

Arjun Subramonian (they/them)

UCLA Computer Science Department

University of California, Los Angeles

405 Hilgard Ave

Los Angeles, CA 90095 USA

Phone: 408-859-2148

Email: arjunsub@cs.ucla.edu

Website: arjunsubramonian.github.io

Google Scholar: [publications](#)

LinkedIn: [linkedin.com/in/arjuns22](https://www.linkedin.com/in/arjuns22)

Medium: medium.com/@arjunsuboo

GitHub: github.com/ArjunSubramonian

Semantic Scholar: [publications](#)

Research Interests

graph machine learning, large language models, fairness, privacy, ethics

Education

PhD in Computer Science; University of California, Los Angeles (2021-), Eugene V. Cota-Robles Fellow

Advisors: Yizhou Sun, Kai-Wei Chang

BS in Computer Science; University of California, Los Angeles (2018-2021), *GPA:* 3.927, Summa Cum Laude

Work Experience

Machine Learning Researcher, UCLA Scalable Analytics Institute (2019-)

Location: Los Angeles, California

Machine Learning Researcher, UCLA NLP (2020-)

Location: Los Angeles, California

Research Engineering Intern, AllenNLP, Allen Institute for Artificial Intelligence (2021)

Location: Seattle, Washington

Description: I developed [AllenNLP's fairness library](#), which makes fairness metrics, training-time fairness algorithms, bias mitigation algorithms, and bias metrics accessible to researchers and practitioners of all levels. I also wrote a [guide chapter](#), [documentation](#), and a [blog post](#) to communicate my work and make usage of the fairness library accessible.

Privacy Research Intern, Snap, Inc. (2021)

Location: Los Angeles, California

Description: Researched algorithms to improve safety of friend suggestions for underage users while preserving privacy. Contributed to development of Snap's Responsible AI principles. Machine learning for ads and monetization.

Software Engineering Intern, Microsoft Corporation (2020)

Location: Sunnyvale, California

Description: I crafted a peer-to-peer-anonymous, secure backend technical design for a feature to report harassment on Microsoft Teams.

Software Engineering Intern, Get Heal, Inc. (2019)

Location: Los Angeles, California

Description: I engineered full-stack integrations of mechanisms used every day at Heal that enhance the automated routing of medical providers, like automated triaging, doctor-assistant match prevention, and phone number verification. I also adapted Heal's automated routing algorithm to optimally schedule telemedicine visits, which greatly benefits patients during the COVID-19 pandemic.

Deep Learning Engineer, Sike AI (2018-2019)

Location: Los Angeles, California

Description: I designed, implemented, and trained the in-house deep learning model for working style-analysis from video with TensorFlow.

Publications

Dev, Sunipa, Masoud Monajatipoor*, Anaelia Ovalle*, **Arjun Subramonian***, Jeff M Phillips, and Kai-Wei Chang. "**Harms of Gender Exclusivity and Challenges in Non-Binary Representation in Language Technologies.**" Accepted to EMNLP 2021.

Subramonian, Arjun. "**Fairness and Bias Mitigation: A practical guide into the AllenNLP Fairness module.**"

Zhang, Shichang, Ziniu Hu, **Arjun Subramonian**, and Yizhou Sun. "**Motif-Driven Contrastive Learning of Graph Representations.**" Accepted to SSL@WWW2021.

Subramonian, Arjun. "**MOTIF-Driven Contrastive Learning of Graph Representations.**" Accepted to Undergraduate Consortium @ AAAI 2021.

Brown, Calvin, Derek Tseng, Paige M. K. Larkin, Susan Realegeno, Leanne Mortimer, **Arjun Subramonian**, Dino Di Carlo, Omai B. Garner, and Aydogan Ozcan. "**Automated, Cost-Effective Optical System for Accelerated Antimicrobial Susceptibility Testing (AST) Using Deep Learning.**" ACS Photonics 2020 7 (9), 2527-2538 DOI: 10.1021/acsp Photonics.0c00841

Crandall, Sara, Graeme H. Smith, **Arjun Subramonian**, Kelly Ho, and Evelyn M. Cochrane, "**Estimating the Ages of FGK Dwarf Stars Through the Use of GALEX FUV Magnitudes.**" Astronomical Journal 2020 160, 217, DOI: <https://doi.org/10.3847/1538-3881/abb77d>

QueerInAI, Organizers of, Ashwin*, William Agnew*, Juan Pajaro*, and **Arjun Subramonian***. "**Rebuilding Trust: Queer in AI Approach to Artificial Intelligence Risk Management.**" NIST AI Risk Management RFI.

