

Ganesh Palanikumar

550 E Weddell Drive, Apt 1208, Sunnyvale, CA 94089

Email : ganeshpkumar93@gmail.com

Mobile : +1 (408) 581-1410

EXPERIENCE

- **Omneky Inc.** San Francisco, CA
Data Engineering Lead *Jan 2023 - Present*
 - **Data Module Redesign and Implementation:** Redesigned and optimized the data importer module that is used to populate internal database. Designed and implemented Data Access API to regulate access to database from backend and AI servers.
 - **ELT Data Ingestion Pipeline using Airbyte and Python:** Lead a team of 3 engineers to transition from home-grown ETL pipeline to Airbyte-backed ELT pipeline. Was in charge of design, setting milestones and implementation of the new pipeline. Developed an observability suite to monitor the pipeline and alert in case of issues.
 - **Low-latency Interactive Data Platform:** Lead a team of 3 engineers to design and develop the backend for a customer-facing analytics dashboard. Involved design and implementation of the full system, database schema as well as setting up aggregates. Full setup was developed and deployed on AWS.
- **Apple Inc.** Cupertino, CA
Software Engineer (On-Device Frameworks/Big Data Application) *Feb 2019 - Oct 2022*
 - **CoreML Privacy:** Designed and developed on-device framework for private federated learning solution demonstrated in NeurIPS 2019. Involved generation of test data, training of model and iOS app creation.
 - **Turi Trial - Daemon:** Involved in ideation, design and development of on-device framework to enable adopters to integrate with our API. Also contributed to Trial daemon to enable multiple features, including download and unpacking of compressed assets required for running A/B tests in memory-sensitive nature.
 - **Turi Trial - Analytics:** Create, deploy and maintain Spark jobs that process analytics data from device. End-to-end design, development and testing of dashboard to show inferred metrics to customers. Work with customers to design and implement statistical tools to get more insights into experiment data.
- **Samsung R&D Institute** Bengaluru, India
Senior Software Engineer *Aug 2015 - May 2017*
 - **Samsung Pay Spain:** Focussed on developing back-end servers to support Samsung Pay Spain. Used REST APIs in Spring to create Admin Portal.
 - **Samsung Pay India:** Spearheaded ideation, design and development of the Value Added Services module for Samsung Pay India. Also supported performance testing of servers using Dynatrace and HP Loadrunner

KEY PROJECTS

- **Statistical Allocation Health Metrics for Experiment Validation:** Worked with adopters to understand pain points in using the product and designed a solution to address the concerns. Drove the project from design and ideation all the way till incorporating the feature in our product. Involved coordinating efforts with multiple teams and stakeholders at Apple.
- **Analytics Dashboard:** Owned the analytics dashboard product including data processing using Spark and development/testing of user-facing dashboard. Involved consistently meeting different stakeholders, understanding business needs and creating features to address the needs.
- **Benchmarking ACAI/Fairhair.ai:** Undertaken as Capstone project under guidance of Prof. Nyberg at CMU. Project aimed at developing a set of realistic benchmark challenges for natural language processing (NLP) tasks, with an initial focus on Named Entity Recognition (NER). Involves exploration of various data, features and models that are currently available in the NER domain.
- **Sharing-aware Last Level Cache Replacement Policy** Done as B.Tech project under the guidance of Prof. Madhu Mutyam. Developed a dynamic, sharing-aware last level cache replacement policy targeted at multi-threaded workloads. Algorithm proposed performed at least as good as Least Recently Used (LRU) policy.

EDUCATION

- **Carnegie Mellon University** Pittsburgh, PA
Master of Computational Data Science *Sept 2017 - Dec. 2018*
- **Indian Institute of Technology, Madras** Chennai, India
Bachelor of Technology in Computer Science and Engineering *Aug 2011 - May 2015*

PUBLICATIONS

- [Best Paper, SIGMETRICS 2020] Rateless Codes for Near-Perfect Load Balancing in Distributed Matrix-Vector Multiplication ([link](#))