

SRINATH SURESH KUMAR

404-630-5415 skumar780@gatech.edu linkedin.com/in/srinathsureshkumar github.com/sureshkumarsrinath

Education

Georgia Institute of Technology

Master of Science in Computer Science GPA 4.0/4.0

Aug. 2024 – May 2026

Atlanta, Georgia

Relevant Coursework: Big Data Systems & Analytics, Adv Database Systems, Distributed Computing, High Perf. Comp.
College of Engineering Guindy

Aug. 2018 – May 2022

Bachelor of Engineering in Computer Science & Engineering (Hons) CGPA 9.29/10

Chennai, India

Experience

Tesla

Software Engineering Intern

May 2025 - Present

Fremont, CA

- Optimized data pipelines for Energy Devices data, moving from an Entity-Attribute-Value (EAV) tall tables to Apache Iceberg wide tables, reducing storage by 85% and pipeline run times by up to 20 seconds
- Refactored database connection layer to enable seamless migration of data pipelines to new compute clusters with zero code changes, by adopting to Environment-as-a-Parameter patterns, reducing onboarding effort by 30%
- Developed and deployed data pipelines to capture the diagnostic accuracy and service proficiency in Energy Devices that powered critical dashboards

Arcesium India Pvt Ltd

July 2022 – July 2024

Senior Software Engineer - Aquata Data Platform

Bengaluru, India

- Reduced the metadata fetch time to less than 100 milliseconds from 7-10 seconds by building distributed caches
- Developed a low/no-code platform to build ETL pipelines reducing the end-to-end build and deployment time by 90%
- Orchestrated an event-driven data pipeline triggering system, with 30% increased resource efficiency of Spark jobs
- Implemented distributed rate-limiting for the micro services in the platform using Bucket4j with less than 10ns latency
- Built microservices to query/edit data in the data lake, reducing the edit time from 2-3 minutes to under 30 seconds
- Benchmarked Spark jobs & implemented dynamic allocation of resources for ETL jobs, reducing the cloud costs by 35%

CEG Tech Forum

November 2021 – June 2022

Backend Developer / DevOps

Chennai, India

- Deployed Node APIs for registration and payments with subsecond latencies & payment success rate greater than 99.8%
- Scaled the system to handle a traffic of close to 1 million incoming requests a day with 100% uptime with auto-scaling

Technical Skills

Languages: C/C++, Java, Kotlin, Python, Scala, JavaScript, TypeScript, SQL, HTML, CSS, Go, Bash

Technologies/Frameworks: Spring Boot, NodeJS, Flask, ReactJS, Hazelcast, Express, Redis

Devops: Nginx, Git, Docker, Kubernetes, CI/CD Pipelines, Amazon Web Services, Terraform, Argo, Airflow

Data Engineering: Apache Spark, PySpark, Hadoop, Presto(Trino), Snowflake, PostgreSQL, Iceberg, Delta Lakes

Data Science: Scikit-learn, TensorFlow, Keras, Pytorch, HuggingFace

Projects

Fault Tolerant Distributed Sharded Key-Value Store

January 2025 - May 2025

- Developed a highly-available, sharded key-value store using 3-way Paxos replication for fault tolerance, successfully passing concurrency tests designed to simulate thousands of concurrent operations per second across multiple nodes
- Verified the system's resilience to more than 5,000 randomized failure scenarios, with strong consistency guarantees and no data loss

Concurrent Operations on B-plus Trees

August 2024 - December 2024

- Developed lock-based and latch-free B-plus trees, to support concurrent inserts, deletes, and look-ups, ensuring correctness under multi-threaded and multi-processor execution
- Increased throughput by up to 5 times by reducing lock contention using fine-grained node-level locking and using lock-free compare-and-swap (CAS) operations with near-linear scalability to 16 cores

Ensemble of LLMs for Sentiment Analysis on Neologisms

September 2024 - December 2024

- Implemented an ensemble for Large Language Models (LLMs) to detect sentiment expressed by Neologisms on Social Media and achieved an accuracy of 87% on Twitter Neologisms using Hugging Face Transformers and PyTorch
- Studied the performance of task based fine tuning techniques and resource efficiency PEFT techniques like LoRA, DoRA and AdaLoRA against full model fine tuning of the GPT-2 and BERT based LLMs

SDN-based Intrusion Detection System For Mitigating Automated Malicious Traffic Flows

May 2022

- Build a traffic classification engine based on Random Forest algorithm to detect anomalous flows with F1 score of 95.3%
- Designed SDN controller rules to identify, verify real human users and serve the requested resources only to legitimate requests increasing file download speeds by 30%