QueerInAI, Organizers of, A Pranav, MaryLena Bleile, **Arjun Subramonian**, Luca Soldaini, Danica Sutherland, Sabine Weber, Pan Xu, William Agnew, Michael McKenna, and Nyx McLean. "**How to Make Virtual Conferences Queer-Friendly: A Guide.**" Accepted to WiNLP 2021 Workshop @ EMNLP 2021.

Subramonian, Arjun. "**Queer | Inclusive | Badass.**" Accepted to Resistance AI Workshop @ NeurIPS 2020.

Talks and Panels

2021 - Harms of Gender Exclusivity and Challenges in Non-Binary Representation in Language Technologies, EMNLP 2021

2021 - Safer Privacy-Preserving Friend Suggestions, Snap, Inc.

2021 - Machine Learning Justice, Catalysts for Change

2021 - How Can I Make My Hackathon Queer-Inclusive? (Slides, Video), Hackcon IX

2021 - Intersectionality Panel, NAACL 2021

2021 - Queer in AI Inclusive Conference Guide DEI Update, Allen Institute for Artificial Intelligence

2021 - Queer in AI Panel, UCLA

2020 - Fair Machine Learning, Microsoft Garage Brown-Bag

2019 - An Automated and Cost-Effective System for Early Antimicrobial Susceptibility Testing Using Optical Fibers and Deep Learning, UCLA HHMI Day 2019

Honors and Awards

2021 - MLH Top 50 Class of 2021
2021 - UCLA Samueli School-Wide Outstanding Bachelor of Science
2021 - UCLA Chancellor's Service Award
2021 - UCLA Samueli Engineering Achievement Award in Student Welfare
2021 - Eugene V. Cota-Robles Fellowship, UCLA
2021 - Graduate Research Assistantship, UCLA
2021 - Boeing Company Scholarship, UCLA
2021 - Brian J. Lewis Endowment, UCLA
2020 - Computing Research Association Outstanding Undergraduate Researcher Honorable Mention
2020 - AAAI Undergraduate Consortium (1 of 14 accepted out of 82 applicants)
2020 - IBM Quantum Challenge (1 of 574 winners out of 1745 participants)
2020 - Out for Undergrad Tech Conference (1 of 300 accepted applicants)
2020 - Google Queer Tech Voices Conference (1 of 32 accepted out of hundreds of applicants)
2019 - 3rd Place Award for Best Hack @ Rose Hack, Major League Hacking
2018-2021 - Dean's Honors List
2017 - Siemens Competition Regional Finalist (1 of 101 finalists selected from 4092 entrants)
2016 - Award of Achievement, Association for Computing Machinery, San Francisco Bay Area Professional Chapter

Research Projects

Species Reidentification via Graph Neural Networks (2021-)

Collaborators: Sara Beery, Peter Kulits, Jake Wall

Location: Computational Vision Laboratory

Description: I am applying graph neural networks to elephant co-sightings to contextualize and improve elephant reidentification.

Analyzing Biases via Attention Flows (2021-)

Collaborators: Sunipa Dev

Location: UCLA-NLP

Description: I am exploring how attention flows in large language models correlate with bias metrics and are affected by bias mitigation.

Expressive Graph Transformers (2020-)

Collaborators: Yizhou Sun

Location: UCLA Scalable Analytics Institute

Description: I'm empirically and theoretically studying the effect of different types of handcrafted and adaptive relational information for relation-aware self-attention on improving the expressiveness and performance of graph Transformers, particularly on NP-hard graph problems. As part of this project, I implemented and train a multi-GPU graph Transformer model using PyTorch. I'm also proving the logical expressiveness of graph Transformers.

Learning Fair Node Representations for Social Networks (2020-)

Collaborators: Kai-Wei Chang, Yizhou Sun

Locations: UCLA-NLP, UCLA Scalable Analytics Institute

Description: I am researching interpretable methods to provably debias pretrained graph convolutional networks (GCN). I am also theoretically analyzing bias amplification in message-passing networks (e.g. GCN) when the minority-class nodes are outliers in terms of node connectivity. Finally, I am interested in learning fair node representations without demographics and researching the effects of partially debiasing GCNs on fairness in downstream tasks.

