

Investing in AI Startups - Some Thoughts

AI is the hot topic of the moment, and I thought I would share a few thoughts on it.

I recently participated in a panel (picture at the end) titled “AI Investing: Where Are the Commercial Opportunities and Where Not to Invest”. The panel was organised by SuperInvestor, one of the largest private equity conferences, which this year took place in Monaco (note: how convenient it is to be attending a conference in Monaco - as I now write this on my way back from Slush in Helsinki).

Below are some high-level points I made as well as some extracts from an interview with Sam Altman, co-founder of Open AI, and a panel discussion with some of the biggest names in the industry, giving a glimpse into where they see AI going.

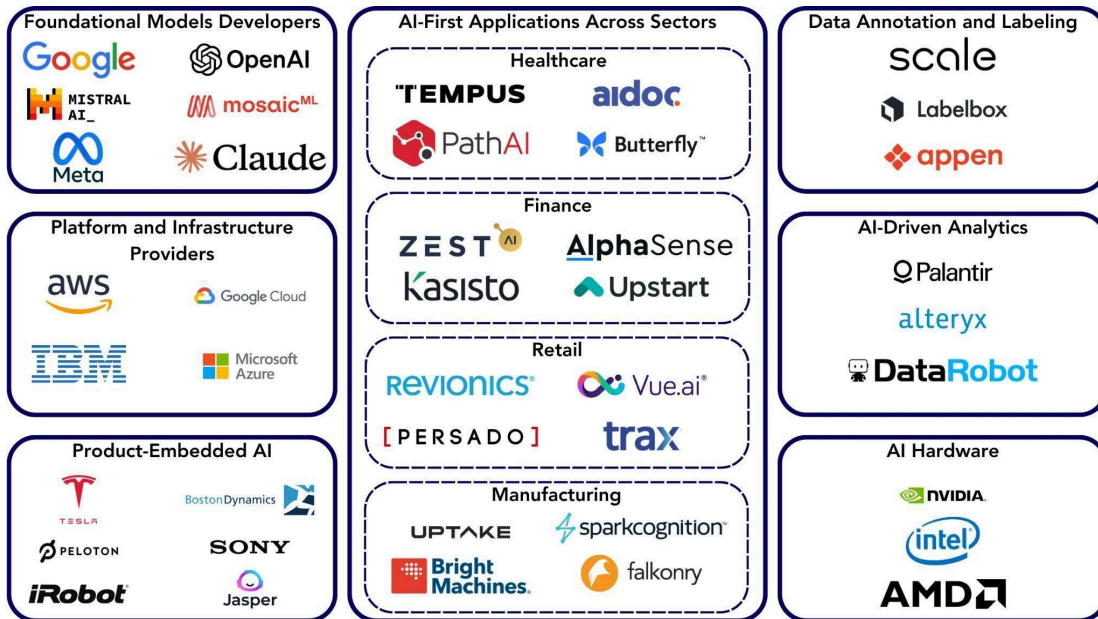
1. Is it a fad or a game-changing technology?

This is a question I ask myself whenever I see the hype around a particular sector. Having spent my earlier career on trading floors, I have seen my share of boom and bust and “this time it’s different”. Venture capital tends to get excited even more often than the financial industry does, with a new sector becoming the new focus of the tech industry every 18 to 24 months. So as an investor, I tend to err on the cautious side when I see such a surge of interest in a sector. However, for reasons that become obvious when simply using some of the AI applications that have sprung up recently, and going through a multitude of AI-focused startup decks, I believe that AI will fundamentally change how we live and work.

2. Where should we invest?

When we talk about investing in AI companies, one first needs to define what an AI company is. AI is related to a wide range of companies, ranging from Nvidia to AWS or the latest video-generating startup. As an illustration, below is a non-exhaustive list of large companies that are exposed in one way or another to the growth in AI.

As far as Monte Carlo Capital is concerned, so far we have invested in several verticals: foundational models (eg Klay), AI-driven analytics (eg Cognovi Labs) as well as AI-first applications (eg Borderless) and companies building AI to operate hardware (eg Algorized).



3. Evaluating AI startups

So, how do we look at AI-first startups?

3.1 Keep your head on

To begin with, it's important not to succumb to the hype, and we use the same high-level filter as with any other business:

- Are they solving a real problem,
- Is the market big enough to generate outsized returns to investors,
- And, most importantly, whether we believe the team has what it takes to deliver.

