

PARTH VIPUL SHAH

San Francisco, California | (661) 514-7214 | parthvshah@gmail.com | <https://parthvshah.me>

EDUCATION

University of Southern California (USC)

Master of Science in Computer Science; GPA 3.67

Los Angeles, CA

Jan 2022–Dec 2023

PES University

Bachelor of Tech. in Computer Science and Engineering; GPA 4.00 (9.06/10.00)

Bangalore, India

Aug 2017–Jul 2021

SKILLS

- Technologies: Perf, Dtrace, Linux, PostgreSQL, Docker, AWS, Azure, GCP, Airflow, React, MongoDB, .NET, Git, TensorFlow, Neo4j
- Languages: Bash, C/C++, Python, JavaScript, C#, HTML/CSS, SQL, Solidity

EXPERIENCE

Oracle

Redwood Shores, CA

Linux Performance Engineer

Feb 2024–Present

- Conducting root cause analysis into performance regressions and improvements observed in Oracle's Unbreakable Enterprise Kernel (UEK). Using Linux tools such as perf and dtrace.
- Employing statistical methods to analyze performance data and make recommendations to improve UEK and other Oracle products. Using Python libraries such as numpy, scipy, pandas and matplotlib for analysis and plotting.
- Contributing to new hardware platform development by characterizing the performance of benchmarks on a platform.
- Developing kernel performance measurement infrastructure used routinely to measure the performance of UEK.
- Creating tools and methodologies to characterize, analyze and profile workloads used to characterize the performance of UEK.

Commvault

Bangalore, India

Associate Software Engineer

Jan 2021–Dec 2021

- Spearheaded development of the PostgreSQL data agent - multi-stream file system/dump-based backups and restores.
- Enabled protection of on-prem and cloud (AWS, Azure, GCP) PostgreSQL databases for **25+** environments by working on **90+** enhancements and defects.

Samsung R&D Institute

Bangalore, India

Research Intern

May 2020–Jul 2020

- Created a deep learning model with TensorFlow Lite to detect device overheating as part of the On Device AI team.
- Processed raw handset images to determine ambient temperature - average RMSE of **3.79** (CNN) and **2.32** (LSTM) for single and multi-image prediction.

RESEARCH

Information Sciences Institute, Los Angeles, CA

Nov 2022–Dec 2023

- The unequal opportunities of large language models: Examining demographic biases in job recommendations by ChatGPT and LLaMA. Studying the biases in Large Language Models (LLMs). **20 citations**. Published at ACM's conference titled Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO'23). dl.acm.org/doi/abs/10.1145/3617694.3623257
- Creating a Bayesian model to represent Ranked Choice Voting responses. Can predict the outcome of elections in the United States and segment populations based on survey responses.

PES University, Bangalore, India

Jun 2020–Sep 2020

- Prediction of the Peak, Effect of Intervention and Total Infected by the Coronavirus Disease in India: Forecasted using the SEIR compartmental model. **4 citations**. Published with the Cambridge University Press. doi.org/10.1017/dmp.2020.321.

SSCU, Indian Institute of Science (IISc), Bangalore, India

Jun 2019–May 2020

- Conceived 3 parallel algorithms for computing correlation functions using the MPI standard in C - achieved **super linear speedups**. Released package on a 120 node HPC cluster, academic paper is under review.

PROJECTS

- Video Indexer: Indexes a video into scenes, shots and sub-shots based on visual and audio features. Achieved **99%** accuracy.
- Nutritional KG: A recommendation system that uses a knowledge graph with **10K+** entities to help minutely alter your diet.

- Converting Black-box Neural Networks into Interpretable Decision Trees, Explainable AI: Processed using layer-wise relevance propagation and perturbations. Model and data agnostic methods.
- Naturalization of Text by the Insertion of Pauses and Filler Words: Used bigram frequency and an RNN. **55%** convincing.
- Database as a Service: Created using Docker, RabbitMQ, ZooKeeper. High availability and scalability. Tested on AWS. **100%** uptime and supported **500+** concurrent reads/writes.