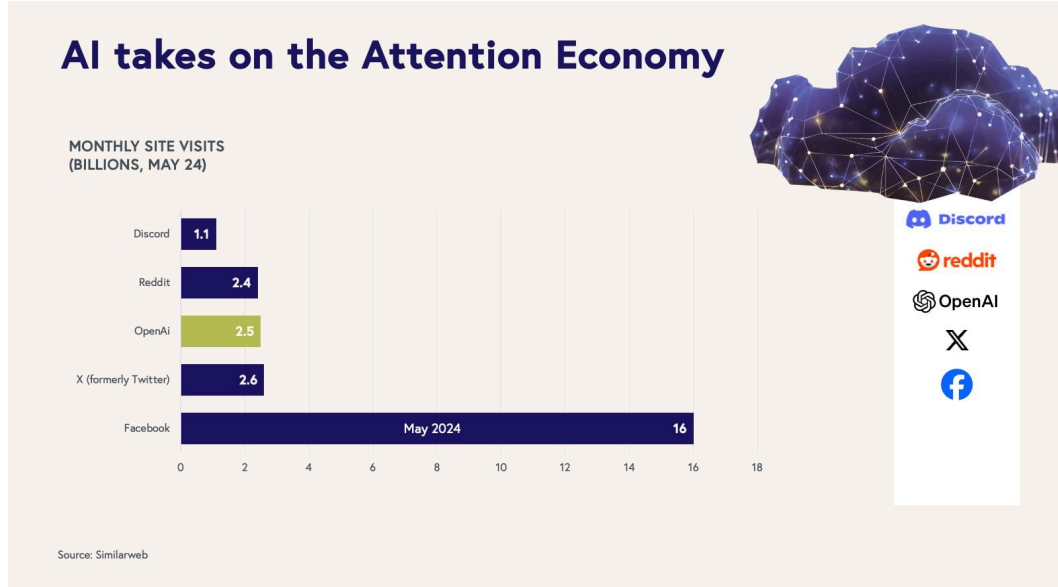
Gen AI technology has seen fast adoption since ChatGPT's first release in November 2022.

# Path to 100 Million Users (stylized)



Source: Sequoia Capital

ChatGPT has competed with major social media platforms for the users' attention.

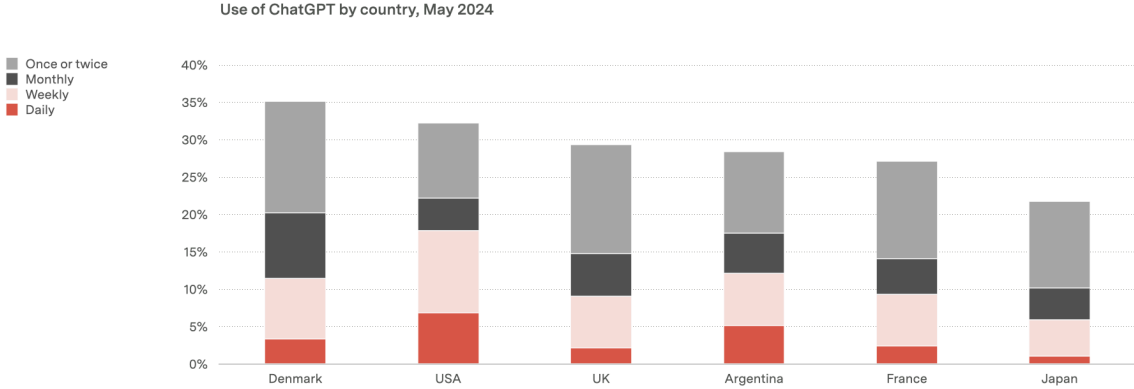


Source: [Bessemer State of the Cloud 2024](#)

One significant point of weakness is, however, that **both retention and engagement of gen AI consumer apps has been relatively unimpressive so far**. While the median 1-month retention of best-in-class traditional consumer apps is above 60%, for gen AI apps is only just above 40%, as per Sequoia's [recent report on gen AI](#). Some of the best consumer companies have 60-65% DAU/MAU; WhatsApp's is 85%. By contrast, gen AI apps have a median of 14%.

## Huge interest, limited use so far

Glass half full: a quarter of people have tried it after just 18 months!  
Glass half-empty: but most of them didn't find it useful - yet




Source: Reuters Institute

Benedict Evans -- July 2024

4

At least for now, **consumers are not finding enough value in gen AI products to use them consistently.** The main exceptions are gen AI “companion apps” (virtual friendships, romance, casual chat), such as Character AI, recently acquired by Google.

## Top Gen AI Consumer Mobile Apps: Engagement by Category



Charts are for informational purposes only and should not be used for investment decisions. Past performance is not indicative of future results. None of the above should be taken as investment advice; see [a16z.com/disclosures](https://a16z.com/disclosures).



