Rohan Varshney

rohanvar@gmail.com | **rohanvarshney.github.io** | linkedin.com/in/rohanvarshney | github.com/rohanvarshney

Education

Georgia Institute of Technology, Atlanta, GA ***Graduated May 2021***

- B.S. in Computer Science (Modeling / Simulation & Intelligence Concentrations)

Experience

Lyft, SWE, New York City, NY

Data Fabric & Data IDL (Analytics Events, Protobuf, Kinesis, ETL), ***Software Engineer I*** ***July 2021 - August 2022***

- Drove the migration from CJSON to Google's Protobuf format for 450+ services across Data Platform's 2400+ unique analytics events by pioneering new on-the-fly (un)marshal library functionality and refactoring fleet-wide, saving ~15% CPU
- Utilized a native schema-retrieval API service to adjust the methodology used in analytics events proxy sidecars of all Lyft services that emit events to reduce their memory footprint up to ~80%, reducing EC2 provisioning costs

Data Persistence (GraphQL, Hive & Spark, Golang, Apache Flink), ***Software Engineer II*** ***August 2022- May 2023***

- Served as tech lead (and sole oncall) of T0 services, GSchema & DDBPersister. Saved 100s of engineering hours over 20+ table onboardings with engineered automated processes. Reduced the SLA breach frequency, monthly PagerDuty page count, and failed record redrive quantity all by 75+%.
- Facilitated work of 6 Data Platform contractors to track, confirm, and monitor significant PRs to migrate from & deprecate Hive query engine for Lyft's biggest DAGs, personally unblocking complicated use cases w.r.t new Spark syntax and rollouts.

Streaming Compute (Feature Service, TecTon, Apache Flink, Golang), ***Senior Software Engineer*** ***May 2023 - Present***

- Implemented novel architecture supporting the storage & low level latency of Embeddings via OpenSearch ANN/KNN vector search capabilities, unlocking ML <> LLM use cases at Lyft (i.e. M1 HAP Search Bot for customer support, RAG platform).
- Led KTLO initiatives that enabled cost-saving org-level optimizations: deprecation of Flyte ($150K+/yr), Redis cache and EKS cluster downsizing ($50K+/yr), preliminary implementation of TTL + backing stores consolidation ($200K+/yr).
- Evaluated managed solution for Feature Service, authoring 200+ pages of documents detailing current system capabilities, use cases, customers, hypothetical migration plan design, & evaluation criteria. Led meetings with Director, Staff, & TPM.

Microsoft, Software Engineering Intern, Redmond, WA (Virtual) **May 2020 – August 2020**

- Extended an internal chatbot that resides in Microsoft Teams with 8 new deployment-related features via an architecture that enabled greater bot utilization in the XDeployment team, cutting down Time to Investigate & Response for on-call to answer FAQs for clients regarding Azure Storage deployments & tenant-based information by up to 4-6 hours
- Achieved end-to-end integration for all coded features using the backend command client of KQL functions in the XStore connection of Kusto, appropriate response objects coded in C#, & the frontend dialogs that handle user input & parameter validation through extensive communication & assistance from a team in Shanghai (EST+12:00)

Lyft, Software Engineering Intern, San Francisco, CA **January 2020 – May 2020**

- Invented 10 prototype dashboards using Mode Analytics by amalgamating Presto SQL queries, Python Notebooks, & HTML/CSS in order to generate insight reports about the performance of various Basemap quality metrics, thereby increasing confidence in the adoption of newly updated maps as part of the Scorecard project
- Engineered novel heatmaps using (geo)pandas, road network segments as well-known text (wkt) geometries, & Geohash weights through DataFrames to provide spatial visualization of the quality metrics on a region-level, allowing for immediate validation of metric heuristic iterations & pertinent map changes made in OpenStreetMap (OSM)

Amazon.com, Software Development Engineering Intern, Seattle, WA **May 2019 – August 2019**

- Created multiple new features (autocomplete, bookmark migration for favorites) for the Amazon Enterprise Access (AEA) Android app for AWS that improved usability & robustness for the 5500+ (and exponentially growing) unique monthly Amazon employees (as of August 2019) in accessing corporate services from their mobile devices
- Integrated dynamic service catalog retrieval to substantially reduce the onboarding duration for service access from 1 month to near real-time (less than 1 hour) by levying multiple AWS services (API Gateway, Lambda, DynamoDB, SNS, S3) to remove manual dependencies across multiple project packages & automate the process

Manhattan Associates, Software Engineering R&D Co-Op, Atlanta, GA **January 2019 – May 2019**

Raytheon, Software Engineering Intern, Indianapolis, IN **May 2018 – August 2018**

Skills

Languages:	Java & Kotlin, Python, Golang, C# & C, SQL, Typescript, HTML & CSS &JS, MATLAB, Bash, Terraform, YAML
Technologies:	Jupyter & Python & Databricks Notebooks, Angular, Amazon Web Services (S3, Kinesis, MSK, DynamoDB, Redis), Hive & Spark, Docker, Apache Flink, NumPy & Pandas, Postman, Node.js, LaTeX, Slack