Motif-Driven Contrastive Learning of Graph Representations (2020)

Collaborators: Shichang Zhang, Ziniu Hu, Yizhou Sun

Location: UCLA Scalable Analytics Institute

Description: We propose MICRO-Graph to: 1) pre-train Graph Neural Networks (GNNs) in a self-supervised manner to automatically extract graph motifs from large graph datasets; 2) leverage learned motifs to guide the contrastive learning of graph representations, which further benefit various graph downstream tasks.

Heterogeneous Graph Transformer (2019-2020)

Collaborators: Ziniu Hu, Yizhou Sun

Location: UCLA Scalable Analytics Institute

Description: I adapted the implementation of the Heterogeneous Graph Transformer (HGT) to efficiently embed web-scale knowledge graphs (e.g. YAGO, DBpedia) for link prediction and ran R-GCN baselines. Additionally, I prepared an OGB leaderboard submission in which I applied HGT to the ogbl-ppa dataset.

Automated, Cost-Effective Optical System for Accelerated Antimicrobial Susceptibility Testing (AST) Using Deep Learning (2018-2019)

Collaborators: Calvin Brown, Aydogan Ozcan

Location: UCLA Ozcan Research Group

Description: I designed a neural network that inexpensively and automatically detects bacterial resistance to antibiotics, which shortens the timeline of prescribing antibiotics to patients in resource-limited settings by about 60%, helping to mitigate the rise of global antimicrobial resistance. I implemented and trained the neural network with Python and Keras, tuning hyperparameters and visualizing learning curves, weights, and hidden-layer activations, achieving FDA essential agreement for 99.5% of drugs.

Estimating the Ages of FGK Dwarf Stars Through the Use of GALEX FUV Magnitudes (2017)

Collaborators: Kelly Ho, Graeme Smith, Sara Crandall

Location: University of California, Santa Cruz, Smith Lab

Description: I applied machine learning and statistics to discover a novel method for estimating the age of FGK dwarf stars using GALEX far-ultraviolet (FUV) magnitudes that is more cost and time-efficient, as well as more accessible, than existing methods.

An Empirical Characterization Of Internet Round-Trip Times (2016)

Collaborators: Ishani Karmarkar, Alice Lim, Daniel S.F. Alves

Location: University of California, Santa Cruz, Internetworking Research Group (i-NRG)

Description: I applied machine learning to develop an RTT boundary prediction algorithm which employs online linear regression to predict future RTTs, thereby reducing unnecessary packet retransmissions, delays in retransmission, and overall network congestion.

Course Projects

Twitter Saliency Algorithm: Identifying Unintentional Harms to Gender Non-Conforming Individuals (2021)

Collaborators: Michael McKenna

Description: We attempted to uncover unintentional harms of the Twitter saliency algorithm, e.g. 1) identifies images of potentially-cis or binary-presenting individuals as more salient than those of gender non-conforming folks, 2) identifies undesirable secondary sex characteristics of gender non-conforming individuals that may trigger body dysphoria.

Notes: [Report](#)

Robust Model-Agnostic Meta-Learning for Binary Content Moderation Tasks in Natural Language Processing (2020)

Collaborators: John Dang, Kai-Wei Chang

Location: University of California, Los Angeles

Description: We investigated applying Model-Agnostic Meta-Learning (MAML) to boost performance on binary content moderation tasks in low-resource contexts. Using PyTorch, we compared the ability of a model pre-trained with MAML to adapt to unseen binary content moderation tasks to those of a model pre-trained using traditional transfer learning approaches and a model trained from scratch.

Notes: [Report](#)

On the Complexity and Convergence of Approximate Policy Iteration Schemes (2020)

Collaborators: Shree Kesava Narayan Prasanna, Nikil Roashan Selvam, Justin Yi, Professor Lin Yang

Location: University of California, Los Angeles

Description: We surveyed relevant literature in approximate policy iteration, and provided theoretical proof sketches involved in the analysis of the complexity bounds, convergence guarantees, and rates of convergence for various approximate policy iteration algorithms.

Notes: [Report](#), [Poster](#)

Model-Agnostic Meta-Learning for a Policy Gradient Approach to MuJoCo Continuous Control Tasks (2020)

Collaborators: John Dang, Jonathan Kao

Location: University of California, Los Angeles

Description: We investigated the adaptive power of Model Agnostic Meta-Learning on a policy gradient approach to MuJoCo continuous control tasks.

