

Alok Singh

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

Projects

BFF: Built an iMessage app around ChatGPT with 10,000 users (Dec 2022)

Voted article of the year by The Browser, over 100,000 reads (Nov 2022)

Experience

Research Resident, Redwood Research (2023)

I'm using my reinforcement and deep learning expertise to mechanistically interpret AlphaGo with a team and probe its adversarial exploitability. We find that AlphaGo often **decides its final move very early** in its hidden layers.

Deep Learning Scientist, Lawrence Berkeley National Lab (2019-2022)

My first project at the lab involved leading a team to replace (slow) classical climate simulation models with GANs and super-resolution. Simulation time dropped from 4 hours to 1 second. This work was published at NeurIPS.

Within a year, I led a team of 10 researchers and programmers for a \$1,800,000 **ARPA-E DIFFERENTIATE project**.

I devised and directed us in building a differentiable electromagnetic simulator in Julia. It evolved meshes to have specified desirable optical properties—a hybrid of flexible deep learning and rigid symbolic equations. It guided a **laser** to ablate materials to specific emissivities (relevant e.g., to solar panels).

Machine Learning Consultant, Papert Labs (2018): Several others and I consulted with Silicon Valley financial startups on using machine and reinforcement learning.

Recurse Center (2017): A self-directed programmer's retreat. I taught myself reinforcement learning and **implemented many algorithms** and gave 2 talks.

Talks and Papers

How to differentiate a discontinuous function

Generalization Properties of Machine Learning Based Weather Model Downscaling, ICLR 2020

A reinforcement learning environment for mathematical reasoning via program synthesis. This used LLMs to output math proofs (checked by the **Lean Theorem Prover**) and was improved by reinforcement learning.

Numerical Weather Model Super-Resolution, NeurIPS 2019

Detecting Spiky Corruption in Markov Decision Processes, IJCAI 2019

Education

2013 - 2017 **Bachelor's, Mathematics, UC Berkeley**

Skills

- Python, Math, Reinforcement Learning, Rust, Theorem Proving, Julia, PyTorch, Bash