

Hugh O'Shaughnessy

Mountain View, CA | (650) 625-7191 | hughoshaugh@gmail.com | github.com/hughosh | linkedin.com/in/hugh-o-shaughnessy

Education

University of Wisconsin–Madison

Madison, WI

B.S. Computer Science, May 2026

Selected coursework: Foundation Models, Database Systems, Big Data Systems, Data Science, Computer Architecture, High-Performance Computing, Operating Systems, Compilers, Computer Networks, Algorithms, Data Structures.

Experience

ShiftIQ

2026–Present

Research Scientist Intern

Remote

- Founding research team building governed AI infrastructure around the Causal Execution Engine (CEE).
- Working at the boundary between mechanistic interpretability and engine design: probing residual-stream representations in open-weight models to identify which operations the engine should preserve and govern.
- Running PyTorch hook experiments (activation patching, linear probes, causal interventions) to map which operations a model can reliably execute under steering, feeding results into the engine's design decisions.

Confidential Single-Family Office

2026–Present

Independent Engineer, Agentic Workflow Platform

Remote

- Architecting an agentic workflow platform for the office's investment-research and due-diligence operations; privacy, compliance, and auditability are first-class design constraints rather than retrofits.
- Implementing the canonical compliance path: permissioned retrieval over an ACL-tagged document index, pre- and post-inference classifier layers (PII, prompt-injection, policy violations), and a constitutional system prompt bounding the agent's action space.
- Standing up the oversight layer (replayable traces, tool allowlist with deterministic policy gates, human-in-the-loop checkpoints for high-stakes actions) so the platform can be reviewed end-to-end by compliance.

Summer AI Lab (SAIL) Fellow, UW–Madison CDIS

Jun–Aug 2025

N+1 Institute, OpenAI-Funded, Inaugural Cohort

Madison, WI

- Selected for the inaugural Summer AI Lab cohort; received a \$1,500 research scholarship and OpenAI API credits.
- Co-built a personalized agentic news aggregator on a 3-person research team: web-scraping ingestion, RAG retrieval, and source-conflict adjudication.
- Implemented semantic-similarity and bias-detection modules feeding a real-time credibility-scoring pipeline.

Projects

5-Stage Pipelined MIPS Processor

Verilog, ModelSim

- Designed a 32-bit, 5-stage MIPS CPU in Verilog (IF/ID/EX/MEM/WB) with a register file, single-cycle ALU, and branch resolution in the EX stage.
- Engineered hazard-detection and forwarding units that bypass pipeline registers to eliminate stalls on RAW and load-use hazards.
- Built ModelSim testbenches covering control-flow edges, data dependencies, and the supported MIPS instruction set.

MiniSpark: Parallel Data-Processing Framework

C, Unix

- Multi-threaded analytics engine in C modelled on Apache Spark; runs parallel pipelines through a custom thread pool and DAG scheduler.
- Mutex-guarded shared state and a low-overhead monitor thread tracking per-stage performance and queue depth.
- Implemented core operators (map, filter, join) with stage-boundary materialization and dependency-aware scheduling.

Student Enrollment DBMS

SQL, C++

- Normalized SQL schema (students, professors, courses, enrollments) with referential integrity across join tables; ER model validated against query-performance tests.
- C++ application layer exposing CRUD operations and parameterized query helpers over the database.

Technical Skills

Languages: Python, TypeScript, SQL, C, C++, Verilog, Java

AI / LLM: PyTorch, Hugging Face Transformers, LangGraph, OpenAI / Anthropic APIs, pgvector, RAG, agentic workflows, model evaluation, post-training quantization

Data / Backend: PostgreSQL, MongoDB, Snowflake, Elasticsearch, Apache Spark, Airtable, Node.js, REST APIs, ETL / data pipelines

Web / Infrastructure / EDA: React, Next.js, Linux, Docker, AWS, CI/CD (GitHub Actions), Git, ModelSim, Quartus, Xilinx/AMD Versal