3.2 AI specifics

However, AI-focused startups present several elements that need to be considered, including:

- Lightning-Fast Technological Changes

The industry is evolving at an unprecedented pace, presenting both opportunities (rapid growth leading to a quick exit potential) and risks (the chance of quickly becoming obsolete). It's crucial to invest in companies that become stronger with every release of a new version of a Large Language Model (LLM) rather than risk becoming commoditised. All signs point to models that will keep strengthening, with abilities that will keep stunning us (see comments from Sam Altman further down).

As with any business, and even more so in AI, it's vital to look at companies that can build a strong moat. This can be achieved in a variety of ways, such as:

- **Data:** Companies with unique, high-quality datasets have an edge because they can train models with data that competitors lack. As they grow and collect more data, that moat gets wider. Self-driving companies like Waymo, for instance, have collected millions of miles of driving data, while Google and Facebook have ongoing significant data streams that are hard to replicate.
- **Team and Technology:** Companies like DeepMind, OpenAI, or Nvidia build competitive advantages by assembling top talent and developing cutting-edge technology. However, this is a challenging game to stay on top of, needing considerable resources. How long can a company maintain its tech lead?
- **Business Model:** As with any business, look for traditional moats such as network effects, intellectual property, high switching costs, regulatory barriers, distribution channels, etc.

- Data Acquisition and Model Training

It's essential to understand how models are trained and how the data is sourced. The last thing you want is to face legal challenges for using proprietary data or violating privacy laws. For example, Clearview AI faced significant backlash and lawsuits for scraping billions of images from social media platforms without consent, leading to bans in several countries. Similarly, Google's AI ethics team encountered controversy when concerns about the ethical use of data in training large models were raised, resulting in internal and external scrutiny.

Klay, one of our portfolio companies, is building a foundational model around music and has chosen to work with the music industry, having just announced a partnership with Universal, leading it to build a moat in the process.

- Cost of Building the Business

Traditional SaaS/software investors are accustomed to scalable businesses with limited investment, but AI companies, while often very scalable, also share some of the features of deep tech companies. This means substantial resources are often required, not only in computing power, data acquisition, and model training but also in hiring top talent, costly in terms of both cash and equity.

- Customer Adoption

While the technology is advancing rapidly, user engagement and retention can be a question mark. Many early adopters may be curious, but if the output, user experience, or clear need isn't there, churn rates could be significant. Additionally, while the tech industry often embraces innovation with enthusiasm, more traditional industries tend to adopt new technologies at a slower pace. Despite all the hype around AI, adoption from large companies could be well below expectations if the AI-first service is not offering a significant improvement. There is a real risk that many of the offerings are just nice-to-haves for techies

and that enterprises won't be persuaded to adopt them at scale. Cognovi Labs, one of our historical portfolio companies, is building a business whose focus is the US Department of Defence (DoD), a very challenging customer to crack, with a long sales cycle, which in itself is a moat.

4. Where is AI going?

When we think about AI "taking over" we are usually thinking about AGI, but what is it? Artificial general intelligence (AGI) is a type of artificial intelligence that matches or surpasses human cognitive capabilities across a wide range of cognitive tasks. In other words, AGI represents a stage in which AI systematically performs better than humans across every task, showing a heightened capacity for autonomous reasoning and real-world actions unsupervised.

4.1 YCombinator Interview with Sam Altman

I wanted to share a recent [Interview](#) with Sam Altman, where he discusses a number of topics, including how he built Open AI and his take on the future of artificial intelligence. Here are some of the extracts specifically related to the future of AI.

Altman breaks down the evolution of AI into five distinct levels, each representing a milestone in the journey toward AGI. He further elaborates on these stages addressing their current state and future potential (this [article](#) further dives into those five stages)

Interview extracts:

"We realised that AGI had become this badly overloaded word and people meant all kinds of different things. [...] You have these level one systems which are these chatbots. [Then] there'd be level two that would come, which would be these reasoners; we think we got there earlier this year. [Level] Three is agents – the ability to go off and do these longer-term tasks, maybe like multiple interactions with an environment, asking people for help when they need it, working together – all of that. And I think we're going to get there faster than people expect. [Level] Four is innovators – that's like a scientist [with] the ability to go explore a not well-understood phenomena over a long period of time and understand what's [going on]. And then level five – this is the sort of amorphous-like [level four] but at the scale of the whole company or organisation. That's going to be a pretty powerful thing."

