Sharan Ramjee

sharanramjee.github.io | linkedin.com/in/sharanramjee
765-772-6865 | sharanramjee@gmail.com

EDUCATION

Stanford University
Stanford, CA
M.S. Computer Science
(Concentration: AI)
GPA: 4.04/4.00
2020 – 2022

Purdue University
West Lafayette, IN
B.S. Computer Engineering
2016 – 2020
GPA: 4.00/4.00
Highest Distinction

SKILLS

Programming Languages:
- Python
- C / C++
- JavaScript

Technologies/Frameworks:
- Transformers / LLMs
- PyTorch / TorchVision
- TensorFlow / Keras
- XGBoost / Scikit-Learn
- PyG / NetworkX
- SHAP / LIME / XAI
- Mongo / Postgres
- SQL / Trino / Presto
- PySpark / Hadoop
- Airflow / Flyte / WandB
- Kafka / Flink / S32K
- Docker / Kubernetes
- AWS / Azure / GCP

TEACHING ASSISTANT

- CS 221: AI (Spring '22)
- CS 142: Web Dev. (Winter '22)
- CS 230: [Head TA] DL (Fall '21)
- CS 107: Computer Organization & Systems (Spring '21)

RELEVANT COURSES

- CS 329T: Trustworthy ML
- CS 329S: ML System Design
- CS 231N: Computer Vision
- CS 231A: 3-D Computer Vision
- CS 224W: Graph ML
- CS 224N: NLP
- CS 230: Deep Learning
- CS 229: Machine Learning
- CS 221: Artificial Intelligence

ORGANIZATIONS

- Stanford TreeHacks

LINKS

- Google Scholar
- GitHub

INDUSTRY EXPERIENCE

Stripe
Jul. 2022 – Present
Machine Learning Engineer
San Francisco, CA
- Building ML models to balance Stripe's losses and UX as an MLE on the Fraud Discovery team
- Led a team of engineers, strategists, and ops specialists to propose and build Scorpion - Stripe's first time-series transformer for Risk Detection (Fraud, Credit, ToS, etc.) - saving ~$8M / year in losses
- Tech stack: SQL, Presto, PySpark, Airflow, Kafka, Flink, Flyte, Databricks, PyTorch, LLMs

Stripe
Jun. 2021 – Sep. 2021
Software Engineering Intern
San Francisco, CA
- Built ML models for merchant fraud detection as a SWE on the Fraud Intelligence team
- Proposed and built FI-Explain - an ML explainability tool to probe Stripe's merchant fraud models to gain insights on the decisions they make to trade-off losses and UX during fraud incidents
- Tech stack: PyTorch, XGBoost, transformers, SHAP, LIME, Git, Bazel

Google
Sep. 2019 – Dec. 2019
Software Engineering Intern
Seattle, WA
- Built ML explainability for models on Google Cloud as a SWE on the Google Cloud AI team
- Added Model Distillation capabilities to convert black-box ML models deployed on Google Cloud into interpretable tree-based models (Soft Decision Trees, Random Forests, Gradient Boosted Decision Trees)
- Tech stack: TensorFlow, Keras, scikit-learn, Fig, Blaze

Qualcomm
May. 2019 – Aug. 2019
Machine Learning Intern
San Diego, CA
- Built ML models for power-efficient Qualcomm chips as an MLE on the ML application analysis team
- Proposed and built a time-series LSTM model to estimate QoS parameters that trade-off performance and power depending on the Snapdragon chip's use-cases (AR/VR, Gaming, Multimedia, etc.)
- Tech stack: PyTorch, scikit-learn, Git

RESEARCH EXPERIENCE

Stanford University
Sep. 2020 – Jan. 2021
Graduate Researcher
Stanford, CA
- Researched robot learning for intuitive human-robot interaction using Computer Vision at the Stanford Vision and Learning Lab where I worked on parallelized reinforcement learning on iGibson

Purdue University
May. 2018 – May. 2020
Research Assistant
West Lafayette, IN
- Researched Deep Learning for Signal Processing applications with Prof. Aly El Gamal where I worked on feature selection and subsampling for Wireless Modulation Classification and Interference Identification

JOURNAL AND CONFERENCE PUBLICATIONS

- **[J3]** **Ramjee S.**, Ju S., Yang D., Liu X., El Gamal A., Eldar Y.C., "Ensemble Wrapper Subsampling for Deep Modulation Classification". IEEE Transactions on Cognitive Communications and Networking, Aug. 2021 [LINK]
- **[J2]** Wang X., Ju S., Zhang X., **Ramjee S.**, El Gamal A., "Efficient Training of Deep Classifiers for Wireless Source Identification using Test SNR Estimates". IEEE Wireless Communication Letters, Apr. 2020 [LINK]
- **[C1]** Zhang X., Seyfi T., Ju S., **Ramjee S.**, El Gamal A., Eldar Y.C., "Deep Learning for Interference Identification: Band, Training SNR, and Sample Selection". IEEE Signal Processing Advances in Wireless Communications, Jul. 2019 [LINK]
- **[J1]** **Ramjee S.**, Ju S., Yang D., Liu X., El Gamal A., Eldar Y.C., "Fast Deep Learning for Automatic Modulation Classification". IEEE Machine Learning for Communications Emerging Technologies Initiatives, Jan. 2019 [LINK]

PROJECTS

FLITE: Focusing LITE for Memory-Efficient Meta Learning | Meta-Learner — Dec. 2021
- FLITE is a meta learner wrapper that enables the formation of better gradient estimates on a random subset of the support set during meta training for fast and memory-efficient learning [LINK]

AdaLA: Adapting Gradient Estimation by Looking Ahead | ML Optimizer — Jun. 2021
- AdaLA is an optimizer for neural networks that is a simple modification to AdaBelief that uses "looking-ahead" to adaptively estimate the step sizes of the gradients to be taken during gradient descent [LINK]

HONORS AND AWARDS

- Best Product Opportunity Assessment (POA) Award | Stanford University | Nov. 2021
- Tech Fellowship | Stanford TreeHacks | Jun. 2021