

Amogh Patankar

408-597-2478 | apatankar@ucsd.edu | [linkedin.com/in/apatankar22](https://www.linkedin.com/in/apatankar22) | github.com/apatankar22 | apatankar22.github.io

EDUCATION

- University of California, San Diego** September 2023 – June 2025
M.S., Computer Science & Engineering; Concentration: AI & Machine Learning La Jolla, CA
- University of California, San Diego** September 2020 – March 2023
B.S., Data Science La Jolla, CA

PROFESSIONAL EXPERIENCE

- AI Applications Development Intern** September 2024 – December 2024
Advanced Micro Devices (AMD) San Jose, CA

- Developing high-level software for **RyzenAI Neural Processing Unit (NPU)**; specifically, optimizing performance and benchmarking with respect to Apple, Nvidia, and Qualcomm products.
- Optimizing AMD hardware to efficiently execute generative AI workloads, and enabling hardware to achieve optimal performance and power efficiency. Workloads include transformers, large language models, convolutional neural networks (CNN), and diffusion models.
- Optimization and benchmarking may be applied to other CPU/GPU products- **AMD Radeon, Instinct, EPYC**.

- Data Scientist (Generative AI) Intern** June 2024 – September 2024
Marvell Technology Inc. Santa Clara, CA

- Used generative pretrained transformer (GPT) models to generate synthetic data, leveraging parameter efficient fine-tuning (PEFT) techniques such as low-rank adaptation (LoRA).
- Implemented reinforcement learning (RL) algorithms like q-learning, and deep q-networks (DQN), as well as deep learning methods for DSP parameter optimization.
- Integrated large language models (LLM) and retrieval augmented generation (RAG) to automate hardware modeling process.
- Developer tools include AWS S3, EC2, and Sagemaker, as well as Tableau, Streamlit and SnowflakeDB.

- Researcher** June 2023 – August 2024
Stanford University School of Medicine Palo Alto, CA

- Developed statistical packages for multiple biomedical research teams in Python and R for data analysis. Packages composed of chi-squared and Fisher tests amongst other statistical utilities.
- Led research teams mentored by Dr. Gross, Dr. Palaniappan, and Jin Long to analyze opioid mortality and glycemic control in normal weight, type-2 diabetics.
- Opioid research published in British Journal of Anaesthesia, and diabetes research in preprint at Journal of Asian Health.

- Software Development Engineer Intern** June 2022 – September 2022
Amazon Web Services (AWS) Seattle, WA

- Improved latency of AWS Lex ASR (Automatic Speech Recognition) Services and AWS DataHub for conversational AI models by recommending and implementing architectural changes.
- Enabled compliant storage of critical and non-critical customer data in DataHub by designing and enhancing Lex ASR and DataHub schemas using AWS S3, Kinesis, and Lambda.
- Optimized ASR Service by up to ~75%, allowing faster resolution of customer requests; tracked using AWS CloudWatch.

- Research Intern** June 2021 – August 2021
Scripps Research Translational Institute La Jolla, CA

- Developed an R library to estimate genetic regulatory variation using a confidence interval estimation method, applying simulation-based techniques like parametric bootstrapping to assess the reliability of the estimates.
- Implemented statistical concepts such as binomial distributions, and conducted sensitivity analysis to evaluate the impact of different assumptions on the genetic data from the Genotype Tissue Expression Project (GTEx).

SKILLS

- **Languages:** Python, Java, R, C++, SQL, JavaScript
- **Frameworks and Tools:** PyTorch, Tensorflow, Keras, numpy, sklearn, pandas, AWS S3, EC2, SageMaker, Kinesis, Lambda, SnowflakeDB, Tableau, Streamlit, git

PROJECTS

- **Capstone Project:** Active Learning with Neural Processes for Epidemiology Modeling
- Prediction of Causes of Patient Readmission using Large Language Model(s) (**OpenAI, Llama-7B, BART-Large**)
- Data Science Interview Tool (**GPT-3, PyTorch, Python**)
- Autonomous Vehicle Trajectory using Deep Learning (**Argoverse 2 Dataset, PyTorch**)