

SOHAM GADGIL

sgadgil@uw.edu | <http://www.sohamgadgil.com> | <https://github.com/sgadgil6>

EDUCATION

UNIVERSITY OF WASHINGTON

- Ph.D. in Computer Science and Engineering

Seattle, WA

Sept. 2022 - Present

STANFORD

- M.S. in Computer Science GPA: **4.056**
- Coursework: Deep Learning, Natural Language Processing, Computer Vision

Stanford, CA

Graduated: June 2021

GEORGIA INSTITUTE OF TECHNOLOGY

- B.S. in Computer Engineering, Minor in CS GPA: **4.0 (Faculty Honors)**
- Selected in the China Summer Program (CSP) for study abroad in Summer 2016

Atlanta, GA

Graduated: May 2019

PUBLICATIONS

Estimating Conditional Mutual Information for Dynamic Feature Selection [\[Paper\]](#) [\[Code\]](#)

*Soham Gadgil**, Ian Covert*, Su-In Lee

arXiv Preprint

Fostering transparent medical image AI via an image-text foundation model [\[Paper\]](#) [\[Code\]](#)

Chanwoo Kim, *Soham Gadgil*, Alex J. DeGrave, Zhou Ran Cai, Roxana Daneshjou, Su-In Lee

medRxiv Preprint

Combining Expert Annotations with DNN-generated Saliency Maps for X-ray Segmentation [\[Paper\]](#) [\[Code\]](#)

*Soham Gadgil**, Mark Endo*, Emily Wen*, Andrew Y. Ng, Pranav Rajpurkar

MIDL 2021

Spatio-Temporal Graph Convolution for Functional MRI Analysis [\[Paper\]](#) [\[Code\]](#)

Soham Gadgil, Qingyu Zhao, Adolf Pfefferbaum, Edith V. Sullivan, Ehsan Adeli, Kilina M. Pohl

MICCAI 2020

Solving The Lunar Lander Problem under Uncertainty using Reinforcement Learning [\[Paper\]](#) [\[Code\]](#)

Soham Gadgil, Yunfeng Xin, Chengzhe Xu

IEEE SoutheastCon 2020

RESEARCH INTERESTS

Computer Vision, Clinical and Explainable AI, Generative Modelling, Interpretability, Foundation Models

RESEARCH EXPERIENCE

Lee Lab of AI for bioMedical Sciences (AIMS) at UW – *Research Assistant*

Sept 2022 - Present

Advised by Dr. Su-In Lee

- AI in Dermatology (In collaboration with Dr. Roxana Daneshjou from Stanford)
 - Developing diffusion-model based XAI methods to manipulate specific attributes in dermatology images
 - Analysing AI-specific signals enabling classifiers to detect protected attributes with high performance
 - Developed an image-text foundation model to automatically annotate concepts in dermatology images
- Dynamic Feature Selection for Emergency Medicine (In collaboration with Dr. Nathan White from UW Medicine)
 - Formulated DIME, a novel information-theoretic approach to estimate conditional mutual information
 - Network learns which features to collect for each patient for efficient and cost-effective diagnosis
 - DIME allows per-sample budgets, enables non-uniform costs, and can be extended to any architecture
 - Our method consistently performed better than recent dynamic and static feature selection methods

Stanford Computational Neuroscience Lab (CNS^{LAB}) – *Research Assistant*

Sept 2019 – March 2020

Advised by Dr. Kilian Pohl

- Used deep learning techniques to perform sex classification from functional-MRI scans
- Formulated the non-stationary nature of functional connectivity within the context of spatio-temporal graphs
- The model beat previous approaches with an accuracy of **83.7%**, accepted into MICCAI 2020

Stanford Machine Learning Group (AI for Healthcare) – *Research Assistant*

Sep 2020 – June 2021

Advised by Dr. Andrew Ng and Pranav Rajpurkar

- Developed CheXseg, a semi-supervised method for multi-pathology segmentation
- CheXseg leverages expert annotations and saliency maps generated by image classification models
- Compared to weak supervision, CheXseg reduces the mIoU gap with radiologists by **71.6%**

Mixed Signal Design Group at Georgia Tech – *Research Assistant*

Aug 2016 - May 2017

Advised by Dr. Madhvan Swaminathan, sponsored by Intel and Semiconductor Research Corporation

- Used machine learning to create behavioural circuit models by macro-modelling of transistor models
- Designed and trained a feed forward neural network using Python to reduce complexity up to 14x

