

James Scanga

jas1337@pitt.edu | (724) 715-2132 | linkedin.com/in/james-scanga-05a707381 | github.com/jscanga

EDUCATION

University of Pittsburgh

B.S. in Computer Science & Data Science (Double Major), GPA: 3.5/4.0

Pittsburgh, PA

Aug 2025 – May 2028

- Completing both degrees on an accelerated three-year track via AP and transfer credit

EXPERIENCE

Federated Hermes

Software Engineering Intern

Pittsburgh, PA

Jun 2026 – Aug 2026

- Shipped production data pipelines on **Databricks** (Python, SQL) processing large-scale financial datasets, with CI/CD and version control managed through **Azure DevOps** in a regulated environment
- Profiled and tuned **15+** underperforming SQL queries across tables ranging 10M–500M+ rows via execution plan analysis, index design, and join restructuring, cutting average query latency **~40%** (~12s to ~7s)
- Led development of Paperclip (Next.js, TypeScript), a role-scoped document search tool cutting per-query AI token costs **99.2%** (\$0.18 to \$0.0014 on a 68-page doc); won **1st of 8 teams (64 interns)** at Federated Hermes' hackathon, now under evaluation for production

Keygrad

Founder & Full-Stack Engineer

Pittsburgh, PA

Aug 2025 – Present

- Founded and shipped keygrad.com, a live SaaS platform (Next.js 14, TypeScript, Supabase, Vercel) that maps **43,000+ course equivalencies** across **1,800+ source colleges** into **14 Pennsylvania universities**
- Engineered a plug-in ingestion pipeline and runtime ETL layer that resolves institution and course identity across sources with no shared key, normalizing inconsistently formatted scraped records into a queryable schema without requiring code changes to onboard new institutions
- Built end-to-end authentication, residency-aware pricing, and server-side rendered search; deployed to production on Vercel with a custom domain
- Extended course-matching with a semantic similarity layer (**Claude API**) to catch equivalent courses missed by exact/substring matching across differently-phrased titles between institutions

West Deer Web Design

Co-Founder

Pittsburgh, PA

Jun 2024 – Present

- Co-founded a freelance web development business serving local small businesses end-to-end (scoping, design, development, deployment), growing to **10+ clients**, **\$2,000+** in total revenue, and **\$490** in recurring monthly revenue

PROJECTS

Paperclip | *Next.js, TypeScript, Vercel, Claude API*

- Designed a role-scoped, mobile-responsive document search tool (React) indexing SharePoint documents into small searchable passages, enforcing existing access permissions so retrieval never surfaces content a user isn't cleared to see
- Cut per-query AI token costs **99.2%** versus full-document scans (~60k tokens/\$0.18 to ~468 tokens/\$0.0014 on a 68-page document) by retrieving only relevant passages
- Built a multi-agent routing layer where Claude Haiku classifies query complexity and dispatches to the right-sized model (Haiku, Sonnet, or Opus), balancing cost against response quality
- Estimated **\$3.9M+** in annually recovered employee time (60-person survey); won **1st of 8 teams (64 interns)** at Federated Hermes' hackathon, now under evaluation for production

Chess Blind Spots Analyzer | *Python, Flask, SQLite, Stockfish, REST APIs*

- Built a full-stack chess analysis web app that ingests a player's Lichess/Chess.com history via REST APIs and scores every move by engine centipawn loss, surfacing recurring weaknesses by game phase, color, and severity
- Designed a position-and-depth-keyed evaluation cache making repeat analyses **99.9% cache-served (~130x faster)** behind a background-threaded pipeline with configurable depth (15–700ms/position)
- Applied the tool to **100 of my own games (2,427 moves)** and found endgame play, not opening theory, was the dominant source of rating loss (**443.3 vs. 164.3** avg centipawn loss), with blunder rate (9.4/100 moves) outweighing all other error types combined

TECHNICAL SKILLS

Languages: Python, SQL, TypeScript, JavaScript, Java, C++, HTML/CSS

Frameworks/Libraries: React, Next.js, Flask, Node.js, Tailwind CSS

Data & AI/ML: Databricks, PyTorch, Pandas, NumPy, Scikit-learn, Claude API, Claude Code, GitHub Copilot, LLM Integration, Semantic Search, REST APIs, ETL Pipelines

Databases/Infra: PostgreSQL, SQLite, Supabase, Vercel, Git/GitHub, VMware ESXi, Linux, CI/CD

Other: Data Structures & Algorithms, OOP, Agile/Scrum, Financial Data Analysis