

Aatif Jiwani

aatifjiwani.me

EDUCATION

University of California, Berkeley

Bachelor of Science in

Electrical Engineering & Computer Science

GPA: 3.848

August 2017 - May 2021

Honors at Graduation (Top 20%)

- **Relevant Coursework:** Deep Reinforcement Learning, Deep Neural Networks, Machine Learning, Advanced Probability, Computer Graphics, Robotics, Digital Signal Processing, Operating Systems, Efficient Algorithms, Computer Architecture
- **Awards & Certifications:** Eta Kappa Nu Membership (*IEEE*), Machine Learning by Stanford University (*CourseEra*), Deans List (*2018, 2019*)

PUBLICATIONS *(See website for more details)*

“A Semantic Segmentation Network for Urban-Scale Building Footprint Extraction Using RGB Satellite Imagery”

Aatif Jiwani, Shubhrakanti Ganguly, Chao Ding, Nan Zhou, David M. Chan

arXiv:2104.01263, preprint 2021 - In Submission to IEEE Transactions on Geoscience and Remote Sensing

ACADEMIC EXPERIENCE

Undergraduate Researcher

Lawrence Berkeley National Lab (*Energy Technologies Area*)

Berkeley, CA

April 2020 - May 2021

- Worked on generating high-quality segmentation masks with distinct building footprints from only 3-channel satellite imagery by employing modified boundary losses and an adaptation of the F-Beta Measure as a tunable objective.
- Developed a full-scale PyTorch pipeline for training, testing, and performing inference on the DeepLabV3+ network, and used Weights & Biases and the Lawrencium HPC to track and analyze experiments on over 500 different models.

Undergraduate Researcher

CannyLab (*Assistive AI*) @ Berkeley Artificial Intelligence Research

Berkeley, CA

Jan 2020 - May 2021

- Built a graph-based feature extraction pipeline capable of extracting/processing optical flow and ResNet3D features, as well as dense trajectory feature descriptors for various video datasets such as MSR-VTT and YouCook2.
- Rebuilt video captioning networks like the GRU-EVE and Hierarchical-RNN in PyTorch from reading their publications.
- Modified the state-of-the-art captioning networks to experiment with various latent video representations to achieve higher performance benchmarks and allow the network to capture finer details from the actions in the video.

Tutoring Officer

Eta Kappa Nu

Berkeley, CA

May 2020 - Dec 2020

- Spearheaded the adoption and deployment of a Flask ticketing application to help facilitate online HKN tutoring for over 100 students every week in the midst of the pandemic.
- Tutored students weekly on courses ranging from Data Structures to Deep Neural Networks and Advanced Probability.

Course Tutor (CS 61C: Machine Structures) *Previously Senior Mentor, Junior Mentor*

UC Berkeley College of Engineering

Berkeley, CA

Jan 2019 - May 2020

- Taught 2-3 weekly 1-hour sections with 5-6 students each that consisted of presenting a mini lecture on the week's material and facilitating practice problems while resolving lingering conceptual and practical questions.
- Hosted review sessions and office hours to help students on topics including C programming, the usage of OpenMP and AVX for parallelization, and the abstraction of virtual memory to aid students in coursework and preparing for exams.

WORK EXPERIENCE

Machine Learning Engineer

C3 AI

Redwood City, CA

July 2021 - Present

- Contributing to the model development framework within C3 AI's proprietary model-based AI application development framework which is the core of all C3 AI's pre-built applications. Currently responsible for improving the performance and experience of architecting and deploying productionizable, heterogeneous ML Pipelines, and upgrading the AutoML framework to allow for more flexible hyperparameter optimization and parallel execution strategies.

Software Development Engineer Intern

Amazon (Amazon Web Services)

Remote (Seattle, WA)

Jun 2020 - Aug 2020

- Used the AWS Cloud Development Kit to deploy 2 AWS Lambda functions built using AWS Java SDK to receive on-demand internal requests from AWS service teams and store processed request into respective DynamoDB databases.
- Contributed to an existing Ruby on Rails application by adding new backend services to communicate with various AWS services including IAM, Lambda, and S3 and to process IAM managed policy documents for certain disclosed data.
- Tied the AWS Lambda functions and the Rails backend services together to create an automated, self-service web feature that allows internal service teams to request updates to external-facing IAM managed policies.

