

Gokul Swamy

gokul.swamy@berkeley.edu · gokul.dev · (858) 705-9269

EDUCATION

University of California, Berkeley

B.S. Electrical Engineering & Computer Science, GPA: 3.84

August 2016 - May 2019

M.S. Computer Science

August 2019 - May 2020

Select Classes: Machine Learning, Artificial Intelligence, Computer Vision*, AI Safety*, Optimization, Probability, Information Theory*, Algorithms, Networking, Operating Systems, Security (* = graduate)

WORK EXPERIENCE

SpaceX, Data Engineering Intern ([link](#))

May 2018 - August 2018

- Built set of NLP utilities for querying, cleaning, vectorizing, and reducing dimensionality of text.
- Built word2vec-backed semantic similarity engine to autofill fields and detect anomalous entries.
- Used engine to build approximate nearest neighbor algorithm to estimate operation duration.
- Created computer vision algorithm to detect not-for-flight tags and other flight risks.

InterACT Lab (part of Berkeley AI Research), Research Assistant ([link](#))

January 2018 - Present

- Working for Prof. Anca Dragan investigating the relative performance of theory-of-mind, model-based, and model-free deep-RL agents in self-driving context. Published in HRI 2019 ([arXiv](#)).
- Created driving simulator, implemented various human driver models with learned response parameters, and built policy gradient-based (PPO) agent that interacts with human models.

Intuit, Software Engineering Intern ([link](#))

May 2017 - August 2017

- Used Latent Semantic Analysis, CNNs, and Character CNN's to classify product reviews into complaint groups to identify areas of consumer dissatisfaction.
- Created iOS UI component framework with theming, grouping, and live prototyping in Swift.

SELECT PROJECTS (see [GitHub](#) and [website](#) for full list)

Generative Models for Pose Transfer, Computer Vision Class Project ([link](#))

- Built algorithm to transfer action (e.g. dance) from one subject to another (e.g. person to person).
- Used pose skeleton detection algorithm and k-NN to match frames to train a GAN (pix2pix).

Take a Picasso, CalHacks 3.0 Project ([link](#))

- A robotic sketch artist that draws a physical portrait of user based on image taken by smartphone.
- Detects a voice command, performs Canny edge detection in OpenCV to create a drawable image, and uses a Traveling Salesman approach to create a vectorized image for the Arduino to draw.
- Won Best Hardware, Best 3D Printed Hack, and a grant from Peter Thiel's 1517 Fund.

CS101: Application Design and Development for iOS, MOOC

- Helping create online course to be released to the public through Berkeley's edX page.
- Giving lectures on Swift, MVC, UI/UX, Networking, Augmented Reality, and Machine Learning.

ACTIVITIES

Machine Learning @ Berkeley, DeCal Facilitator ([link](#))

August 2016 - Present

- Teaching course on the impact of AI on society and potential responses to dilemmas posed
- Previously lead MOOC development, managed NLP consultant project for Intuit, gave workshops on machine learning for mobile devices, and consulted for IBM on NLP and Unity on Deep RL.

HONORS

- Member of Eta Kappa Nu (EECS Honor Society)
- 2016 & 2017 WWDC Scholar (granted by Apple based on iOS development project)
- Valedictorian of Del Norte HS Class of 2016
- Regents and Chancellors Scholar (granted to top 2% of each incoming UC Berkeley class)

SKILLS

- Languages: Python, Swift, SQL, Java, HTML, CSS, Javascript, C, RISC-V, Go
- Tools: Tensorflow, scikit-learn, XGBoost, PowerBI, React, Bootstrap, CoreML, Firebase