On the enterprise side, optimism is high as to AIs general usefulness, however, putting it into practice hasn't been straightforward. The main obstacles are as follows:


- **Hallucinations.** Essentially inaccurate and inexplicable model predictions or “excessive creativity”, a feature of how gen AI models function. In regulated environments such as healthcare or finance, companies cannot afford to have “fake” data being generated. There are ways to minimize hallucinations but it's unlikely we can eliminate them fully



- *High cost and uncertain ROI.* Running gen AI systems is expensive and considering additional complications of systems integration, the ROI is quite unclear
- *Low trust and transparency.* “Black box” perception of AI and lack of ability to explain model outputs. AI systems will likely have to outperform human systems by a very significant margin before AI-generated mistakes become tolerable
- *Low data quality.* A lot of work today revolves around data collection, transformation, storage, tracking and monitoring. Data scientists end up spending the majority of their time bringing data quality up to scratch
- *Data security and privacy.* Companies are worried about leaking data and being in breach of compliance rules. In addition, cumbersome IT systems integration and unwieldy workflow integration is needed to get it fully working
- *Talent.* There is a shortage of people who understand AI sufficiently

According to [recent McKinsey research](#), **only 11% of companies which tested gen AI pilots have adopted them at scale**. Given so many new model releases and applications built on top, it also makes sense for enterprises to wait until a reference architecture has been accepted by the market so as to minimize the experimentation they need to do. **In summary, the current LLM models still lack sufficient quality to be adopted en-masse by conservative enterprises (especially in financial services and healthcare).**

We believe **these barriers are ultimately surmountable in the vast majority of use cases, however the timing of resolving such issues carries with it a lot of uncertainty**. Progress on cost can be made relatively fast following [Wright’s Law](#), which relates to predictions of cost declines as a function of cumulative production. The ongoing “cloud wars” among the major cloud providers should drive the “cost of intelligence” down further as well. When GPT3.5 launched in November 2022 it cost \$0.0200 for 1000 tokens. It’s \$0.0005 now – 2.5% of its original pricing in just 1.5 years!



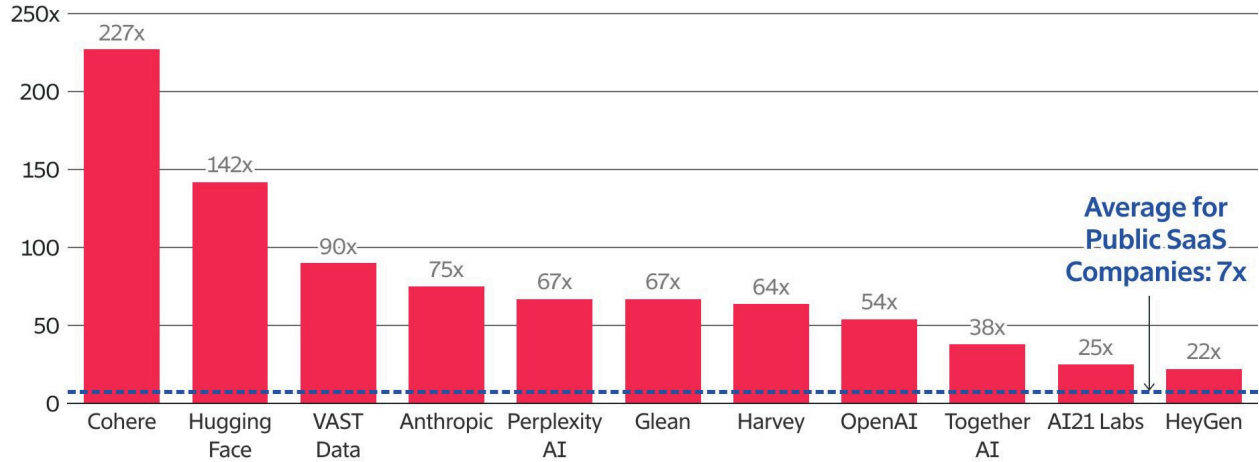
Source: [SAP](#). RCU is Read-Copy-Update – a way of implementing synchronization between multiple processes, ie a unit of processing

**Valuation multiples observed in the market are extreme.** It’s a classic case of an extreme power law, where the mortality rates of startups are very high but those that make it will likely be immensely valuable. While investors don’t know which ones will be the winners, they value each

company as though they were going to become the extreme winners. As Coatue reports, in the US YTD, the average round size is 6x larger for AI-deals than non-AI deals and valuations are 5x higher.


