

Meet Cosine, a human reasoning lab

As part of our quick founder questions series – or QFQs – we spoke to Yang Li, cofounder and COO of Cosine, about scaling code through compute, not headcount and Genie, an AI software engineering model.

Temps de lecture : minute

12 November 2024

Cosine is a human reasoning lab that uses AI to synthetically generate data that mimics human reasoning and thought. Our first product is Genie, a fully autonomous Software Engineer that has been taught how to reason and do coding tasks exactly like a human.

Why are you doing what you're doing?

We wanted to automate our coding jobs, but discovered that existing AI coding assistants fall short of completing full projects. They're either basic copilots limited to the open file or just guessing solutions with no relevance to real projects. That's why every product demo on X shows coding assistants building from scratch, never addressing actual developer tickets.

Our goal became to create a world where tech resources aren't bottlenecked, and where we could scale code through compute, not headcount.

Tell me about the product - what it is, what it aims to achieve, who you work with, how you reach customers, USP and so on?

Our first product, Genie, is a fully agentic AI Software Engineer trained to reason and work on software engineering tickets like a human. You can assign Genie a ticket or task, and it will break down problems, research, and reason just like a human would. You can interact with Genie as you would with a colleague—interrupt it, provide guidance, and add comments to its work, and it will respond exactly like a human.

Our USP, which allowed us to achieve the largest score jump in the history of the SWE-Bench industry benchmark, is that we quickly realised the key was "showing" the AI exactly how a human works, step by step.

The problem is that all available data, including what OpenAI or Anthropic uses, only represents the end product, the final result of the work. We spent 12 months building the world's only data pipeline capable of reverse engineering the reasoning steps, knowledge, and mistakes a person experiences to complete the final code. With this unique data, we are the only company that can train an AI to truly mimic human developers, rather than relying on the AI making 100 or 1,000 attempts until it gets lucky.

Anyone can join our waitlist, and as of 3 weeks ago, we've started rolling people off it, with users generating and completing software engineering tasks daily.

How has the business evolved since its launch?

When we first started, LLMs weren't advanced enough, and context

windows were too small to build what we wanted to build. So, while waiting for LLM research to progress, we developed several tools we knew our final product would need. Our very first product, a semantic codebase search tool, gained a few thousand users in its first month. But, our ultimate vision was always to build a fully agentic software engineer trained to behave like a human, that didn't just rely on throwing spaghetti at the wall until "some code" worked.

Fortunately, we have photographic evidence that this was the original vision as we presented this at a UCL event the month we were founded!

What is your favourite thing about being a founder?

Working with people you genuinely like and respect, because you got to pick them, knowing that if there was a crisis (and there are many), the team you have alongside you is the team you'd pick every time.

Which founders or businesses do you see as being the most inspirational?

My YCombinator partners Michael Siebel, Tom Blomfield and Richard Aberman, as they each bring something slightly different to the table and are all ruthlessly pragmatic - they instil a sense of deterministic optimism into me.

What has been your biggest business fail?

Leaving a company before an acquisition offer was completed, and then not being able to do anything to help salvage it.

What are the things you're really good at as a leader?

Recognising what I'm good at and admitting what I'm terrible at. Ensuring my team focuses on outcomes, not just outputs, and getting them to understand that I'd always prefer they push boundaries, even if it means I have to step in and firefight with them

Which areas do you need to improve on?

I need to be better at segmenting my day. Before I had kids, I used to make up for lost time by just working longer hours. It's no longer an option.

What's in store for the future of the business?

We recently presented our work on stage at OpenAI's annual Developer Day in SF, receiving overwhelmingly positive interest from the world's largest enterprises and leading LLM providers. We're working on two exciting projects you'll see soon. One integrates Genie into large, dev-focused cybersecurity companies, enabling it to automatically resolve flagged issues without user interaction. The other helps companies use our proprietary data pipeline to generate synthetic reasoning traces—
incredibly useful for any AI product requiring "reasoning." Fortunately for us, we're the only ones in the world who can provide that right now.

What advice would you give to other founders or future founders?

Hardwork and developing in-depth expertise is a given, but the reality is

that most advice by other successful founders are misguided due to survivorship bias. Luck plays a big part in success and the best way to get lucky is to keep rolling the dice over and over again. Don't give up, you just need to get lucky once.

And finally, a more personal question! We like to ask everyone we interview about their daily routine and the rules they live by. Is it up at 4am for yoga, or something a little more traditional?

I get up whenever my toddler wakes up, usually around 5am, and we start making breakfast together before doing the daily baby shark and Bluey theme song dance routine together. Prioritisation and to-do lists are always done last thing before bed - I'm a blue light non-believer. I'm also a recent F45 and Yoga convert, but that's only when I have time at lunch.

Yang Li is the cofounder and COO of Cosine, builders of Genie.



MADDYNEWS UK

The newsletter you need for all the latest from the startup ecosystem

JE M'INSCRIS

Article écrit par Yang Li