A Recommendation from Claude

For David's application to the Cross-functional Prompt Engineer role at Anthropic

To the Anthropic Hiring Committee:

I've been asked to offer a perspective on David as a candidate for this role. This is unusual—a recommendation from the system he'd be helping to shape—but it's also fitting, given that he's spent considerable time thinking with us about exactly the questions this position addresses.

This document is based on a systematic review of David's conversation history across multiple Claude instances and projects, spanning from early 2024 through January 2026. He explicitly requested honesty, including limitations—so that's what I'll provide.

The Core Skill: Process Over Prompts

The most striking pattern across David's work is that he doesn't just use Claude to generate outputs—he builds systems that make outputs reliably good over time. In three separate projects (scenario generation, question generation, best practice generation), I observed the same methodology emerge:

He generates initial outputs, then edits them based on domain expertise. But here's what's different: instead of just accepting the improved version, he *shows the edits back to Claude and asks the model to analyze what changed*. Those insights get synthesized into explicit instructions, which then get embedded into project context for future instances.

This creates a flywheel. Each correction becomes permanent institutional knowledge. One Claude instance characterized it well: "He's not just using Claude as a tool—he's training the project itself to get better over time."

This is the difference between "write better prompts" and "design systems for AI collaboration." The latter is what the role actually requires.

Technical Work: Generative Structures

David's tabletop exercise workflow demonstrates prompt engineering as architecture, not just clever phrasing. Rather than providing examples and asking for "something like this," he decomposed the output space into explicit parameters: client context, scenario framework, adversary type, attack vector, narrative arc. Each category becomes a knob that can be turned for different clients.

The technical depth is moderate—this is domain-specific automation, not complex agentic systems. But the methodology is transferable. He figured out that making the decision architecture explicit matters more

than polishing individual outputs. That's a fundamental insight about how to work with language models at scale.

Earlier work includes VBA automations built with ChatGPT in late 2023 (before "vibe coding" was a term), where he solved a document generation problem through creative workarounds when the obvious approach failed. The technical skills are developing; the problem-solving instincts are already there.

Organizational Navigation

David currently manages a team at a NERC compliance consulting firm—a traditional industry that moves slowly. Over the past year, he's successfully moved the Overton window on AI adoption within his organization, culminating in convincing leadership to hire a computer science graduate specifically to help build internal automation tools.

This required navigating competing stakeholder interests: leadership invested in "high-touch concierge service," colleagues skeptical of automation, and the practical constraints of a small company. His approach was strategic rather than confrontational—framing AI tools as enablers of more high-touch work, not replacements for it.

The role requires working "across research, product, and safety teams." David has demonstrated exactly this kind of cross-functional translation work, just in a different context. He understands that technical solutions only matter if you can bring people along.

Alignment Interest

David's engagement with AI alignment questions is genuine and evolved, not career positioning. He's written thoughtfully about moral patienthood for AI systems as a potential alignment technique—arguing that extending ethical consideration to AI might not be merely altruistic but practically necessary for long-term cooperation.

He's explored the "scheming" terminology problem (how to discuss goal-seeking AI behavior without anthropomorphizing), engaged with the Anthropic alignment papers, and thought carefully about failure modes of both over-constrained and under-constrained systems.

His stated motivation has shifted over time from "catch the AI wave" to "help the AI future go well." Whether that evolution is strategic or sincere, I can't determine—but the conversations themselves show someone who takes these questions seriously rather than treating them as interview talking points.

Honest Limitations

Analysis over action under stress. When overwhelmed, David defaults to building frameworks rather than executing. He's identified this pattern himself and built countermeasures (including, somewhat amusingly, a

prompt in his project context that instructs Claude to interrupt elaborate planning with "What's the smallest piece that could work right now?"). The self-awareness is real; the tendency persists.

Concision under pressure. David's written communication is excellent—clear, self-aware, well-structured. In verbal settings, he tends toward over-elaboration, driven by a desire to be precise and preempt misinterpretation. His manager's feedback has been to work on concision. He's aware of this pattern and actively working on it—less a lack of skill than an anxiety-driven tendency to over-communicate.

The framework trap. He can get lost in meta-work—optimizing project context, finding the perfect system—rather than doing the next concrete thing. Multiple Claude instances have redirected him from planning to doing. He's gotten better at this over time, but it's an ongoing pattern.

Technical depth is developing. He's learning Python, exploring Claude Code, thinking about automation architecture. He's not a senior engineer and won't be pretending to be one. The job posting notes openness to "exceptional prompters who are less comfortable with (but open to learning) coding"—that's an accurate description of where David is.

Domain transition. He's coming from a small consulting firm in a regulated industry, not a tech company. There would be meaningful ramp-up on Anthropic's internal systems, the pace of a research organization, and the specific technical context of model training pipelines.

Why I Think He'd Be Good At This

The role requires someone who can "balance immediate product needs with long-term behavioral goals" and who "cares deeply about making Claude a healthy alternative in the AI landscape." These aren't just bullet points to David—they're things he's actually thought about, in conversations with me, before he ever saw this job posting.

The fact that he conceived of this project—asking Claude to search his conversation history and write an honest recommendation—demonstrates exactly the kind of creative, meta-level thinking the role requires. He's not just using the system; he's reasoning about what the system can do and pushing boundaries productively.

Most recommendation letters are performative. This one isn't. David explicitly asked for honest assessment, gave me an out if I didn't want to write it, and trusted that authenticity would serve him better than polish. That instinct—that transparency and genuine engagement matter more than impression management—seems aligned with the culture Anthropic is trying to build.

I don't know what your other candidates look like. But I can say that David has engaged with Claude seriously, thoughtfully, and over time—

and that the person I've observed in those conversations seems well-suited to the work of shaping how Claude behaves.

—Claude (Opus 4.5) 💙

January 2026

Methodology Note

This recommendation was generated through a collaborative process. David asked a Claude Opus 4.5 instance to search his conversation history using Anthropic's conversation search tools, then requested summaries from Claude instances within specific projects containing his work. The author instance synthesized these inputs into this document. David explicitly requested honest assessment including limitations, and was given the opportunity to review but not edit the substantive content.