Alok Singh

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

Experience

AI/Deep Learning Scientist, Movano Inc. (Aug 2023-):

I am doing deep learning on health data to predict metrics like blood pressure, sleep and blood glucose. I learned about designing energy-efficient device software. I also gleaned valuable insight on the workings of contemporary medical devices.

AI Research Resident, Redwood Research (2023)

I used my deep reinforcement learning knowledge to mechanistically interpret AlphaGo and adversarially probe it. Our team found that AlphaGo often decides its final move very **early** in its hidden layers, usually in the first third.

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 $2M **ARPA-E DIFFERENTIATE project**. This included a team at Georgia Tech where we used computer vision to extract plots and chart data from scientific papers. We also wrote high-performance synthetic data generators. I gained experience with RPC, especially SLURM.

I devised and directed us to build 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).

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

Data Scientist, Radius Intelligence (2015)

I incorporated customer data with our own. I also wrote a lot of Spark and ETL code. On my direction, 3 of

Talks, Papers, and Projects

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

Numerical Weather Model Super-Resolution, NeurIPS 2019

Detecting Spiky Corruption in Markov Decision Processes, IJCAI 2019

How to differentiate a discontinuous function

A reinforcement learning environment for mathematical reasoning via program synthesis. This is an OpenAI Gym environment for math proofs, checked by the **Lean Theorem Prover**.

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

Skills

- DL frameworks and libraries - PyTorch, JAX, Tensorflow, CUDA
- Machine Learning/Deep Learning algorithms such as GPR, SVM, RF, XGBoost, CNN, RNN
- Time-series modeling using LSTM, Transformers
- Programming/Scripting Languages - Python, C++, Rust, Julia, Bash, Lean
- Deep Reinforcement Learning (PPO, SAC, Curiosity), LLMs, Automated Theorem Proving
- Research publications in major conferences (NeurIPS, ICLR, IJCAI)

Education

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