

ROHIT MITTAPALLI

rohit.mittapalli@gmail.com | github.com/rohitmittapalli | rohitmittapalli.com | 630-777-4728 | US Citizen

PROFESSIONAL EXPERIENCE

Applied Intuition

May 2021 - Present

Software Engineer | Data Platform

Mountain View, CA

- **Apache Spark:** Reduced query latency on simulation, drive, and map data by over 50x. Introduced and fully deployed Spark in a Kubernetes environment across multiple production clusters across all major cloud providers.
- **Apache Hudi:** Improved query performance on simulation data by over 100x by leveraging data compaction, index creation, and clustering in Apache Hudi. Deployed Hudi in Kubernetes at a production scale. Maintained the Applied fork of the Hudi repository with numerous patches to add additional schema flexibility for specific Applied use cases.
- **Customer growth:** Managed an end-to-end customer relationship, overseeing the entire process from contract sign to production usage with 20 daily users. Integrated the Applied product suite into the customer's release and production workflows. Facilitated regular office hours to provide support and gather user feedback to inform product planning.
- **Database maintenance:** Achieved a 70% reduction in necessary storage by leading two company-wide migrations to move data out of a monolithic POSTGRES database. Reduced database costs, increased our retention capabilities and upgrade POSTGRES major versions 11 to 14. Reviewed large database updates as a database admin.
- **Generic async job execution:** Created the initial framework for all asynchronous jobs supporting 20+ production tasks. Built the framework using Kubernetes and RabbitMQ with out-of-the-box auto-scaling in KEDA, monitoring with Prometheus/Grafana, and custom alerting. Framework supports both cloud and local ships of the Applied product suite.
- **Technical mentorship:** Served as the company expert for Spark/Hudi and led several "How Does our Code Work" sessions. Ran the database and GraphQL onboarding session required for all engineering new hires at a monthly cadence.
- **On-call rotation:** Set up the very first PagerDuty rotation. Laid out best practices for handling on-call issues. Created workflows for customers and application engineers to submit on-call incidents. Conducted on-call training for all new employees joining the team's rotation. Served as primary escalation for the team for over a year.

Uber

Sept 2019 - Dec 2019

Software Engineering Intern | Search Infrastructure

Palo Alto, CA

- Optimized performance of weighted queries by 150x by building a custom WAND (Weighted AND) Lucene operator
- Improved performance on spatial queries by creating a Lucene operator to iteratively expand geographic search area

Microsoft

May 2019 - Aug 2019

Software Engineering Intern | PowerAI

Seattle, WA

- Built an expandable ML pipeline to visualize the performance of feature engineering and selection methods

RESEARCH EXPERIENCE

Adaptive Learning Lab

Jan 2020 - May 2021

Student Researcher | Dr. Judy Hoffman

Atlanta, GA

- **Transfer learning:** Led code development for the published paper (NeurIPS 2021) titled "Scalable Diverse Model Selection for Accessible Transfer Learning". Introduced a model bank, benchmarks, and a new method (PARC) to improve model transferability estimation in transfer learning intended to support low-resource training.
- **Adversarial learning:** Published paper "Likelihood Landscapes: A Unifying Principle Behind Many Adversarial Defenses" on adversarial defenses in machine learning at AROW@ECCV. Developed the scalable test bench to run experiments and visualizations at scale on Georgia Tech compute.

SKILLS

Languages: Python, C++, JAVA, SQL, Javascript, Typescript

Technologies: Apache Spark, Apache Hudi, Kubernetes, Terraform, POSTGRES, React.js, Azure, AWS, PyTorch, Lucene

EDUCATION

Georgia Institute of Technology

May 2021

B.S. in Computer Science (Intelligence & Theory)

GPA: 4.00 / 4.00