

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. 38 repos, 3,300+ commits, 5,000+ tests authored.

REPRESENTATIVE PROJECTS

Policy-control evidence review

Audit stage · per-control verdicts against matched evidence, append-only audit log · 40 tests · MIT

LLM & RAG evaluation release gate

Evals stage · exits non-zero when a case regresses against baseline · 54 tests · MIT

Agent orchestration reliability harness

Deploy stage · unsafe-action gate, six-class failure taxonomy · 45 tests · MIT

LLM reliability gateway

Cross-provider failover, from-scratch HNSW · 118 tests · 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 · 25+ 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, 3,300+ commits authored through it, 25+ skills and agents of my own
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: role-based fuel-delivery operations platform, DTC premium grooming brand, independent investigative case work, AI-engineering services collective. DOT/FMCSA-certified trainer, safety-critical regulated freight.

PUBLICATIONS

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

EDUCATION

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

CERTIFICATIONS

Anthropic Academy, 15 courses complete (Claude, Claude Code, MCP, Cowork, Bedrock, Vertex AI, AI Fluency)

DeepLearning.AI, Generative AI with LLMs | Google, AI Fundamentals

Sophia Learning, 30 verified ACE college-credit courses, Credly profile