

What building an AI tool taught me about venture capital

By Isaac Kwa

ON PAPER, evaluating a startup looks algorithmic. Assess the problem statement. Size the market. Map the competition. Check the traction. Profile the founders. Run the numbers. It is the kind of structured, repeatable process that seems perfect for automation.

So, when I was first tasked with building an artificial intelligence (AI)-powered deal-sourcing tool for our fund, I assumed the hardest part would be the engineering.

I was wrong. What I found, through months of building and testing, is that venture capital (VC) resists automation not because investors are technophobes protecting their turf – quite the contrary – but because the job is far more nuanced than it appears from the outside.

Here is what I learnt about the gap between what AI promises and what the job actually demands.

The data problem

Let us start with the inputs. In VC, data is sparse, private, fragmented and frequently stale. The last funding round recorded in a database might be 18 months old. The most recent Accounting and Corporate Regulatory Authority filings tell you what the company looked like at the previous year's end.

These data points serve as snapshots of where the company was, not signals of where it is going. Real-time momentum lives in founder conversations and reference calls, not in external databases.

This is a problem quant hedge funds know well. Most sophisticated firms had to invent proprietary leading indicators from whatever public market data they could find.

In VC, you have less available information to work with, more qualitative variables, and brutally slow

feedback loops. A bad investment thesis in the public markets gets punished in weeks, but in VC, it might take years.

Then there is the problem of outliers. Investments that define a fund's vintage are, almost by definition, the ones no model would flag. Bitcoin's inception, the Covid-19 acceleration, the GPT moment. Black swans do not backtest.

You cannot train a model to identify the next paradigm shift on data from the last one. And given that the industry largely abides by the power law, professionals are counting on this handful of elusive investments to make good on the entire fund.

Perhaps most importantly, the signal that matters most in early-stage investing is often not quantitative at all. It is the story a founder tells about why this problem, why this solution, why now, why them.

A compelling narrative, one that reframes how you see a market, is worth more than any spreadsheet. Context beats data. Narrative moves markets.

The subjectivity problem

Investment judgement is not the application of a consistent rubric, but a craft developed over years of pattern recognition and hard lessons, and it is deeply personal. The same deal can be a pass for one partner and a conviction bet for another, and both can be right.

When building the tool, we had to make a fundamental decision – how much should we hard-code our firm's criteria, versus how much do we trust the general intelligence of a foundation model?

Hard-code too much, and the rigid filter might miss anything that does not fit a familiar pattern – which might be exactly the kind of contrarian bet that defines the best VC returns. Lean too heavily on the



In venture capital, one has less available information to work with, more qualitative variables, and brutally slow feedback loops. PHOTO: UNSPLASH

model's black-box reasoning, and we might lose explainability.

Then there is also the small data-set problem. How often does a fund actually do deals? Dozens over several years, maybe. That is not a training set – it is barely a sample. Older memos may also be anchored to paradigms that no longer apply. A thesis built on software-as-a-service-era assumptions, for instance, might not be directly relevant when it comes to evaluating AI-native companies.

The relationship moat

If the data and subjectivity problems could theoretically be solved

with enough engineering, the relationship problem cannot.

There is a paradox here: if your AI can see the deal, so can everyone else's AI. The information edge disappears the moment it becomes systematised.

The best deals in VC do not surface through data – they travel through trust. Founders choose which investors to call before the round is public. A co-investor saves a seat for a fund with a track record of being genuinely useful. A reference check is only as good as the relationships that make honest feedback possible.

These are not inefficiencies

waiting to be disrupted. They are features of a market built on asymmetric information, where access is earned and reputation is currency.

Relationships give access. And access, in this industry, is most of the job.

Where AI actually wins

None of this means that AI has no place in VC. In practice, it has already started to reshape parts of the workflow, and firms pretending otherwise will eventually fall behind.

AI can scan, flag and process opportunities very efficiently. One

clear application is qualifying inbound leads, ensuring fewer interesting companies fall through the cracks simply because no one had time to look.

The surrounding workflow is also being transformed: data room agents that summarise and surface critical information; deep research tools that compress market analysis from days to hours; Excel, PowerPoint and investment memo agents built around a firm's house style that help develop a first cut, meeting transcription and outreach automation.

These are not trivial gains. They free up human bandwidth for the parts of the job that matter most.

So where does that leave us?

Since I first embarked on this project, AI has already moved beyond the co-pilot stage. Parts of the workflow are no longer just augmented, but automated. As frontier models continue to reach new heights, maybe they will eventually creep into higher echelons of judgment.

And maybe that is fine. Maybe it is even a good thing.

As the leaders of my firm constantly remind us, investing is the first and easiest part of the job. Everything after – the board work, the difficult conversations with founders, the judgment calls in a down round, the long-term bet on a person – is where the real work lives.

If AI eventually handles the sourcing, screening and first-pass analysis, what is left is the part that was always most important anyway. Maybe that is not something to resist. Maybe that is the whole point.

The writer is an associate, investment at Vertex Ventures South-east Asia and India