## Dividing Multiples

The most recent valuation multiples for select AI startups



Note: Valuation multiples are based on forward-looking revenue.  
 Chart: Shane Burke • Source: The Information reporting

Observing the space we operate in, the data from YC batches is clear. (YC is the world’s most prestigious SF-based accelerator of pre-seed stage companies). The most recent cohort had >75% of their companies building in AI.



Source: Y Combinator

Benedict Evans -- July 2024

**The good news is that engineering talent from the CEE region is more than pulling its weight in gen AI.** In OpenAI’s early stages, for example, there were about 10 Polish employees out of 50 in total. There are good examples across other top organizations.

## Value Capture: Battle for primacy

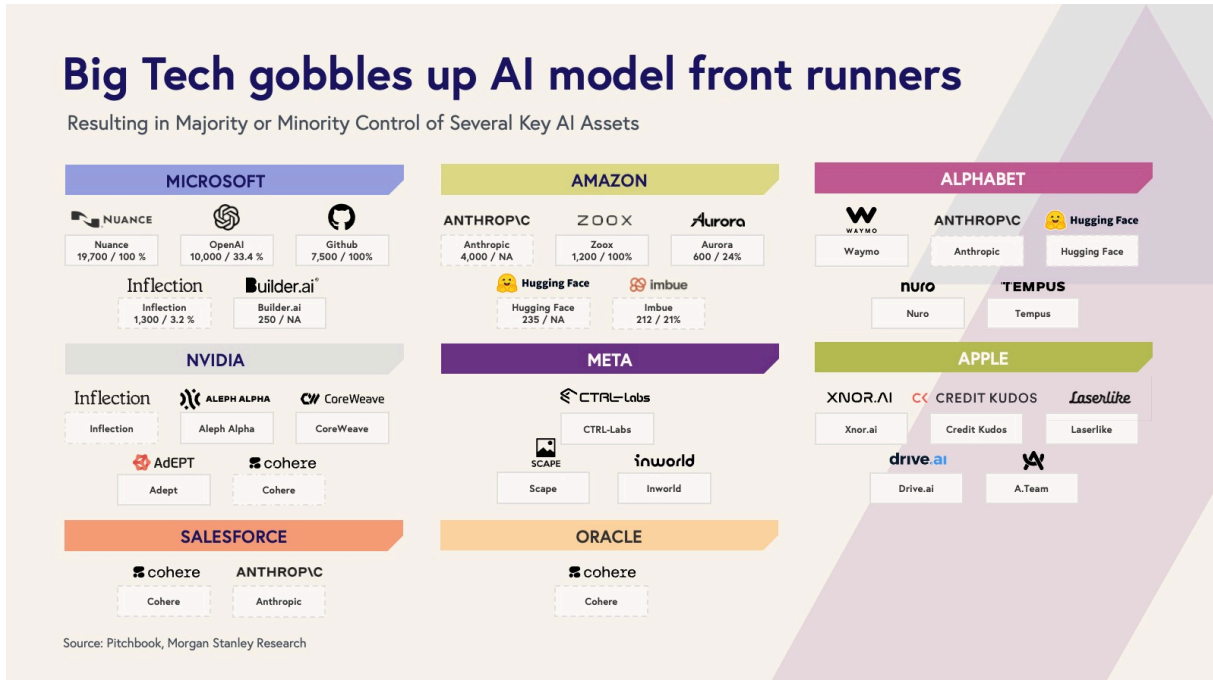
There are several plausible scenarios as to where and how the long-term value creation will occur. Some of the key unanswered questions that will shape the evolution of the sector are as follows:

- *AI scaling constraints.* Does LLM quality scale exponentially, linearly, or asymptotically? Research suggests that hardware supporting frontier models could cost around [\\$250bn of capex a few years out](#). We believe that even if future foundational model versions may become “too expensive” to generate significant quality gains, at current model quality levels, the sector will flourish and over time transform the way software and hardware enables higher productivity.
- *Frontier vs. open source quality.* If the frontier models are able to establish a significant performance advantage over open-source such that people prefer closed-source for most things, then substantial economics will be captured by Big Tech. However, in a world with relatively undifferentiated models and open-source being good enough, value accrues to applications that leverage one of the chosen foundational models.
- *Hallucination control.* Excessive model creativity is a major barrier for enterprise adoption and will need to be much better managed to unlock value in high-cost-of-failure sectors.
- *Multi-step reasoning.* This process can break down complex tasks into smaller, more manageable steps and then execute those steps in a logical sequence. Progress on this front will unlock higher levels of AI capabilities.
- *Fundamental research.* There is a non-trivial chance that some breakthrough happens in the way machine learning is conceived of through new insights from mathematics, for example energy-based models (favored by Yann Le Cun, Meta’s Chief AI Scientist)

We will go out on a limb to say that **no one has a firm grasp of the ultimate answers or even progress towards solving these topics more than 3-6 months out**. The pace of change is too high. In such an environment, we would be loath to make bold predictions, so will remain at the level of categorization and frameworks, that is, “seeing the present clearly”.

