

# Krish Patel

 [krisapa](#)  [krisapa](#)  [krishspatel.com](#)  [kspatel8@icloud.com](mailto:kspatel8@icloud.com)  U.S. Citizen

## EDUCATION

---

### University of North Carolina at Chapel Hill

Dec. 2026

*B.S. Computer Science*

*Current GPA: 3.87*

**Courses:** Data Structures & Algorithms, Operating Systems, Compilers, Distributed Systems, Computer Networks, Cryptography, Database Systems, Digital Logic, Linear Algebra, Probability, Artificial Intelligence

## EXPERIENCE

---

### Susquehanna International Group

Jun. 2026 – Aug. 2026

*Incoming Trading Systems Engineering Intern*

*Philadelphia, PA*

### Tesla

Jan. 2026 – Apr. 2026

*Autopilot Engineering Intern — AI Inference Performance*

*Palo Alto, CA*

### UNC NLP

Aug. 2025 – Dec. 2025

*Machine Learning Research Assistant — MURGe Lab*

*Chapel Hill, NC*

- Researching evolutionary reasoning algorithms that use neural tree search to iteratively mutate action plans
- Integrating execution feedback, symbol grounding, and mutator ablations into ranking policies
- Evaluating iterative reasoning on long-context tasks to compare accuracy and efficiency with full-context models

### UNC Department of Computer Science

Aug. 2025 – Dec. 2025

*Undergraduate Teaching Assistant — Cryptography*

*Chapel Hill, NC*

- Creating labs on security games and reductions for PRG/PRF distinguishers, hybrid arguments, and IND-CPA

### Fidelity Investments

Jun. 2025 – Aug. 2025

*Software Engineering Intern — Machine Learning Infrastructure*

*Durham, NC*

- Launched GraphRAG-powered code intelligence platform that reduced onboarding time by ~60%
- Built AWS ETL pipelines to parse and embed ASTs, dependencies, and symbols from 120+ codebases into Neo4j
- Tuned graph indexing and caching for semantic queries achieving 45ms p95 latency
- Evaluation harnesses (golden datasets, regression, policy, red-team) enabled auto-gated LLM deployments

## TECHNICAL PROJECTS

---

### CAN-Cuda Logger | C++, CUDA, SocketCAN, LZ4

- GPU-accelerated CAN logger that batches and compresses frames in real-time with CUDA kernels
- Overlapped PCIe transfers and kernel execution to keep latency under 2ms while sustaining 90% GPU utilization

### Chrome Dino on FPGA | SystemVerilog, MIPS Assembly, Xilinx Vivado

- Built complete 32-bit MIPS computer on a Nexys A7 FPGA with VGA graphics and memory-mapped I/O
- Wrote the Chrome Dino game in MIPS assembly with smooth graphics, sound, and progressive speed scaling

### PeerBeam | Go, SvelteKit, TypeScript

- Cross-platform peer-to-peer file sharing app using WebRTC data channels for 90MB/s transfers without a server
- Tuned ICE negotiation and UDP parameters with Wireshark to reduce packet loss by 20% on congested networks

### STL-MLP | C++

- Built a STL-only feedforward neural network with matrix operations and backpropagation written from scratch
- Achieved 94% accuracy on UCI Seeds dataset using 5-fold cross-validation

## SKILLS

---

**Languages:** C++, C, Python, Go, CUDA, SystemVerilog, SQL, Java

**Systems & Networking:** Linux, TCP/IP, Wireshark, gRPC, perf/trace

**ML & Data:** PyTorch, ONNX Runtime, TensorRT, NumPy, Pandas

**Cloud & Infra:** AWS, Docker, PostgreSQL, Redis, Neo4j