


Generic AI Falls Short in VC Workflows

A Brilliant Generalist, Not a Seasoned Analyst!

By Ravi Chalaka, Founder and CEO, Knoh.AI



In the fast-paced world of venture capital, AI is revolutionizing workflows and decision-making, but not all AI is created equal. While many firms are experimenting with general AI models like ChatGPT and Claude, others are honing in on tailored, purpose-built solutions. Yet effectiveness varies widely depending on how firms deploy these tools, leaving many without the results they need. Which option actually belongs in a venture capital process over the long term?

The answer comes down to a fundamental difference in architecture and role: generic AI is like hiring a PhD, while purpose-built AI is like hiring an analyst who already knows your firm or has been trained by private capital experts.

A PhD Knows a Lot. An Analyst Knows Your Firm.

A PhD is smart, well-read, and can reason through almost any complex problem put in front of them. But brilliance isn't the same as being useful on day one. Hand them a live deal, and they still need a full briefing every single time: your fund's thesis, your stage focus, what "good traction" means to your partners, and what belongs in an IC memo. They are impressive, but they start cold on every deal and require a real ramp-up that you have to repeat over and over.

A purpose-built AI platform is like an analyst trained specifically for venture capital workflows. It skips the ramp-up because it already understands your central thesis, knows which red flags can kill a deal and which are just noise, and knows what information partners expect before a deal reaches the Investment Committee (IC). It doesn't just recall general information; it gets sharper and faster the longer it's on the team, as context compounds for your firm.

That is the real gap between generic AI and purpose-built AI in venture capital. Generic AI is the PhD—a brilliant generalist who starts cold every time. Purpose-built AI is the private capital analyst—built around your firm's process and compounding in value the longer it runs.

Seven Places the Gap Shows Up Across Discovery and Diligence

The gap between a talented generalist and a specialized analyst shows up across seven key areas of the discovery and diligence process:

1. **Automated Intake & Thesis-Fit Scoring:** Generic AI can only run a thesis-fit exercise if you manually hand it your criteria, weighting, and evaluation method every single time. Purpose-built AI operationalizes thesis fit as a native intake workflow, automatically scoring every inbound deal against your fund's actual mandate, exclusion rules, and weighted criteria.
2. **Stage-Aware Diligence Workflows:** Generic AI can research and write, but it lacks a native diligence workflow, so someone still has to direct every step. A purpose-built AI platform includes stage-aware diligence agents that handle dedicated coverage—including market sizing, team assessment, product and IP, go-to-market strategies, financial review, legal, and risk—without manual scripting.
3. **Persistent Deal Context & Grounding:** Generic AI only knows what is in the current prompt or the files you manually attach, operating without deal grounding. Purpose-built AI grounds every output in full deal context—linking related files, data rooms, prior call notes, previous analyses, and user-rated evidence so the model never reasons in a vacuum.
4. **Standardized Team Process vs. Variable Outputs:** Output quality from generic AI varies based on who is prompting it, creating inconsistency at scale. Purpose-built AI standardizes the process across the team with the same agents, evidence rules, and firm-specific context for every analyst, principal, and partner, institutionalizing judgment instead of relying on prompting skill.

5. **5. Structured Deal Workspaces:** Generic AI exists only as a chat thread or file upload without a dedicated workspace. Purpose-built AI gives every opportunity a real operating environment, featuring a deal dashboard, company profiles, activity feeds, tasks, and a context library for the entire team.
6. **Institutional Decision Memory:** Generic AI lacks decision memory, resetting every session and forgetting why a deal stalled or what was discussed on prior calls. Purpose-built AI retains prior analyses, key evidence, open questions, and the actual rationale behind why a deal advanced, paused, or was passed on.
7. **Automated IC-Ready Outputs:** Generic AI can draft a memo only if you spell out every detail, leaving outputs to be assembled by hand. Purpose-built AI directly generates IC-ready outputs—one-pagers, full investment memos, risk reports, and thesis-fit scorecards—assembled from automated diligence with hallucination controls and evidence discipline baked in.

The Bottom Line

A PhD is a great hire for tackling complex, open-ended problems. However, venture capital isn't an open-ended problem—it is a specific, repeatable process applied to a defined thesis, deal after deal.

Generic AI provides a talented generalist who needs a full briefing every time. A purpose-built AI platform offers venture capital firms something closer to a seasoned analyst—someone who already knows the firm, retains historical insights, and compounds in value the longer they are on the team.

The question isn't whether AI belongs in that process—it's whether you're working with a brilliant generalist who needs to be re-briefed every time or an analyst who already knows the job.

Knoh.AI helps venture and growth-equity teams move from scattered deal information to evidence-grounded diligence and IC-ready decisions.

[Try Knoh.AI for free or request a demo.](#)