Notes: [Report](#)

Quantum Programming Algorithms (2020)

Collaborators: Vaishnavi Tipireddy, Siddarth Chalasani, Jens Palsberg

Location: University of California, Los Angeles

Description: We implemented Deutsch-Jozsa, Bernstein-Vazirani, Grover's algorithm, and Simon's algorithm using PyQuil and Qiskit. We then evaluated the implementations and modern quantum compile and runtime capabilities using the Rigetti and IBM quantum simulators and IBMQX quantum devices.

Notes: [PyQuil Report](#), [Qiskit Report 1](#), [Qiskit Report 2](#)

MovieLens Recommender System (2019)

Collaborators: Amit Mondal, Bryan Chiang, John Dang, Jyun-Yu Jiang, Wei Wang

Location: University of California, Los Angeles

Description: We created a recommender system to predict the binary rating for 4M unseen UserID-MovieID pairs in the MovieLens dataset. We surveyed the performance of content-based (e.g. TF-IDF, genre-based decision tree, etc.) and collaborative filtering (e.g. SVM, SVD, element-wise matrix factorization, tabular matrix factorization, hybrid matrix factorization, etc.) methods. **We achieved the third highest ROC-AUC on the test set in our data mining class.**

Notes: [Report](#)

Service

Core Organizer, Queer in AI (2021-)

Location: Virtual

Description: I organize workshops and socials at AI conferences (e.g. [AAAI-21](#), [ICML '21](#), [NeurIPS 2021](#)), as well as the [undergraduate mentoring program](#), which gets junior queer and trans folks involved with AI research and aids them in [applying to graduate school](#). Additionally, I advise AI conferences on [diversity and inclusion and accessibility issues](#). The work I do with Queer in AI has been featured by [500 Queer Scientists](#).

Accessibility Chair, NAACL 2022 (2021-)

Location: Seattle, Washington

Description: I am serving as an Accessibility Chair on [NAACL 2022's Diversity and Inclusion committee](#).

Queer and Trans in STEM Representative, UCLA Samueli Standing Committee on Diversity (2021-)

Location: University of California, Los Angeles

Description: I am working towards dropping the GRE requirement for graduate school admissions.

UCLA Engineering Scholarship Application Reviewer (2021)

AllenNLP Hacks Organizer, AllenNLP (2021)

Location: Seattle, Washington

Description: I helped organize [AllenNLP Hacks](#), a hackathon to connect with marginalized students, welcome them into AllenNLP's

open-source community, bring their perspectives to AllenNLP's research, and encourage them to apply to intern and work with [AllenNLP](#).

Organizer, UCLA Computer Science Summer Institute (2021)

Location: Los Angeles, California

Description: I interviewed and recruited a diverse group of Undergraduate Tutors for the inaugural [UCLA Computer Science Summer Institute \(CSSI\)](#) to lead interactive coding and problem-solving sessions with the high school students.

Outreach Director, ACM AI at UCLA (2019-2021)

Location: Los Angeles, California

Description: I strive to make an [AI education](#) accessible to everyone. I created, led, and taught open-source, accessible [machine learning](#) and [AI ethics](#) classes at Title I schools in LA, through in-person visits, virtual sessions, and educational technology (e.g. [mean-squared error](#), [convolutional filters](#), [biases in machine learning](#), etc.) I also created and produced the "You Belong in AI!" podcast, which empowers marginalized youth to pursue AI opportunities through inspiring interviews with researchers. The podcast has been featured by the [Daily Bruin](#) and [UCLA Samueli Newsroom](#).

Co-Founder and Organizer, QWER Hacks (2019-2021)

Location: Los Angeles, California

Description: I co-founded and organized Major League Hacking's first-ever LGBTQIA+ event and the first student-run, collegiate [LGBTQIA+ hackathon](#) in the US. QWER Hacks has been featured by the [Daily Bruin](#) and the [UCLA Samueli Newsroom](#).

Undergraduate Learning Assistant (2018)

Location: Los Angeles, California

Description: I led weekly recitation sections of 20 students for the introductory computer science class (programming in C++), walking through practice problems and actively applying pedagogy techniques (e.g. open questioning, fostering belonging, etc.)

Coursework and Skills

Graduate-level Coursework: Fairness, Ethics, Accountability and Transparency in Natural Language Processing; Neural Networks and Deep Learning; Reinforcement Learning Theory and Applications; Adversarial Robustness in Machine Learning; Quantum Programming

Relevant Skills: Python, PyTorch, PyTorch Geometric, git, shell scripting, LaTeX