

Mohammed Nihad

(917) 365-0877 | mnihad1107@gmail.com | github.com/mnihad000 | [LinkedIn](#) | [Portfolio](#)

EDUCATION

City College of New York

August 2024 - May 2028

GPA: 3.93

- Bachelor of Science in Computer Engineering
- **Awards:** 5x Hackathon Winner, 3x Dean's List
- **Relevant Coursework:** Data Structures, Algorithms, Software Engineering, Software Design, Circuit Analysis,

Cornell University, Break Through Tech AI Program

March 2025 - May 2026

- Selected from a competitive applicant pool (~20% acceptance rate) to complete ML and agentic AI coursework

Skills

- **Languages:** Python, C++, TypeScript, Java, MATLAB, SQL
- **Frameworks:** FastAPI, Spring Boot, React.js, React Native, PyTorch, scikit-learn, Next.js
- **Tools:** Git/ GitHub, AWS, Docker, Figma, PostgreSQL, MongoDB, REST APIs, Kubernetes. Claude Code, Codex

EXPERIENCE

Bloom Energy, Software Engineering Intern

June 2026 – August 2026

- Incoming SWE intern Building scalable backend tooling and automated data pipelines for real-time monitoring of 30,000+ distributed energy systems, leveraging Kubernetes, Grafana, and FastAPI

Buildify, Software Engineering intern

December 2025 – March 2026

- Improved 30-day returning users by building an agentic Retention Autopilot that ingests analytics, segments users, and launches approved A/B-tested push + rewards campaigns, cutting campaign setup time by 30%.

STEMKasa, Software Engineering Intern

June 2025 – August 2025

- Led 3 interns to rebuild a 17-file monolith as a MERN app, implementing component-based design, REST APIs, and integration tests to preserve 100% functionality.
- Architected and deployed a full CI/CD pipeline using GitHub Actions, automating build, test, and deployment workflows to eliminate manual deployment errors across staging and production environments.
- Engineered an AI tutoring layer with GPT-powered adaptive flashcards and lessons, then integrated secure auth and Stripe subscriptions, driving a 25% engagement increase across 3 schools

Research Paper

June 2025 – December 2025

- Co-authored a Research paper on a self-critiquing LLM pipeline (expert feedback + distillation + RL, Pareto optimization) that improved accuracy while cutting infra costs up to 90% with open-weight models.

Universacare, Software Automation Intern

June 2023 – September 2023

- Built a Python Selenium web scraping workflow to extract daily Home Health Aide data, automating manual work and improving productivity by 11%.
- Developed 15 reusable React components and integrated a role-aware online application form into the production website, improving UI consistency/access control and saving 2+ hours per submission.

Projects

Squeeze AI (Morgan Stanley 2026 Hackathon winner)

FastAPI | LangChain | LangGraph | Nextjs | Supabase

- Built a multi-agentic event management platform with autonomous event creation and promotion agents, acquired by Lemontree (Hackathon partners) — now serving 25k+ users.

RhetoriQ ([Source Code](#))

Kafka | Flink | Kubernetes | Neo4j | pgvector | LangChain

- Architected a narrative intelligence agent detecting hot political trends across 100K+ daily documents using Apache Kafka/Flink, reducing latency to <500ms.
- Built and deployed a Kubernetes microservices backend (Terraform) with a polyglot data layer integrating Neo4j for graph traversal, pgvector for semantic search, and Redis for caching to trace political narrative chains.
- Developed an agentic LangChain investigation pipeline leveraging HuggingFace NLP models for semantic similarity detection, with Prometheus/Grafana observability for 15+ microservices.

Autonomous Strategy Agent ([Source Code](#))

Python | PyTorch | Claude API | FastAPI | MCP

- Built an MCP-based autonomous agent using Claude for strategy and PyTorch PPO for fast execution (<200ms p95), with typed tools and opponent behavior modeling for Clash Royale.

YOLOSynth

- Built “Autonomous Dataset Agent,” a hackathon-winning AI/ML systems pipeline that converts raw YouTube videos into deployable YOLOv8 object detection models through automated video ingestion, frame extraction, AI-generated bounding-box labeling, augmentation, training, evaluation, and retraining.
- Automated dataset lifecycle using vision LLMs and Codex subagents — 5x expansion, zero manual labeling.