**Broadly speaking, company formation falls into three main categories: LLM Models, Infrastructure Tools, and Applications.**

**LLM Models.** There are half a dozen to a dozen great foundational models. In 2023, foundational model companies captured the lion’s share of venture funding in AI, accounting for over 60% of total AI dollars, [according to Bessemer](#). This influx of capital was primarily driven by corporate VCs, who represented 90% of private gen AI fundraising in 2023. Big Tech companies such as Microsoft, Google, Amazon, NVIDIA, and Oracle, now have significant stakes in foundational model companies.




A now-famous internal Google memo—[We Have No Moat. And Neither Does OpenAI](#) claims that technological moats in this area are thin. Note that this doesn't mean owning these "commoditized" models won't be valuable, just as Exxon, Shell, and others supply a commodity while being hugely valuable companies. **What we are witnessing at the moment is effectively an "arms race" by Big Tech where the risk of under-investing is significantly higher than the risk of over-investing.**

### Big tech is spending more than ever — thanks to AI

Standardized quarterly capital expenditure, in billions

Amazon Alphabet Microsoft Meta



Updated Q2 2024

Source: FactSet

Sherwood

This means there is less concern on return on investment in the near term. OpenAI, for example, is expected to make a loss of ca. \$5bn this year despite a large user base. The LLM leaderboards are constantly evolving, with the open-source Llama having made compelling progress recently. Mark Zuckerberg decided to give away Llama's capability for free.

Math		Coding		Instruction Following	
Model	Score	Model	Score	Model	Score
1st Claude 3.5 Sonnet	96.6	1st Claude 3.5 Sonnet	1165	1st Llama 3.1 405B Instruct	90.35
2nd Llama 3.1 405B Instruct	95.6	2nd GPT-4 Turbo Preview	1131	2nd Claude 3.5 Sonnet	90.17
3rd Claude 3 Opus	95.2	3rd GPT-4o	1125	3rd GPT-4o	88.77
4 GPT-4 Turbo Preview	95.1	4 Llama 3.1 405B Instruct	1123	4 GPT-4 Turbo Preview	87.95
5 GPT-4o	94.9	5 Gemini 1.5 Pro (May 2024)	1085	5 Llama 3 70B Instruct	85.65
6 Claude 3 Sonnet	93.3	6 Claude 3 Opus	1024	6 Gemini 1.5 Pro (May 2024)	85.20

Source: Scale AI

**We do believe that there will be a separation between the application layer companies and foundation model providers**, with model companies specializing in scale and research and application layer companies specializing in product and UI.

So can an investor make money in the LLM-space? **We believe that if one can build a domain-specific LLM with access to data that others cannot obtain, with the idea of building focused applications on top of that own specific model**, then one will have the ability to create significant value. Our investment in [Yoneda Labs](#) (Czech-Polish founding team) follows this thesis of differentiation, as they are building their own robots to perform a large set of pre-selected chemical reactions that will form the basis of the unique dataset based on which they look to [build the "OpenAI for Chemistry"](#). This captive dataset which can self-enhance through usage, leading to a better and better outcome for the user, is not replicable from common data sets. We would love to find more businesses like this.

***Infrastructure Tools.*** The infrastructure layer includes all of the tools that developers will use to build applications and deploy them to the real world (Here a good [overview by Bessemer](#)). This includes traditional cloud vendors like Microsoft and AWS who are well positioned, all the many AI monitoring, orchestration, and observability tools, frameworks to build and deploy agents, vector databases, etc.. There will be some large winners here as incumbents may find it hard to get the developer traction needed for this new category, but it's quite hard to tell who will win. Nevertheless, we have made a well-considered bet on AI agent workflow solutions with [E2B](#), which is a Czech team, [building a next-gen code interpreter](#) for AI apps.

**Founder Focus.** We think the successful founder profile in this space will be someone who is extremely methodical and keeps up with the latest developments and intentions of Big Tech, as Google/ Microsoft/ Amazon are moving very fast in this space and it's hard to just follow one's gut instinct and "solving one's own problem" at a particular time. They are a talent magnet for other first-rate engineers.

We believe with E2B we have backed a team with such qualities, and they've all left Prague and settled in SF to be close to the pulse of the market and customers. (They will soon announce a large \$9m seed round led by a reputable AI-orientated Bay Area investor [Decibel VC](#), whose deal partner is also a "superuser" of E2B's solution).

