Andrew Zagula

(908) 392-6724 | andrewzagula800@gmail.com | github.com/andrewzagula | andrewzagula.com

Education

California Institute of Technology — Pasadena, CA
B.S. in Computer Science — *Expected Jun 2030*

Bridgewater-Raritan High School — Bridgewater, NJ
Salutatorian | GPA: 5.0 | SAT: 1570 — *Sep 2022 – Jun 2026*

- **Awards**: 2x USAMO Qualifier, 5x AIME Qualifier, USACO Gold
- **Coursework**: Calculus III, Differential Equations, Linear Algebra, AP Calculus BC, AP Statistics, AP Physics C, AP Chemistry, AP Biology, AP Macroeconomics, AP Computer Science A, Data Structures (Rutgers University)

Experience

Alumhub — Bridgewater, NJ
Co-Founder — *Mar 2025 – Present*

- Founded student–alumni social networking platform; led team of 10+ in development and growth, scaling to 5K+ users in 250+ U.S. high schools; secured $50K+ in cloud and platform credits from Google, Vercel, and Notion
- Led development of full-stack Next.js architecture with AI user-matching engine, combining structured LLM outputs and deterministic fallback ranking, real-time Firebase messaging, and Stripe billing with webhook feature gating

University of California, Berkeley — Remote
Student Researcher — *Jul 2024 – Dec 2025*

- First-author AutoAdv (NeurIPS Lock-LLM 2025), an adaptive adversarial prompting pipeline for multi-turn LLM jailbreaking; achieved 99% attack success rate on Qwen3-235B with 24% increase over baseline within $\leq$6 turns
- Designed 2-phase system-prompting framework for Grok 3 Mini; developed 4 score-driven temperature adjustment strategies and pattern memory logging to capture 28 successful jailbreak techniques across attacks
- Integrated provider-agnostic API routing and 10-worker parallel execution with per-turn tracking of token usage, latency, and cost; created GPT-4o mini scoring function assessing 3 success indicators for jailbreak detection

Boston University — Boston, MA
Research Intern (RISE) — *Jun 2025 – Aug 2025*

- Co-author BabyVLM-V2 (CVPR 2026), a developmentally inspired framework for sample-efficient vision-language models; pretrained 1.4B BabyLLaVA-V2 on 1M+ multimodal pairs and instruction-tuned on 110K+ samples
- Constructed DevCV Toolbox, a 10-task benchmark adapted from NIH Baby Toolbox; tasks generated from 478 hours of infant-perspective (SAYCAM) video data to assess spatial reasoning, memory, and vocabulary in VLMs
- Designed OpenCV video processing pipeline for frame sampling and sliding-window segmentation; implemented YOLO for object detection and tracking, filtering 3.5K+ candidate clips to 2.2K+ ($\sim$63%) usable samples via Label Studio

Projects

PaperTrail | *FastAPI, Next.js, TypeScript, OpenAI API, LangChain, LangGraph, ChromaDB*

- Built self-hosted arXiv research assistant supporting LLM-guided ranking for paper discovery and 6-part structured paper breakdowns, powering multi-paper comparison, research ideation, and starter code generation
- Developed section-aware LangChain retrieval with GPT-4o mini for citation-grounded Q&A and 3 GPT-4o LangGraph multi-step workflows, backed by SQLite for persistent storage and ChromaDB for chunked vector search

AUDIT | *Python, OpenAI API, ChromaDB, SQLite, Typer, PyInstaller, Node.js*

- Built CLI security scanner for detecting vulnerabilities in code repositories using retrieval-augmented LLM analysis over 50-pattern CWE/OWASP-aligned knowledge base to produce structured JSON reports
- Architected 5-stage scanning pipeline spanning candidate extraction, heuristic prefiltering, ChromaDB vector retrieval, parallel execution, and incremental SQLite caching, reducing LLM-bound candidate volume by $\sim$70%

Technical Skills

Languages: Python, C++, JavaScript, TypeScript, Java, SQL
Frameworks: React, Next.js, Node.js, FastAPI, Flask
Developer Tools: Git, Docker, PostgreSQL, Firebase, SQLite, ChromaDB, OpenAI API, Stripe
Libraries: PyTorch, NumPy, Pandas, scikit-learn, OpenCV, Matplotlib, Tailwind CSS