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, Y Combinator Startup School 2026
- **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 — *Jan 2025 – Mar 2026*

- Co-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 platform credits from Google, Vercel, Notion, and others
- Led full-stack Next.js development; built user-matching engine with schema-validated LLM generation and deterministic fallback ranking, real-time Firebase messaging, and Stripe billing with pro-tier feature access

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

- First-authored AutoAdv (NeurIPS Lock-LLM 2025), an adaptive adversarial prompting framework for black-box multi-turn LLM jailbreaking; achieved 99% attack success rate on Qwen3-235B, improving 24% over 6 turns
- Implemented two-phase prompt rewriting framework for Grok 3 Mini; developed temperature manager adjusting LLM sampling via 4 strategies, and pattern manager logging 28 effective rewriting techniques to guide future attacks
- Integrated provider-agnostic API routing and 10-worker parallel execution with per-turn tracking of token usage, latency, and cost; designed GPT-4o mini scoring function evaluating 3 success indicators for jailbreak classification

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

- Co-authored 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
- Developed OpenCV pipeline with frame sampling and sliding-window segmentation; applied YOLO for object detection and tracking, filtering 2.2K+ usable clips (∼63% yield) from 3.5K+ candidates via Label Studio

Projects

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

- Built local arXiv research assistant with arXiv URL/PDF ingestion, LLM-generated multi-query paper discovery, structured paper breakdowns, and 3 LangGraph workflows for comparison, ideation, and implementation planning
- Developed provider-agnostic RAG pipeline using LangChain for multi-turn Q&A with section-level citations; integrated ChromaDB top-5 retrieval and SQLite storage for papers, chats, and generated research artifacts

AUDIT | *Python, OpenAI API, ChromaDB, SQLite, Typer, PyInstaller, npm*

- Built CLI security scanner to detect vulnerabilities in repositories with RAG over 50-pattern CWE/OWASP knowledge base, generating JSON reports with file/line evidence, confidence scores, and remediation guidance
- Engineered canning pipeline with high-signal candidate extraction, heuristic prefiltering, ChromaDB semantic retrieval, parallel LLM review, and content-hash SQLite caching, reducing LLM-reviewed candidates by ∼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, Hugging Face, NumPy, Pandas, scikit-learn, OpenCV, Tailwind CSS