Elsewhere we backed [Artificial Agency](#) together with our scout [Martin Schmid](#), the founder of [EquiLibre Technologies](#). Artificial Agency is largely composed of an ex Google DeepMind team that is building a generative behavior engine based on an own model that integrates runtime decision-making into any aspect of video game development. They recently announced a large [\\$16m seed round](#). Here, we believe that the team's world class expertise across AI and gaming development experience will make game character development faster, more individualized, and more realistic and that their integration with top game developers will serve as a defensible moat as they'll be an infrastructural tool the top game developers can't live without.

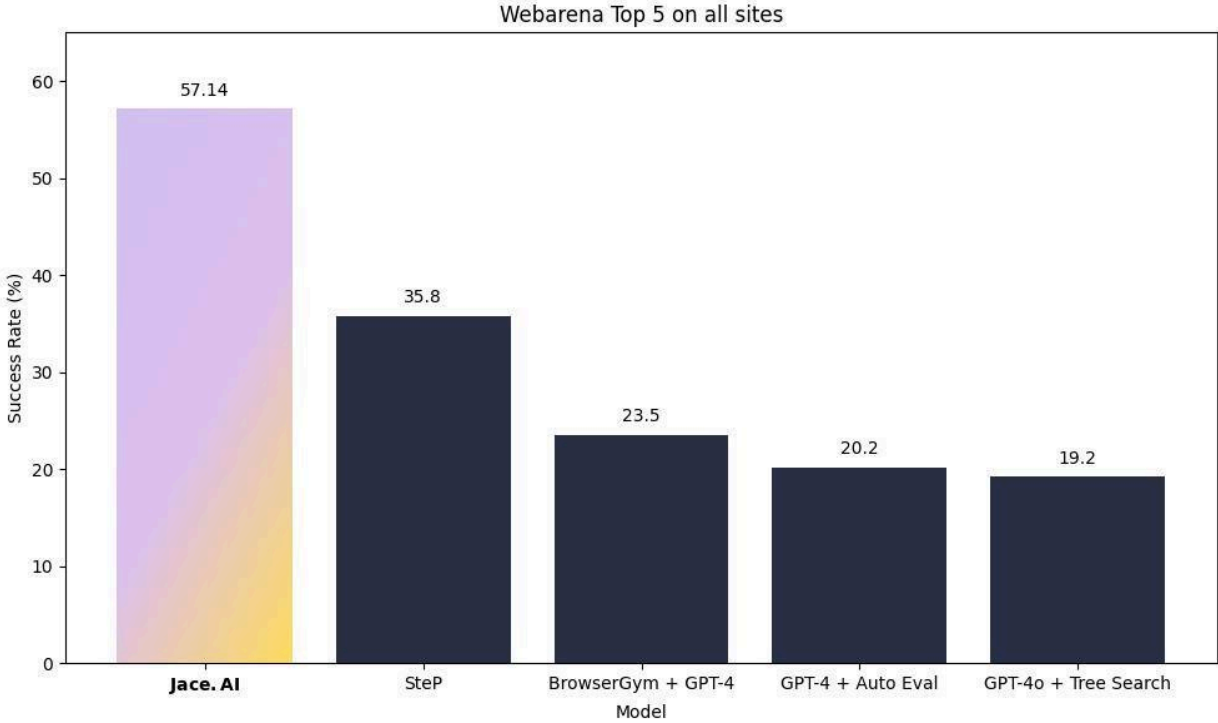
**Applications.** Over the past year, it seems **we have settled into a few common archetypes of AI applications**, although startups are still moving between them as the technology improves or as they realize that one of them is the right approach for go-to-market in the market they're going after.

- [\*AI co-pilots\*](#), where AI functionality helps assist workers perform some core work or operational related tasks
- [\*AI colleagues\*](#), where AI agents autonomously handle some subset of tasks for human workers, with human workers "managing" them
- [\*AI native services\*](#), where companies provide an end-to-end service to other companies using a mix of AI agents and humans on the backend

We believe all archetypes can become successful businesses, while the former two are more likely to be scalable and achieve tech-like valuations. In most instances, the AI will not necessarily be the differentiating factor; it will rather be the overall product superiority and "unfair" distribution advantage a company can generate.

[\*AI co-pilots\*](#). We have backed [Upheal](#), (Slovakian founding team based in Prague) which [helps therapists with the most repetitive tasks](#) connected to them working with patients and managing their practices. We are under no illusion that this is not a competitive space, however we have strong belief that **Upheal is building a special product that nails the value proposition better than anyone else in the market**, and the momentum with customers has been outstanding. (Upheal is in the final stages of closing a \$10m Series A which had been preempted by a prestigious transatlantic fund with a 30-yr history).

