

Alok Singh

alokbeniwal@gmail.com • 408-421-5658 • alok.blog • github.com/alok

2042 Hearst Ave Apt C, Berkeley, CA 94709

Education and Coursework

2013-2017 **BA, Mathematics**; University of California, Berkeley

Deep Reinforcement Learning, Machine Learning, Functional Analysis, Real Analysis, Topology, Measure Theory, Artificial Intelligence, Algebraic Topology

Publications

Spiky Corrupt Markov Decision Processes (forthcoming)

Relevant Experience

Machine Learning Consultant, Papert Labs (2018)

Consulted companies on how to use ML. Worked on time series forecasting and automating our ML pipeline, cutting runtime by 2 hours.

Ray, AMPLab (2018)

Maintained PyTorch support for distributed ML computing library **Ray**.

Recurse Center (2017)

Worked on deep reinforcement learning. Implemented DAgger, DQN, and policy gradient methods. Also worked on metaheuristics such as genetic algorithms and simulated annealing.

Data Scientist Intern, Radius Intelligence (2015)

Did EDA using Apache Spark. Integrated customer data with our own, increasing total dataset size by 3x.

Relevant Projects

Network Compression

Implemented model compression to test the conclusion of the paper *Understanding Deep Learning Requires Rethinking Generalization*. Blog post and code [here](#).

Parallel PPO

Implemented advanced policy gradient algorithm and tuned to achieve superhuman performance on Atari.

DAgger

Implemented and gave a talk about DAgger at the Recurse Center.

Deep Q-Network

Implemented DQN with optimizations to play Atari games to superhuman levels.

Skills

- Python, PyTorch, TensorFlow, Bash, UNIX, Haskell