Georgia Tech Bionics Lab – Research Assistant

Aug 2017 – May 2018

Led by Dr. Ghovanloo, sponsored by NSF and the National Institutes of Health

- Worked on developing an assistive technology for speech impaired people
- Developed an algorithm to track lips in real time using the picamera
- Optimized video frame transfer from the camera to the system by 5 fps using the Micro-USB OTG port

WORK EXPERIENCE

Microsoft – Data Engineer

Jul 2021 – August 2022

- Worked on a small, fast-paced team to provide automation tooling for Windows experience
- Served ~1500 customers with a RESTful smart service to spin up secure cloud VMs for OS development
- Saved over 300 hours of developer time spent on setting up machines, access policies, and repo cloning

Microsoft – Software Engineering Intern

Jun 2020 – Sep 2020

- Worked in the Windows Toolkits team on an Azure hosted web portal to automate backporting bug fixes
- Developed a RESTful web API using .NET Core 3.1 with a Model-View-Controller (MVC) design pattern
- Automated web-app deployments using custom CI/CD pipelines and improved team productivity by 15%

Microsoft – Program Manager Intern

May 2019 – Aug 2019

- Worked on the Windows Updates team to analyse telemetry data in the order of petabytes
- Designed and implemented quantifiable metrics to optimize release decisions for newer Windows builds
- New system reduced token utilization in compute resources by 10%

Goldman Sachs – Technology Analyst Intern

June 2018 – Aug 2018

- Developed a backend application in Java to merge legal entities with metadata in over 100 tables
- Established a stream to collect, validate, and process data events by making REST compliant API calls

Bank of America – Software Development Intern

June 2017 – Aug 2017

- Worked on full stack development for the migration of an internal application
- Established necessary DDL, DML trigger scripts using T-SQL and developed requisite data models
- Developed SSIS packages to transfer over 1,000,000 records from Oracle to SQL server

TEACHING

School of Computer Science and Engineering at University of Washington

Jan 2021 – June 2021

- Introduction to AI (CSE 473): TA for undergraduate course with ~150 students
- Created a new homework assignment and held office hours every week

School of Computer Science at Stanford

Jan 2021 – June 2021

- Trustworthy ML (CS 329T): TA for a new course being offered for the first time with ~50 students
- Created and graded homework, developed materials for the lab sections, and managed course logistics

School of Mathematics and ECE at Georgia Tech

Aug 2017 – May 2019

- Math 1554 - Linear Algebra: Taught two 50-minute recitations each week (Aug 2017 – Dec 2017)
- ECE 3056 - Computer Architecture, Concurrency, and Energy in Computation (Jan 2018 – Dec 2018)
- ECE 2035 - Programming for hardware and software systems (Jan 2017 – May 2017)
- Collaborated with the lead instructor for grading and held weekly office hours

PROJECTS

Evaluating and Mitigating Geographical Bias in CLIP

June 2023

- Demonstrated that CLIP embeddings are substantially biased toward countries with high incomes
- For zero-shot classification, top-5 countries with high performing images have **49.5x** higher income than bottom-5
- Explored techniques like prompt engineering and orthogonal projection to reduce Gini Index by ~16%

Customized Prompts for Zero-Shot Concept Generation in Dermatology Images

Dec 2022

- Trained a CLIP model for concept classification using image-caption pairs from PubMed articles and textbooks
- Improved the text by passing the raw captions through a LLM (GPT-3) fine-tuned using the field's several textbooks
- New model beat the original model with raw captions by **2%** with a mean AUROC of **0.648** across all concepts

Using Deep Learning for ICD-9 Code Classification of Medical Notes

Dec 2019

- Analysed different methods to mitigate inherent geographical biases in SOTA image classification models
- Quantitatively presented this bias in ImageNet by fine tuning two popular models, VGG16 and ResNet-18
- Used methods like focal loss and uniform sampling to reduce variation in model accuracy from **9%** to **5%**

Home Depot Deep Learning Hackathon

April 2017

- Developed a convolutional neural network in TensorFlow to categorically sort product images
- Used transfer learning to retrain Google's Inception-v3 model using the new dataset with **80%** accuracy

SKILLS

- **Languages:** Python, Java, C++, C#, MATLAB, T-SQL, PL/SQL, HTML, CSS, JavaScript
- **Frameworks:** PyTorch, NumPy, Pandas, Scikit-Learn, Hugging Face, Tensorflow
- **Workload Manager:** Slurm