AI colleagues. We have backed [Zeta Labs](#) (Polish founder) which helps SMEs automate certain tasks with autonomous agents end-to-end. Autonomous agents are entities that are able to conduct certain types of tasks absolutely independently. And we consider them as a very important shift in the sense of how software products will be built. Once Zeta Labs or any other startup working in this field figures how to effectively plug the agents into the companies' workflows, **these startups will no longer be selling enhancements of the productivity of employees but rather a whole AI employee itself.** We can expect that for many new software solutions we will no longer see a standard pricing model focusing on price per seat per time unit (as was natural for tools enhancing employees productivity) but rather price for the workload completed. Jace, built by Zeta, significantly outperforms current state-of-the-art across a range of ca. 800 tasks.



AI-native services. We invested in [TrueClaim](#) (Czech founder) that helps companies save on healthcare by streamlining billing processes with the help of AI. Their business model may eventually transition to a fully automated healthcare cost optimization solution and aim to replace other human-staffed Third Party Administrators (TPAs) of which there are several players with \$1bn+ in revenues. The opportunity to capture revenue from these analogue players who have low customer satisfaction is very large and exciting for us.

Our bet in the app space [Medisearch](#) is keenly aware of the potential distribution disadvantage vs. incumbents and we're working with them to mitigate these factors and turn them into an advantage in the hyper-competitive, supercharged "Darwinian" environment of AI-powered search. There are several players successfully competing in the broader search space so far.

**Founder focus.** In general, we have found that non-technical founders in the applications space overestimate the technical defensibility they have, but technical founders in the space overestimate what they need to get to defensibility (i.e. superior distribution can win over a product edge or head start time-wise). This space is particularly prone to opportunistic founders, so we look for someone who is obsessed with increasing the delta between their app and a new app that could pop up or a Google/ Microsoft integrated app. Founders in this segment need to be both feature-focused and very creative with respect to distribution. They truly need to want to be entrepreneurs for the right reasons.

There are some things that don't change about business-building in this new AI-infused environment. **AI is a technology and not a product.** So where in the stack a company plays is less important vs. thinking about products and the user and how the initial experience leads to a long-term advantage. So whether LLMs, infra or app, we look for a product (or beginnings of it) that feels magical, that feels 10x better to the end user whether a developer tool, infrastructure monitoring, or something else, and such a company will have a chance to succeed. Incremental solutions will not do and that magical experience is likely powered by a data moat but also involves the basics perfected, such as UI, UX, etc.

The moats are in the customer experience, not necessarily the data, barring a few exceptions such as Yoneda Labs in our portfolio which we pointed out. It's fair to assume that the "data moats" are on shaky ground: the data that application companies generate does not create an insurmountable moat, and the next generations of foundation models may very well obliterate any data moats that startups generate. Rather, **workflows and user networks seem to be creating more durable sources of competitive advantage.**

In summary, we believe in the following medium-term implications:

- The "hype" enhancing the "Darwinian forces" that ultimately catalyze winners and losers faster and hence allows less margin for error when building companies
- AI will become more commoditized over time and the best product will still win, but it will need to create an advantage against incumbents with pre-existing relationships and customer networks. Incumbents have proven to be very nimble in many areas
- Old rules of investing broadly still apply; heightened focus on solving pressing problems for users, rather than just "any" problems. Talent drives product edge, customer love drives overall moat
- Success will not be concentrated in Big Tech, there is an immense opportunity for new unicorn and decacorn formation outside of foundational models

Given hyper-competitiveness and less visibility on how the startup cohort evolves, we are more active in pre-seed with exploratory checks and lean in as soon as we see de-risking and the initial signs of a differentiated product and distribution. If we were to "come from behind" into a competitive seed and Series A rounds, we think that's a less advantageous place to start.




## Traditional Software: Thriving of the fittest

“Traditional” SaaS software companies are also adapting to the new possibilities offered by gen AI. Many posts have come out on the theme of “Software is Dead” as a consequence of gen AI’s rise and people passionately argue about software engineering getting automated (for example, [here](#)). Such brave new world is nicely encapsulated by a recent tweet from Andrej Karpathy (who is of Slovakian heritage, incidentally), a co-founder of OpenAI and former Director of AI at Tesla:



Full thread [here](#)

How far away is such a world? Is that vision simple or simplistic? These are hard questions to answer. Contrast that statement with the stock price of the thriving software giant Oracle, which was founded in 1977 and successfully adapted its offering time and time again, most recently to remain relevant in the SaaS era, and now the (gen) AI era.



Source: Google Finance. Data as of 6/14/2024

Currently, the combined market capitalisation of SaaS companies is ca. \$1.3 trillion. Annually, it is estimated that we spend \$1 trillion on software engineers globally.

