

Arvind Rajaraman

arvind.rajaraman@berkeley.edu | [linkedin.com/in/arvind-rajaraman](https://www.linkedin.com/in/arvind-rajaraman) | github.com/arvindrajaraman | (732) 397-2736

EDUCATION

University of California, Berkeley

Aug 2020 – Dec 2023

Dual Bachelor of Science – Electrical Engineering and Computer Science (Honors), Business Administration

- Management, Entrepreneurship, and Technology Program (M.E.T.) – less than 1% acceptance rate
- **CS Coursework:** Deep Reinforcement Learning · Algorithmic Human-Robot Interaction · Computer Vision · Machine Learning · Artificial Intelligence · Operating Systems · Efficient Algorithms · Computer Security · Computer Architecture · Data Structures
- **Math Coursework:** Optimization Modeling · Control Theory · Probability and Random Processes · Linear Algebra · Discrete Math
- **Organizations:** Machine Learning at Berkeley · Berkeley Consulting · Accel Scholars · Neo Scholars · EECS Honors Program

EXPERIENCE

Atlassian – Machine Learning Scientist Intern

May 2023 – Aug 2023

- Researched and developed a diverse re-ranking algorithm for search results, which involves generating embeddings, setting up a graph representation, performing a random walk on candidates, and maximal marginal relevance scoring to trade-off novelty with relevance.
- Trained a SQL schema relevance model, which determines the most relevant SQL columns for a natural language query, by fine-tuning several HuggingFace LLMs and a task-specific feed-forward head. Achieved F1 score of 0.91, beating previous score by 21 pp.
- Developed an RLAIIF (reinforcement learning with AI feedback) system to develop better candidate queries on Atlassian domain data.
- Created a service that automatically validates and deploys PyTorch models to Amazon SageMaker, increasing model uptime by 80%.

Berkeley Artificial Intelligence Research – Machine Learning Researcher

Oct 2022 – Present

- Researching methods for robots to understand human internal biases, goal preferences, and world dynamics using reinforcement learning (for robotic influence) and Transformers (for modeling), with Profs. Anca Dragan (at Berkeley) and Andrea Bajcsy (at CMU).
- Developed a closed-loop training system that combines LQR (linear-quadratic regulators) with a Transformer to perform human internal state estimation and a PPO (proximal policy optimization) agent to perform robot influence on the human.
- Extended previous SoTA work by co-evolving the robot's influence policy with its understanding of the human in closed-loop.
- Created a robust experimentation platform utilizing PyTorch, CUDA, OpenAI Gym, and Wandb to achieve SoTA accuracy.

Nuro – Machine Learning Engineer Intern, Fleet Infrastructure

May 2022 – Aug 2022

- Built an end-to-end HEVC video compression system for thermal cameras, saving Nuro \$900,000 per year in cloud storage costs.
- Engineered a deep learning model using R-CNNs and PointNets to determine the difficulty of identifying car behaviors for labelers.
- Created CI/CD system to monitor model regressions through simulations of edge case scenarios, both real and synthetically generated.
- Developed parallelized, thread-safe data processing infrastructure in C++ using locks and shared pointers to resources.

NVIDIA – Software Engineer Intern, Autonomous Vehicles

Aug 2021 – Dec 2021

- Created image synchronization algorithm that extrapolates missing image frames for in-car cameras, decreasing misalignment by 95%.
- Spearheaded Docker containerizer and gRPC API to serve models at scale with the associated data preprocessing infrastructure.
- Developed a Unix/Linux and C++ script that automates hardware-aware hyperparameter testing of deep CV models on edge devices.

LEADERSHIP AND PROJECTS

Berkeley EECS Department – Head Teaching Assistant

Jan 2022 – Present

- 5-time TA, with 3 as Head TA. Taught CS 189 (Machine Learning) and CS 188 (Artificial Intelligence), which had a max of 900 students and 30 course staff during a semester, with Profs. Jitendra Malik, Stuart Russell, Jonathan Shewchuk, and Jennifer Listgarten.
- Administered course-wide reforms, gave guest lectures/recitations, and received one of the highest teaching ratings among course staff.

Machine Learning at Berkeley – Vice President of Education

May 2022 – Dec 2022

- Developed and instructed CS 198-126 (Modern Computer Vision and Deep Learning), which piloted in Fall 2022. Got sponsorship by Prof. Stuart Russell and received official departmental and university approval. Taught 100 students and had 300 auditors.
- Led educational initiatives in a 120-member organization, managed >\$100,000 in club finances, and organized talks with researchers.

Origin – LLM browser extension that retrieves knowledge from browsing history. Won Stanford TreeHacks & \$125,000 uncapped SAFE.
Unscrambit – predicts algorithms used in program by parsing code and observing inputs/outputs. Won Jump Start Hackathon (1st Place).
Verbal Coding – verbal code editor that enables Python code formulation by spoken pseudocode. Won HackNYU (3rd Place, Education).

SKILLS

- **Languages:** C++, C, Python, Java, HTML, CSS, JS, React, PyTorch, Pandas, Git, AWS, GCP, Unix/Linux, Docker, Protobuf, gRPC
- **Awards:** USACO Gold Division, ACSL (American CS League) - Top 40 worldwide, Clark Scholars Program (0.7% acceptance rate)