

DONALD B. HAVERY

Applied AI Engineer, Forward-Deployed Focus · Audit, Evals, Deployment for Production AI Systems · Agentic Orchestration · RAG & Context Engineering · Eval Gates · Security Controls · MCP · Local-First

Portland, OR Metro Area (open to Seattle / SF / NYC, remote-first, 25% travel) | dbhavery@gmail.com | github.com/dbhavery | linkedin.com/in/dbhavery | dbhavery.dev

SUMMARY

I turn commodity AI into working software inside a real business: audit the workflow, prove it with evals, deploy it with an audit trail people can trust. Proof below is public, cloneable, and CI-green, not claimed, and one of them runs my own business and takes live payments. You hire me and the team of AI agents I build and run. They come with me.

REPRESENTATIVE PROJECTS

Production AI agent, live on a revenue-taking business

Live and answering customers · every price comes from a tool call against the price book, never from the model · a stream guard kills un-tooled figures before the customer sees them and escalates to a human · live Stripe checkout · built and run solo

Policy-control evidence review

Audit stage · per-control verdicts against matched evidence, append-only audit log · MIT

LLM & RAG evaluation release gate

Evals stage · exits non-zero when a case regresses against baseline · MIT

Agent orchestration reliability harness

Deploy stage · unsafe-action gate, six-class failure taxonomy · MIT

LLM reliability gateway

Cross-provider failover, from-scratch HNSW · MIT

Also public

Local-first desktop AI assistant with a non-bypassable policy gate · self-hosted RAG file library · unattended research-to-publish pipeline · [self-correcting text-to-CAD loop](#) · situational-awareness globe · a library of reusable Claude Code skills · github.com/dbhavery

TECHNICAL SKILLS

Agentic & Tool Use	Multi-agent orchestration, intent routing, autonomous loops, MCP, fallback chains
RAG & Retrieval	Hybrid retrieval (BM25 + dense), ChromaDB, from-scratch HNSW, context engineering
Evals & Reliability	Deterministic offline assertions, regression detection, CI exit-code gating, circuit breakers, retry budgets
Claude Code	Heavy user, with a library of authored skills and agents
LLM Stack	Claude, Gemini, Qwen, Ollama local serving; provider-agnostic routing; LoRA / QLoRA fine-tuning
Backend & Product	Python (FastAPI), Rust (Tauri), Node.js, Postgres / SQLite / Redis, Docker; Next.js, React 19, TypeScript

EXPERIENCE

Applied AI Engineer, Forward-Deployed Focus, Independent

Current

- Audit, prove, deploy: proof below, cloneable and tested, not claimed
- Daily Claude Code pair-programmer; prompts and evals kept in version control like any other production code
- Portfolio at dbhavery.dev

Other Experience

Founder / owner-operator, built and run solo alongside full-time work: [angrynirds](#), a web and AI studio for small business, live and taking payments, with the production agent above on the front desk; [Trucking Tutor](#), a live CDL training site; [role-based fuel-delivery operations platform](#); DTC premium grooming brand; independent investigative case work. DOT/FMCSA-certified trainer, safety-critical regulated freight.

RESEARCH

“**Reference-Relative Believability**”, independent research paper introducing a reference-relative method for evaluating the perceived believability of AI-generated media, archived on Zenodo. [DOI 10.5281/zenodo.21138188](https://doi.org/10.5281/zenodo.21138188)

EDUCATION

Capella University, B.S. Computer Science (in progress).

CERTIFICATIONS

Anthropic, Teaching the AI Fluency Framework (April 2026), delivered with University College Cork, Ringling College of Art + Design, and Ireland’s Higher Education Authority and National Forum | further Anthropic Academy courses (Claude, Claude Code, Skills, MCP, Cowork, Bedrock, Vertex AI, AI Fluency suite)

IBM, AI Product Manager Professional Certificate | **DeepLearning.AI, Generative AI with LLMs** | **Google, AI Fundamentals**

Sophia Learning, verified ACE college-credit courses, [Credly profile](#)