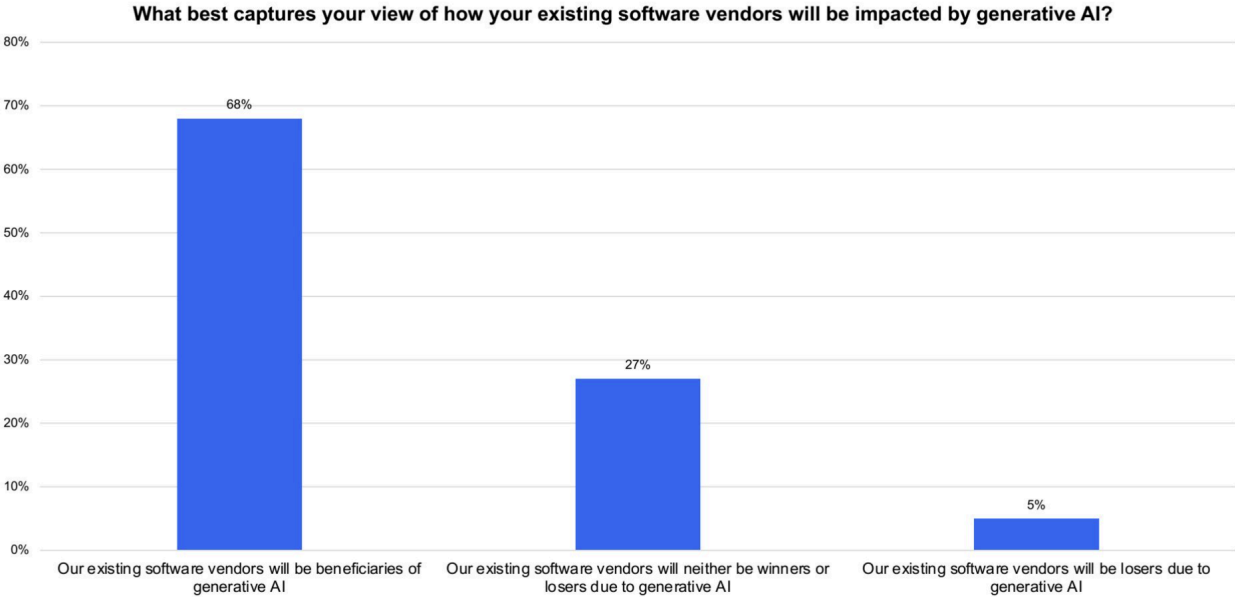
Hard to imagine SaaS companies are facing extinction. It’s true that it has become easier to build a basic SaaS solution and what gave a company an edge 5–10 years ago is now table stakes, but that doesn’t mean that the best entrepreneurs won’t find new ways to differentiate and out-innovate the competition. Thanks to products like [GitHub Copilot](#) more and more code

will be written by AI. How much value unlocks if engineers become 20-30% more efficient? Countless solutions are being built as AI-first companies. A few examples:

- *Root Cause Analysis*. Automated basic and advanced developer workflows
- *Automated Sales*. Prospecting for customers via automated phone calls
- *Automated Compliance*. Solutions that manage companies' compliance obligations
- *SMB Functions*. Customer Service, Scheduling, Marketing done by AI agents
- *Personal Finance*. Finding savings, aggregating demand, negotiating deals, etc.

Those that talk about the end of software argue that software development will become so easy that the market for software products will shrink, and companies will develop more custom software based on exactly what they need. **This may happen in some areas, but we doubt that non-tech companies will suddenly become good at building great software** for their teams. There will be more average software, but premium software will still have a place.

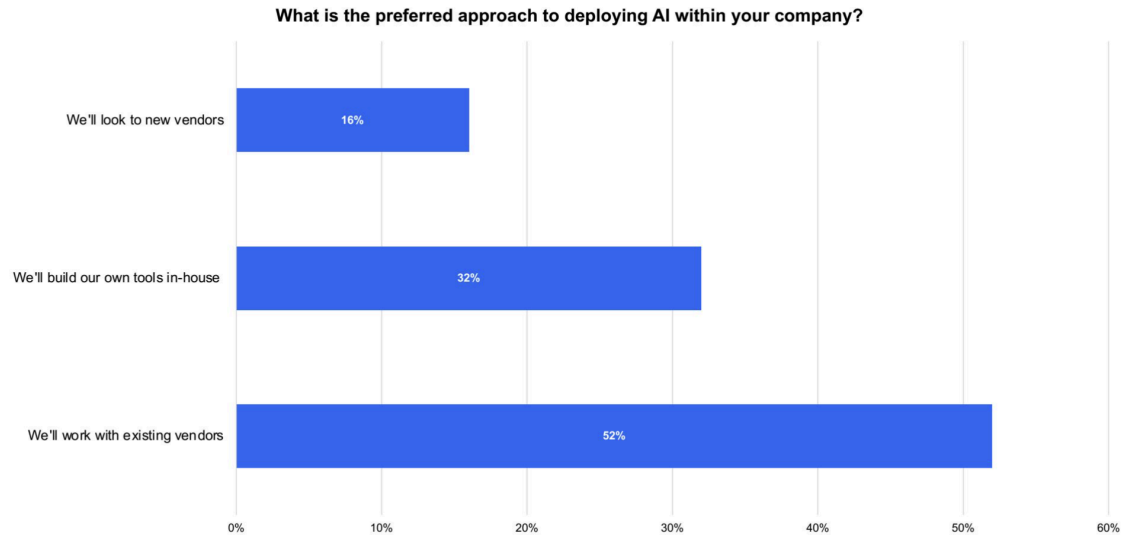
To go from Andrej Karpathy's grand vision to a robust, multi-user workflow solution with integrations, user permissions and complicated business rules, let alone to something that can replace, for example, a CRM, ERP, or HR system in an enterprise, and be fully compliant with GDPR and other regulations, is an extremely large leap. It's also likely that "build" (vs. buy) will always be a headache for companies to implement and not always the way to go. **Big tech won't want to build hundreds of problem-specific and vertical-specific apps.**