"The models are going to get so much better so quickly. What you can do as a startup founder with this versus what you could do without it is so wildly different."

"I can see a path where the work we are doing just keeps compounding and the rate of progress we've made over the last three years continues for the next three or six or nine or whatever. If we can keep this rate of improvement or even increase it, that system will be quite capable of doing a lot of things."

Altman sees a future where AI unlocks solutions to humanity's greatest challenges:

"Fixing the climate, establishing a space colony, the discovery of all of physics, near-limitless intelligence and abundant energy. I do think all of those things, and probably a lot more we can't even imagine, are maybe not that far away."

"You will have companies that make billions of dollars per year and have like less than 100 employees, maybe 50, maybe 20 employees, maybe one."

4.2 Panel on the future of AI

The [Panel](#), which took place in October 2024, brought together some of the biggest names in the tech industry, including Eric Schmidt (former CEO of Google), Ben Horowitz (Co-Founder of Andreessen Horowitz), Travis Kalanick (co-founder of Uber) and Shou Chew (CEO of TikTok)

Here are some of the extracts:

Eric Schmidt: *"it's reasonable to expect that within 6 to 8 years from now, so 2030 right after that maybe 2032, under current growth rates it will be possible to have a single system that is 80 or 90% of the ability of an expert in every field; so 90% of the best physicist, 90% of the best chemist, 90% of the best artist. When you have such a thing you have a nonhuman that is effectively smarter than any human because no human can dominate all of those fields"*

Eric Schmidt: *"I will assert that we as humans are not ready for the arrival of this, we're just not ready for it"*

Ben Horowitz: *"as people have been saying, it's such a gigantic Market. It's the biggest market we've ever seen, [meaning] that there's going to be kind of money to be made everywhere"*

Ben Horowitz: *"big companies can optimize very fast with AI and so as a new company you have to ask the question [of whether] I can get a new application to market faster than the big company can get a good product [out], that's the race. I'd say in some cases it's the biggest private equity opportunity of all time"*

Travis Kalanick: *"workflows, customer support, onboarding sales...[are] all getting automated will make an already strong business model like 'supercharged' and I think for the biggest companies create trillions of dollars of value"*

Jack Hidary: *"the Mantra of AI or die is real, it's not just a phrase, it's happening right now, it's happening on several levels. And certain companies and countries that do not engage will die"*

Jack Hidary: *"there's a whole other side of AI that has had less focus and that's quantitative AI. That's AI based on equations and data: quantitative data, numerical data... Large quantitative models (LQM) are complementary to LLMs [and] allow us to do at scale and speed [what] we do that on the same GPUs but [tweaked] in a very very different way [and trained] on a very very different kind of data. The data doesn't come from the internet, [it] does not come from downloading Wikipedia and Reddit and social media - instead, we generate the data from the actual equations that govern our world. This is a whole new superpower for humans"*

💡 While LLMs are trained on textual and unstructured data sourced from the internet, they are not equipped to handle large and complex mathematical equations and simulations. As such, their strength lies in their contextual understanding and generating human language. On the other hand, LQMs are trained using equations and quantitative data (such as those governing biology, physics, chemistry, etc). Such models are, therefore, not trained on any internet-sourced dataset but rely on data generated from foundational scientific principles and equations. While

LLMs perform well for processing unstructured human-generated data, LQMs excel at simulation and problem resolution using mathematical and physical equations.

Eric Schmidt: *"It is believed that somewhere in around 5 years, no one knows exactly, the systems will be able to write their own code, they literally will take their code and make it better, and of course that's recursive"*

Ruth Porat: *"The art of the possible is fundamentally changed, this is a generational opportunity like we have never seen before"*

Conclusion

The hype surrounding AI is real, but so is its potential to reshape the world. As investors, we can participate in a once-in-a-generation opportunity by focusing on companies solving real problems and building defensible moats, while remaining aware of AI companies' own set of challenges. The next wave of world-changing businesses will emerge from this space, making it an incredibly exciting time to be a seed-stage investor.



SuperInvestor Monaco: 'Panel on AI-Investing'