Purnoor Sharma

469-803-0515 | purnoor.sharma@utexas.edu | linkedin.com/in/purnoor-sharma | github.com/noorps

Education

University of Texas at Austin — Austin, TX
B.S. in Computer Science and Informatics, Minor in Data Science and Business — *Expected May 2028*

Relevant Coursework: Data Structures, Algorithms & Complexity, Software Engineering, Object-Oriented Programming, Operating Systems, Computer Architecture

Experience

Microsoft — Redmond, WA
Software Engineering Intern – Security and Identity — *Incoming Jan. 2027*

Google — Austin, TX
Software Engineering Intern – Google Cloud — *Jul. 2026 – Aug. 2026*

- Cut 6 hrs/wk of manual release-note research for Technical Account Managers with a RAG system on Vertex AI.
- Shipped a pipeline on BigQuery, Cloud Run, and Cloud Functions tracking changes across 200+ GCP services.
- Raised retrieval accuracy 20% with entity extraction and source grounding, shipped to users in Streamlit.

Enterprise Mobility — St. Louis, MO
Technical PM Intern — *May 2026 – Jul. 2026*

- Cut vehicle-authorization penalty risk 22% across 9,500+ branches by automating a manual approval process.
- Turned requirements from operations staff into 500 Java and Drools rules running in the production engine.
- Caught misclassified approvals across 2,000+ rule engine outputs by diffing them in SQL against history.

AfterQuery (YC W25) — Remote
AI Data Engineering Intern — *Oct. 2025 – Jan. 2026*

- Tested generative AI tool calls across 10+ domains by writing multi-step evaluation cases for Manus AI.
- Automated scoring of 3K+ agent runs against researched ground truth with a concurrent Python pipeline.
- Improved production prompts serving 2M+ users by reproducing failure modes and proof-testing each fix.

UT Austin Department of Computer Science — Austin, TX
Undergraduate Research Assistant (Deep Learning & Data Engineering) — *Jun. 2025 – Oct. 2025*

- Built distributed data pipelines cleaning 10M+ meteorological records for ingestion into climate models.
- Trained Transformer models for time-series imputation, shipped into NASA and TxSON's agriculture tool.
- Benchmarked 10+ statistical algorithms to raise data consistency across predictive drought forecast models.

Projects

cuconv | CUDA Kernel Optimization | *CUDA, C++, Nsight Compute, cuDNN, PyTorch*

- Wrote four CUDA convolution kernels and beat cuDNN by 1.2x to 1.6x on five of six batch-1 layer shapes.
- Raised FMA pipe utilization from 11% to 90% by picking each optimization from Nsight Compute counters.
- Matched cuDNN bit for bit on resnet18's first layer across 802,816 outputs, verified on 15 shapes.

tern | Telemetry Ingestion Service | *Java, Spring Boot, Kafka, Jenkins, Docker, Prometheus*

- Ingested 10,000 readings at 1,003 events/sec through Spring Boot services at 56.3 ms p50, 106.6 ms p95.
- Held per-device ordering under 64 concurrent clients, rejecting duplicate ids and stale late readings.
- Ran the stack on Docker Compose with Prometheus counters and a Jenkins gate on benchmark regressions.

Reposter | Shipped Full-Stack Product | *JavaScript, Supabase, PostgreSQL, Stripe*

- Built and launched a Chrome extension on my own, now serving 1,000+ paying small-business customers.
- Took apart Facebook's Lexical editor and drove it through Chrome's Debugger API when DOM methods failed.

Honors & Recognition

Jane Street INSIGHT | *Software Engineering Track* — Aug. 2026
Amazon University Reachback | *Invited Participant* — Summer 2025 & Fall 2026

Technical Skills

Languages: Java, Python, C++, Go, TypeScript/JavaScript, SQL, Rust, C, Scala, Bash/Shell
Backend & APIs: Spring Boot, Kafka, REST APIs, OpenAPI, Node.js, Express, OAuth 2.0, Webhooks
Infrastructure: GCP, AWS, Docker, Kubernetes, Terraform, CI/CD, Prometheus, Linux, Git/GitHub
Data & AI: PostgreSQL, BigQuery, Trino, Spark, Vertex AI, RAG, PyTorch, CUDA, Nsight Compute