Software Engineering Intern

Parnassus Investments

San Francisco, CA

May 2019 - Aug 2019

- Wrote over 300 controller and model validation test cases for the public client website using C#.NET Core and Moq.
- Developed an NLP-based chatbot in 8-weeks using Microsoft's Bot Service .NET SDK, Luis.AI, and QnA Maker to recognize a user's request for fund information, account documents, etc, and provide an accurate response.
- Created a React.js frontend application and configured the chatbot's server response payload to allow users to use both Slack and a browser to interact with the AI-powered bot.

RESEARCH PROJECTS *(See website for more details)*

Physically-Based Lens Flare Distortion Renderer / C++

- Utilized Fourier analysis and ray-tracing to render lens flares in scenes with infinite light sources. By modifying a ray-tracing engine to capture infinite and parallel sources of light, we created a renderer that can produce artificial but physically realistic lens-flare distortions. The implementation was split into 2 distinct parts: (1) we used the discrete-time Fourier Transform to generate the "starburst effect," the effect responsible for protruding rays, and (2) used ray-tracing to determine where incoming rays hit the sensor plane to produce a phenomenon called "ghost reflections."

Distributed Matrix Inversion of $A^T A$ - Open Source Contribution to NumS / Python (Ray)

- Implemented distributed matrix operations to speed up computation for linear least-squares regression. Using the NumS Python framework and Ray as a distributed back-end, we implemented algorithms that distribute computation for upper triangular matrix inversion, and LU and Cholesky decomposition across a cluster of machines. We interleaved these operations to improve the speed of least squares regression on tall-skinny data matrices, where the number of data points severely overwhelms the number of features, versus performing naive matrix-multiply and inversion.

Video Summarization using REINFORCE / Python (PyTorch)

- Employed deep RL to achieve better F-Score benchmarks in video summarization. With inspiration from Kaiyang Zhou, we created a paired unsupervised and supervised pipeline to train a network to select key frames of interest from videos using deep RL. With a custom reward function that encouraged clustering and selection of frames with an abundance of objects, we tested our network on the SumMe and TVSum video summarization datasets and achieved an F-Score of 45.5% and 58.9% respectively, beating the previous benchmark by 4.1% and 1.3% respectively.

Yelp Rating Prediction / Python (PyTorch, HuggingFace)

- Used NLP to predict the star rating of any given Yelp review for any business. Built 3 custom PyTorch models including a Bi-LSTM with a decoder and utilized pre-trained HuggingFace transformers. After training and evaluation on the Yelp dataset, our peak performance neared 80% accuracy on the full dataset and 60% on challenge datasets.

RoBob The Builder / Python (ROSPy, OpenCV)

- Built robotic simulation to recreate simple structures autonomously. Simulated the PAL Robotics TiAGo robot model to perform 2 primary tasks: (1) in an initial world with a desired structure, find the 3D coordinate points of the blocks that make-structure, and (2) autonomously navigate to and pick up randomly placed blocks and place them in a way that replicates the desired structure. Used OpenCV for the imaging simulation and used ROSPy and ArUco for the kinematics simulation.

HIGHLIGHTED SKILLS

- **Programming Languages:** Python (*Flask, SciPy, PyTorch, Tensorflow, OpenCV, Ray, ROSPy, Spark*), Java (*JavaFX*), Javascript (*React*), Ruby (*Ruby on Rails*), C++ (*UPC++, CUDA*), C# (*.NET Core*), C, Swift, SQL, Git, Shell/Bash
- **External Services:** Amazon Web Services (*EC2, S3, Lambda, DynamoDB, Sagemaker, Polly*), Microsoft Azure (*Luis.AI, QnA Maker*), Google Cloud Platform (*Storage, Firebase, Compute Engine*), Docker, Kubernetes