Source: Avenir Survey Data (n=75), June 2024

**AI will change a lot about software**, and many software companies will become casualties. Those traditional software companies that adapt will be more important and valuable rather than less. It will take time, but both legacy and new SaaS companies will become AI-first, abstract away the complexity of deploying LLMs, and "AlaaS" will become the new SaaS. **Many of today's top SaaS vendors will adapt and thrive.**

It's also uncertain if foundational models will ever become great in highly specialized domains where deep knowledge about industry workflows and access to proprietary data to train models are necessary (We believe OpenAI will not compete with our Yoneda Labs). Until we have truly super-intelligent technology, specialized solutions for specific use cases or industry verticals will likely be better. And as mentioned, such time may be many years away.



Source: Avenir Survey Data (n=75), June 2024

Our summary view is that some software will be dead, especially those in areas where LLMs thrive already. More generally, the replacement cycles of average software will shorten, and LTV/CAC metrics will worsen, putting pressure on multiples (software trading at 5x revenues on average may not be “the bottom”).

There is also a strong opportunity to replace the services component of software with pure software (i.e. automation of service tasks), which represents a very large TAM. **We are confident that overall software markets will grow and thrive.**

## Endnote: Beyond software

**Our investment world doesn't end with software.** We see a lot of company building by talented people with strong [STEM](#) backgrounds. They believe **hard science is the best conduit for entrepreneurship and building products with true IP, durability, and a large impact – a mission we love to support.** These sectors will also make strong use of AI capabilities over time and many AI-first companies will be formed. But software and AI will often be a means to an end rather than an end in itself.

***Biotech.*** We invested in [Sensible Biotechnologies](#) who are turning living cells into cost-efficient factories for manufacturing high-quality mRNA to unlock novel therapeutics and vaccines. **They have recently obtained verification that they produce the purest mRNA available on the market,** and consequently they've signed an MOU with their top-of-the-wishlist potential customer. They also secured nearly \$50m of non-dilutive capital. Together with the founder, who also plays the role of a scout with our fund, we are identifying other teams that are innovating in this space and the broader space of biosecurity.

***Systems resilience.*** Energy transition, defense, and new industrial systems similarly have a large capital need and large markets to sell into. They often require software-enabled hardware or elements of offline and online solutions coming together. It's a space we have entered so far only through our investment in [Woltair](#), however, we are evaluating an attractive pipeline of opportunities. Through our investment in [Rohlik Group](#) and others, **we are no strangers to hybrid, non-obvious business models that however translate into extremely resilient moats and enduring qualities of a market leader.**

While we have perspectives on the opportunity set ahead, we are guided by the vision of our founders. Their talent and dedication is exceptional and we are convinced they have a good chance at realizing their ambitious visions. Below are brief profiles of a half dozen of them.

Company Focus	Nat.	HQ	Prior Experience	Education
Automation of SME operations	PL	Warsaw	Meta/Facebook Software Engineer	Oxford, MS Mathematics & Computer Science
LLM for chemical reactions	PL	SF	Jane Street Software Engineer	Cambridge, BA Computer Science
Automated healthcare cost optimization	CZE	SF	Pine Park Health, VP Operations	Stanford, MBA, Wharton, BSc
Runtime environment for AI agents	CZE	SF	Bootstrapped multiple startups	Charles University Prague, MS Computer Science
AI-powered healthcare search	SVK	Bratislava	ML researcher Microsoft, Amazon	Oxford, PhD Machine Learning
Novel mRNA production platform	SVK	Bratislava & Oxford	Research Fellow, Stanford	Oxford, PhD Engineering, Imperial College, MEng