



Andrew Grainger, PhD

Applied AI Systems & Solution Architecture

Agent workflows · retrieval and evaluation · business integration

Rowley, MA

agrainger221@gmail.com

(408) 438-3114

linkedin.com/in/andrew-grainger

github.com/atg221

I architect multi-agent AI solutions, develop rapid prototypes and proofs of concept (PoCs), and turn them into working business integrations. At Cypher Lab, I built orchestration and custom automation infrastructure connecting scientific evidence, specialist agents, review and publication. My work includes personalized vectorized and knowledge-graph-enabled retrieval-augmented generation (RAG), persistent agent memory, and hands-on Google Cloud / Gemini Enterprise Agent Platform experience and multi-agent solution design. A molecular-genetics PhD and published imaging-AI research anchor my evaluation and scientific judgment; commercial experience keeps the engineering focused on real user needs.

PROFESSIONAL EXPERIENCE

Co-Founder & AI Engineering Lead · Cypher Lab

Aug 2025 – present

- ◆ Used rapid prototyping and proof-of-concept development to build Cypher Lab's multi-agent workflow from market intelligence and positioning through specialist-agent content, technical review and Astro distribution. Designed the orchestration and custom automation infrastructure, integrating teammates' shared skills with my own components; documented handoffs and human release decisions.
- ◆ Built a scientific and product-document retrieval layer: corpus staging, passage classification, keyword/vector search, optional reranking and graph context, source-bearing results, regression questions and release checks. The September 2026 scientific database holds 7,715 papers and 216,789 passages; a known-document exam returned the expected item in the first ten results for 35 of 40 queries.
- ◆ Defined agent-to-tool contracts across client-scoped data, team knowledge, review and publishing services. Authored the OpenAlex and LeadLander MCP servers in Node.js; integrated identity-gated CLI services with structured results so a failed call can be handled explicitly.
- ◆ Built and deployed Astro publishing and resource-hub systems, including the PCS Astro/Sanity resource hub live since June 2026. The architecture turns reviewed scientific and commercial material into resources people can find and use.

Senior Manager of Marketing Operations · Precision Cell Systems / ORFLO Technologies

2024 – Oct 2026

- ◆ Led positioning, campaigns and sales enablement for Moxi cell-analysis and Singulator sample-preparation products. Translated customer questions and technical evidence into distinct product messages and reviewed materials.
- ◆ Put AI-assisted website inquiry intake into PCS's live CRM. A model classifies requests; deterministic code owns identity matching, deduplication and writes. I tested record creation, readback and deletion before release, then repaired mismatches and missed writes found after deployment.
- ◆ Built product-support knowledge bases from historical support questions and technical documents, plus a working Slack access prototype through self-hosted Docker/n8n. The prototype makes source material available through a familiar support channel.

ARCHITECTURE IN USE

From inquiry to system of record PCS · deployed

The CRM intake is a practical agent boundary: model classification helps route an inquiry, while code controls identity and writes. Post-release checks exposed incorrect matches and missing records; after repairs, a fresh 21-entry window produced four new opportunities and nine general-information prospect records with zero logged errors. This experience informs how I design human review, rollout and recovery for other business workflows.



Andrew Grainger, PhD

Applied AI Systems & Solution Architecture

Reusable architecture · governed memory · research rigor

Rowley, MA

agrainger221@gmail.com

(408) 438-3114

linkedin.com/in/andrew-grainger

github.com/atg221

INDEPENDENT ENGINEERING · VOIDWRIGHT

VOIDLOOM · VOIDMILL · VOIDSTEAD RAG / MEMORY / ORCHESTRATION

Designed personalized RAG and persistent-memory architecture for multi-agent solutions. Built VOIDLOOM's working SQLite memory, vector/text recall, temporal knowledge graph and session hooks, plus VOIDMILL's source-bearing ingestion and retrieval. Designed VOIDSTEAD's model-agnostic orchestration framework and implemented its foundational substrate. These custom components support iterative prototypes and evaluation.

Google Cloud / Gemini Enterprise Agent Platform PLATFORM EXPERIENCE + SOLUTION DESIGN

Hands-on Google Cloud / Gemini Enterprise Agent Platform experience and multi-agent solution design. Connected managed agent-runtime, tool-integration and governance patterns to my locally built RAG, knowledge-graph and persistent-memory components in cloud architecture designs, with explicit evaluation and cost controls.

EARLIER EXPERIENCE

Field Application Scientist · ORFLO Technologies

2023 – 2024

- ◆ Demonstrated and supported Moxi instruments, joined scientific sales calls and resolved customer technical escalations. Made Moxi training videos and wrote Moxi V IQ/OQ qualification protocols for a customer.

Post-Doctoral Research Associate, Radiology & Medical Imaging · University of Virginia

2020 – 2023

- ◆ Developed and evaluated a deep-learning U-Net for CT fat segmentation with DICOM-to-NIfTI preparation and template augmentation in TensorFlow/Keras. Published Dice 0.94 against expert annotation; coordinated study methods with imaging and clinical collaborators.

Graduate Researcher, PhD · University of Virginia

2014 – 2019

- ◆ Investigated cardiometabolic disease through genetic analysis and laboratory validation; coordinated studies and mentored junior researchers.

HOW I BUILD

Data and evidence METHOD

Preserve source and passage identity through ingestion. Use hybrid search and graph context when useful, then test retrieval on known documents and record failure modes. Keep generated claims open to specialist and human review.

Model development RESEARCH / INDEPENDENT

My CT work trained and evaluated a task-specific segmentation model against expert annotation. In independent work, I built a labeled text corpus and PyTorch dataset for a 15-class classifier, compared pretrained encoders and tuned a language model with 8-bit LoRA. Both projects required careful task design, data preparation and evaluation.

EDUCATION AND TECHNICAL RANGE

PhD, Biochemistry & Molecular Genetics · University of Virginia, 2019

MS, Biology · UC San Diego, 2014 | BS, Molecular Biology · UC San Diego, 2013

10 peer-reviewed publications · ORCID 0000-0003-4131-2775

AI solution architecture Agent workflow design RAG and hybrid retrieval Retrieval evaluation Human review gates Python
Node.js / MCP / REST APIs SQLite and graph context Docker / n8n / AWS Lightsail Cloudflare / Vercel / Railway / Supabase
TensorFlow / Keras Astro / Sanity