# Krish Patel

krisapa | krisapa | krishspatel.com | kspatel8@icloud.com | U.S. Citizen

## Education

**University of North Carolina at Chapel Hill** Dec. 2026
*B.S. Computer Science, B.S. Statistics and Analytics* *Current GPA: 3.89*
**Courses:** Data Structures & Algorithms, Operating Systems, Computer Organization, Database Systems, Cryptography, Computer Networks, Distributed Systems, Algorithms & Analysis, Digital Logic & Computer Design

## Experience

**Fidelity Investments** | *Python, PyTorch, C++, Neo4j, AWS* June 2025 – Aug. 2025
*Software Engineering Intern — Machine Learning Infrastructure*

- Launched LLM agents that automated analysis and code generation, powering company-wide engineering tools
- Built AWS ETL pipelines to parse ASTs, dependencies, and symbols from 120+ codebases, enabling fast (35ms) semantic queries in Neo4j knowledge graphs
- Shipped graph-driven codegen tool that emits production-ready client code, reducing onboarding time by ∼60%
- Developed evaluation harnesses (golden datasets, RLHF, regression tests) that improved LLM safety, catching 25% more failures pre-release and enabling auto-gated LLM deployments with zero manual review

**UNC CS Experience Labs** | *Python, FastAPI, PostgreSQL, Redis* Aug. 2024 – Dec. 2024
*Backend Engineer — Campus Dev Team*

- Built backend services for CS office hours and coworking platform used daily by 100+ students and staff
- Refactored backend data layer and added caching, reducing slow queries by 40% lowering peak DB CPU usage
- Improved reliability and release speed by integrating OAuth SSO, automating deployments, and adding observability (Prometheus/OpenTelemetry)

## Technical Projects

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

- Developed a high-speed CAN bus logger using custom CUDA kernels to batch and compress frames on NVIDIA GPUs (LZ4), sustaining 100 MB/s with zero dropped frames and reducing log size by 40%
- Overlapped PCIe transfers and kernel launches with CUDA streams to maintain over 90% GPU utilization and cap per-batch latency under 2 ms

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

- Designed and built a complete 32-bit MIPS computer system on a Nexys A7 FPGA, integrating custom CPU, VGA graphics, keyboard and accelerometer input, and PWM audio via memory-mapped I/O
- Wrote the full Chrome Dino game in MIPS assembly and optimized hardware logic to meet 100 MHz timing closure for stable graphics and sound

**PeerBeam** | *Go, SvelteKit, TypeScript*

- Created a cross-platform, distributed, peer-to-peer file sharing app (web UI & CLI) enabling fast, private transfers up to 90MB/s, with no relay servers required (WebRTC data channels)
- Diagnosed and optimized UDP networking over Wi-Fi using Wireshark/tcpdump, tuning ICE and keep-alive logic to cut retransmissions by 20% on congested networks

**BGP-Lite Router** | *C, Python, Netlink*

- Built a lightweight BGP routing daemon in C to manage eBGP sessions and program routes into the Linux kernel
- Achieved full convergence of 3,000 prefixes in 1.2 seconds on a 2-vCPU VM, minimizing downtime and speeding up route failover during network events

## Skills

**Languages**: C++, C, Python, Go, CUDA, SystemVerilog, SQL, JavaScript/TypeScript, Swift, Java
**Systems & Networking**: Linux, TCP/IP, Wireshark, scapy, gRPC, perf/ftrace
**ML & Data**: PyTorch, Core ML, NumPy, Pandas
**Cloud & Infra**: AWS, Docker, PostgreSQL, Redis, Neo4j