

Animesh Kumar

+1(858)-241-9504 | ank028@ucsd.edu | [linkedin.com/in/animeshk08](https://www.linkedin.com/in/animeshk08) | github.com/animeshk08

EDUCATION

University of California San Diego

Master of Science in Computer Science. **GPA: 3.92/4.00**

San Diego, CA

Sept 2023 - Dec 2024 (Expected)

National Institute of Technology Karnataka, Surathkal

Bachelor of Technology in Computer Science and Engineering. **CGPA: 9.05/10**

Surathkal, India

July 2017 - May 2021

TECHNICAL SKILLS

Programming Languages: C++, C, Python, Java, GO (Golang), JavaScript, Ruby, SQL, Shell, HTML, CSS

Skills: Docker, Kubernetes, Azure, AWS, GCP, CUDA, PyTorch, TensorFlow, React, Spring, Django, Kafka, Spark, OracleDB, MySQL, Snowflake, Datadog, MongoDB, Elasticsearch, REST, JUnit, Terraform, OOP, Bash, CI/CD, Agile

EXPERIENCE

Shang Data Lab, UCSD

GenAI Student Researcher

Sept 2023 - Present

San Diego, CA

- Developed a claim-dependency graph network to predict the novelty in patent applications, **reducing the timeline by days** and outperforming fine-tuned LLMs, like Llama 2, Mistral, and GPT-4 with an AUC of 0.67. [ACL '24](#)
- Implemented **cross-lingual data contamination detection** on multiple-choice benchmark datasets using LLM-as-a-judge, overcoming limitations in current state-of-the-art contamination detection measures. [EMNLP '24](#)
- Developing **model-merging** methods, resolving redundant and conflicting parameters, to merge domain specific LLM experts and MoEs into a single multitask model without additional training overhead.

AppFolio, Inc.

Software Engineering Intern

June 2024 - Sept 2024

San Diego, CA

- Engineered a rental property recommendation system leveraging **CLIP embeddings**, **V-LMs** and a Llama 3 vision encoder fine-tuned using LoRA adapter, optimizing property ranking, improving user retention by 10%.
- Designed a **React** web app, **SQL** models, **GraphQL** APIs and a **Kafka** event publisher for managing user lifecycle and authorizations, migrating 20k+ business partners from legacy IAM platform to unified user profile system.
- Developed a Python-based unified authentication platform, using Keycloak Identity and Access Management microservice, integrated with SSO, 2FA and auth tokens, supporting 50+ services and over **150k+** daily requests.
- Optimized **Hadoop MapReduce jobs** to extract features from ~ 20TB of user rental search logs, reducing job execution time by 30%, enabling consolidated data feature to be utilized in rental property recommendations.

Oracle

Software Development Engineer (MTS)

July 2021 - Sept 2023

Bengaluru, India

- Integrated a machine learning pipeline in the infrastructure operations workflow using an XLNet transformer for incident classification and root cause prediction, resulting in **80+ hours per week** of reduced engineering time.
- Developed a **multi-threaded Java micro-service**, exposed as a Spring and React web app, automating database lifecycle management for 10k+ infrastructures on OCI Database Cloud, resulting in \$100k+ quarterly savings.
- Architected a **query language engine** to streamline escalated incidents into Kafka events, leveraging **Spark** for data processing and routing a high volume of incidents across teams, reducing operation time by over 30%.

Kubernetes, The Linux Foundation

Open Source Software Engineering Intern (Mentee)

May 2020 - July 2020

Remote

- Implemented gRPC procedures to core [Kubernetes codebase](#), collaborating with Azure Storage team to integrate Docker Container Storage Interface(CSI) drivers for SMB, supporting over **1500 Azure Storage** customers.

SELECTED PROJECTS

Multimodal Temporal Lobe Epilepsy Analysis using GAN and Diffusion Models | [Github](#)

- Developed a StyleGAN-based model for latent space manipulation to analyze and classify the progression of Temporal Lobe Epilepsy (TLE) in MRI scans, helping doctors identify key areas and features for diagnosis.
- Integrated a AltDiffusion model to develop attribution maps highlighting brain anatomy features of TLE patients.

Loma Open MPI: Distributed Automatic Differentiation | [Github](#)

- Developed a source-to-source distributed automatic differentiation compiler converting native Python code base to differentiable CUDA and C code running on multiple nodes and communicating using Open MPI framework.

Cardiac Electrophysiology Simulation using CUDA and NVComp

- Developed a CUDA interface to simulate kinetics of the membrane of a heart cell, on high throughput compressed dataloads using Nvidia NVComp, achieving over 20% runtime improvement over standard OpenMPI platform.