

Nihar Ogale

(248) 564-9353 | niharo@umich.edu | [linkedin.com/in/nihar-ogale](https://www.linkedin.com/in/nihar-ogale) | github.com/niharogale | niharogale.dev

EDUCATION

University of Michigan

Expected Graduation, Dec 2027

B.S.E in Computer Engineering

Ann Arbor, MI

- **GPA:** 3.95
- **Relevant Coursework:** Operating Systems, Distributed Systems, Machine Learning, Computer Networks, Parallel Programming with GPUs, Computer Organization

WORK EXPERIENCE

Amazon Web Services

May 2025 – Aug. 2025

Software Development Engineering Intern

Seattle, WA

- Designed and deployed low-latency, distributed backend services to power automated system code generation, enhancing data flow to reduce end-to-end pipeline execution time and manual engineering overhead by 60%
- Boosted request efficiency and response accuracy by 200% by integrating retrieval pipelines with data indexing
- Architected a resilient, 5-stage microservices CI/CD infrastructure utilizing AWS CloudFormation and automated rollback triggers to ensure fault-tolerant deployments while slashing rollout latency by 55%
- Engineered backend services to handle high-throughput requests, bolstering system reliability under load

University of Michigan - PURE-ECE

Dec. 2023 – May 2025

Artificial Intelligence Researcher

Ann Arbor, MI

- Engineered a device classification system for 1000+ IoT devices, streamlining inference pipelines and vectorized telemetry access patterns to strictly bound latency and improve throughput under large-scale evaluation workloads
- Built shared lookup infrastructure to accelerate evaluation workflows, reducing processing time by 80%
- Designed an asynchronous feature extraction layer to parse massive datasets of device-specific software logs and hardware heuristics, utilizing request batching to feed the AI fingerprinting models under strict latency boundaries

PROJECT EXPERIENCE

High-Performance GPU Convolution Kernel

C++ | CUDA

- Optimized a CNN forward-pass CUDA kernel, slashing time from 13s to 30ms by bypassing memory bottlenecks
- Achieved 90% SM occupancy and 83.8% FMA utilization via spatial coarsening and register-level optimization

User-Level Threading & Synchronization Library

C++ | Linux

- Engineered a preemptive user-level threading library in C++ from scratch, implementing custom CPU context switching and precise register state management to seamlessly multiplex execution across processes
- Architected rigorous synchronization primitives, including mutexes and condition variables, to guarantee strict thread safety and eliminate race conditions, minimizing context-switching overhead during execution

Sharded Paxos Key-Value Store

Go | Distributed Systems

- Engineered a highly available, sharded distributed key-value store utilizing the Paxos consensus protocol to orchestrate robust state replication and ensure consistent data routing across dynamic networked nodes
- Guaranteed strict linearizability and resilient system fault tolerance under heavy concurrent client loads, mitigating network partitions, message drops, and simulated node crashes without compromising data integrity

LEADERSHIP EXPERIENCE

Tau Beta Pi Engineering Honor Society

Aug. 2023 – Present

Professional Development Officer

Ann Arbor, MI

- Orchestrated a premier engineering career fair by driving corporate outreach to secure 300+ industry-leading technology partners, while structuring professional development workshops for over 200 engineering students
- Facilitated a hands-on high-performance computing curriculum centered on custom CUDA kernel development, taking members through advanced GPU parallelization, improving proficiency in hardware-level optimization

TECHNICAL SKILLS

Languages: C/C++, Python, Java, SQL, Typescript, Go, MATLAB, CUDA

Technologies: AWS, Docker, Git, Kubernetes

Systems: Memory Management, Networking (TCP/UDP), Parallel Computing (CUDA)