

Nalin Nagar

github.com/nndev1 | [linkedin.com/in/nalin-nagar/](https://www.linkedin.com/in/nalin-nagar/) | nalinn2@illinois.edu

EDUCATION

B.S. in Computer Engineering and Minor in Math

May 2028

University of Illinois Urbana-Champaign

GPA: 3.89/4.0

Relevant Coursework: Intro to Computing, Computer Systems and Programming, Discrete Structures, Computational Linear Algebra, Intro to Electronics, Data Structures and Algorithms, Numerical Methods, Number Theory, Analog Signal Processing

TECHNICAL SKILLS

Languages: Python, C++, JavaScript, TypeScript, Java

Frameworks and Libraries: PyTorch, LangChain, Apache-Spark, Pandas, Scikit-Learn, PostgreSQL, Dask, Node.js, React, FastAPI, Git, AWS(Sagemaker, Lambda), Azure(CosmosDB, App Service), Docker, gRPC, Terraform, MongoDB

EXPERIENCE

Stripe

June 2026

Software Engineering Intern

San Francisco, CA

- Incoming Summer 2026

Camber (YC W21)

June 2025 – September 2025

Software Engineering Intern

New York, NY

- Built and productionized an internal **AWS + Notion MCP** powered by **Bedrock LLMs**, orchestrated with **InlineAgent** and **FastMCP** to interface with **Glue + Athena**, provisioned via **Terraform**.
- Cut DBT build times from **2+ hrs to 30 min** via SQL optimizations and indexing; improved backend performance by introducing **AWS ElastiCache** and local caching layers, cutting response times for high-traffic endpoints (e.g., Stripe transactions, clinic data) to **sub-100ms**.
- Created a demo environment with **MongoDB, RDS, DBT, and Python** for GTM, anonymizing PHI while preserving referential integrity using a **set + hashmap-based mapping algorithm** that maintained consistent joins across tables, scaled synthetic datasets for SMB and enterprise clinic demos, **generating \$500K+ revenue**.

Farragut

June 2024 – September 2024

Software Engineering Intern

Durham, NC

- Developed production-grade Automated Valuation Models (AVMs) for the N.C. Department of Revenue, enabling data-driven property assessments that **impacted \$100M+ in property valuations**.
- Leveraged Agents with LangChain and Azure AI Foundry to dynamically query **CosmosDB, REST APIs, and 200+** multi-source documents, automating appraisal query analysis and decision-making with **95% accuracy**.
- Built and deployed a **FastAPI and React-based** inference application on **Azure App Service using Docker**, streamlining the valuation process and cutting average appraisal time **by 90%** (from 2 days to under 5 hours).

Anomaly Six

June 2023 – September 2023

ML Engineering Intern

Alexandria, VA

- Processed **5B+ geolocation data points** using **Apache Spark and Dask clusters (500+ nodes)** from a **MongoDB** database to deliver real-time, large-scale data products for federal agencies including FBI and CIA.
- Trained and fine-tuned LLMs with PyTorch, LangChain, and HuggingFace offline/locally and optimized under strict memory constraints using **parameter-efficient fine-tuning (PEFT)**, and **quantization**.

PROJECTS

RapidMatch | C++, TCP, FastAPI

April 2025 – May 2025

- Designed and implemented a high-performance C++ order book matching engine, supporting market, limit (GTC), and fill-or-kill (FOK) orders with **price-time priority** and efficient **real-time trade matching**.
- Achieved a throughput of over **400,000 orders per second** with an average order processing latency of **2.5μs**.
- Built a custom **TCP client** and optimized communication using compact **bitfield-encoded messages** to minimize payload size and reduce network overhead in real-time order operations.
- Built a **FastAPI middleware** component allowing communication of a user with the orderbook through HTTP.