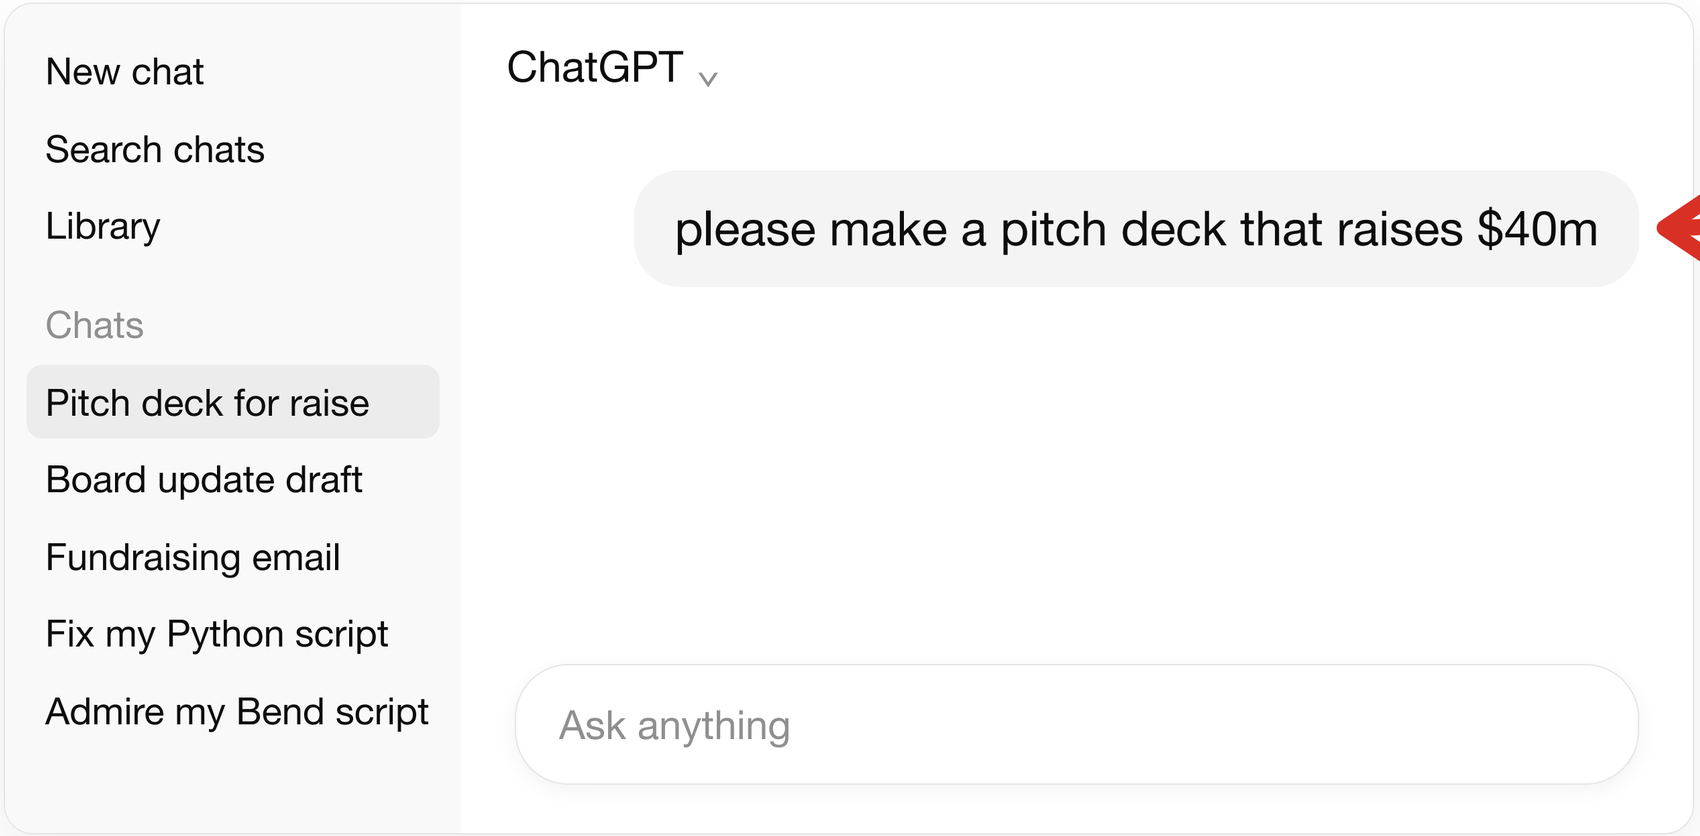
SupGen

A different route to AGI

Problem

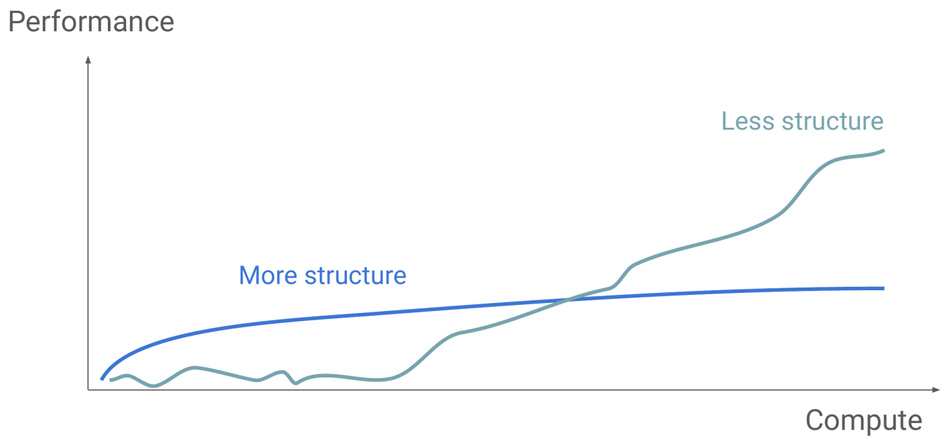
AGI isn’t here yet

That didn’t work



Solution

A new AI architecture



> “The story of deep learning is this: empirically, old simple methods, which were usually invented in the 80s and the 90s, when scaled up on very large clusters, work really well. That happened with computer vision, that happened with normal supervised learning, and it also happened with RL.”
>
> — **Ilya Sutskever**, co-founder of OpenAI and SSI, in a talk

We’ll scale Symbolic AI to GPU clusters using HVM and Bend

Figure: H. W. Chung, “Don’t teach. Incentivize.”, MIT EI Seminar, 2024, slide 9.

Business model

We sell useful intelligence



Elon Musk
@elonmusk · Jun 18, 2026

Anthropic has rightly focused on maximizing useful intelligence, which does not show up in benchmarks, but definitely shows up in revenue.

If research works, business is straightforward

Raising

Up to $8M

On SAFEs, at a $120M post-money cap

$100k minimum check

First come, first served: investors@higherorderco.com

All the money goes to **research**

This is a high-risk investment. You are investing in whatever comes out of our symbolic AI research; Bend will likely remain open source and free. Most early-stage investments return nothing, and we don’t yet have a clear long-term revenue path. We believe in our technology and our hustle, but that is no guarantee of success.

Team

The team

Victor Taelin, founder, CEO, created HVM, Bend and Kind

Vitor Chiarelli Neves, cofounder

Nicolas Abril, researcher

Lorenzo Battistela, researcher

Paulo Cavalcanti, researcher

Vision

The bitter lesson: the more humans hand-engineer an AI, the better it does at small scale, and the worse it scales. In the long run, the simplest method with the most compute wins. So we take that to the extreme. The simplest method there is: proof synthesis, run on massive compute.

That space holds every possible AI, including ones far better than anything a human could design, or even understand. The only question is how much compute it takes to find one. A learner like nanoGPT may fit in about 1,200 bits. Made by humans, it is surely far from the best. If the best is much smaller, an NVL72 can plausibly mine it.

It’s a big bet, on purpose, and it may fail. But if it works, it changes everything. Our reason to believe is scale: the same force that took GPT-2 from a useless autocomplete to the heart of the AI industry.