

Saurabh Garg

Machine Learning Department
School of Computer Science
Carnegie Mellon University

Google Scholar Profile
<https://saurabhgarg1996.github.io>
✉ sgarg2@andrew.cmu.edu

Education

Ph.D. in Machine Learning. 2019 – Present

School of Computer Science, Carnegie Mellon University (CMU)

GPA: 4.19/4.33

Advisors: Zachary Lipton, Sivaraman Balakrishnan

Awards: JP Morgan AI PhD Fellowship, Amazon Graduate Research Fellowship

Bachelors (with honors) in Computer Science and Engineering. 2014 – 2018

Minor in Applied Statistics and Informatics,
Indian Institute of Technology (IIT) Bombay

GPA: 9.51/10.0

Awards: Excellence in Research Award (1 among 100 students), Institute Academic Award

Selected Research & Publications

Overview: Published >15 papers (five competitive oral and spotlight presentations awards) and 2 journal papers in machine learning and its applications in venues such as NeurIPS, ICLR, ICML, EMNLP and MICCAI. Work in my main line of research includes:

RLSbench: Investigating Domain Adaptation Methods Under Relaxed Label Shift. [Saurabh Garg](#), Nick Erickson, James Sharpnack, Alex Smola, Sivaraman Balakrishnan, Zachary C. Lipton

Domain Adaptation under Open Set Label Shift. [Saurabh Garg](#), Sivaraman Balakrishnan, Zachary Lipton. ICML SCIS Workshop, 2022. NeurIPS 2022.

Leveraging Unlabeled data to Predict Out-of-Distribution Performance. [Saurabh Garg](#), Siva Balakrishnan, Zachary Lipton, Behnam Neyshabur, Hanie Sedghi. NeurIPS DistShift Workshop 2021. ICLR 2022.

RATT: Leveraging Unlabeled Data to obtain Generalization Guarantees. [Saurabh Garg](#), Zico Kolter, Sivaraman Balakrishnan, Zachary Lipton. ICLR Robust ML Workshop, 2021. *ICML 2021 (Long Oral)*.

Mixture Proportion Estimation and PU Learning: A Modern Approach. [Saurabh Garg](#), Yifan Wu, Alex Smola, Sivaraman Balakrishnan, Zachary Lipton. *ICML UDL Workshop, 2021. NeurIPS 2021 (Spotlight)*.

A Unified View of Label Shift Estimation. [Saurabh Garg](#), Yifan Wu, Sivaraman Balakrishnan, Zachary Lipton. *ICML UDL Workshop, 2020 (Oral). NeurIPS 2020*.

Selected Awards & Honors

JP Morgan PhD Research Fellowship	2022-23
Amazon Graduate Research Fellowship	2022-23
Top Reviewer Award, NeurIPS	2022
Invited to attend Deep Learning Theory Summer School at Princeton (remote)	2021
Excellence in Research Award (1 among 110 students) from CSE dept, IIT Bombay	2018
Undergraduate Research Award, IIT Bombay	2018
Institute Academic Award, IIT Bombay	2015
All India Rank 93 in JEE Main (out of 1.4 million)	2014
All India Rank 154 in JEE Advanced (out of 126k)	2014

Work Experience

Amazon AWS	Santa Clara, CA (remote)
Student Researcher under Alex Smola	May '22 – ongoing

- Investigating domain adaptation methods under relaxed label shift

Google Brain

Mountain View, CA (remote)

Student Researcher under Hanie Sedghi and Behnam Neyshabur

Sept '21 – Dec '21

Research Intern under Hanie Sedghi and Behnam Neyshabur

June '21 – Aug '21

- Real-world machine learning deployments are characterized by mismatches between the training and test distributions that may cause performance drops. Developed a method for predicting the target domain accuracy using only labeled source data and unlabeled target data.

Samsung Research HQ

Suwon, South Korea

Research Engineer

Sept. '18 – July '19

Research Intern

May '17 – July '17

- Explored AI-based decision making and close loop automation policies for intelligent 5G network deployment. Developed a RL framework for self-learning algorithms that are able to learn the network behaviour.

Microsoft Research

Bangalore, India

Research Intern with Sunayana Sitaram

Dec '17

- Lack of conversational monolingual Hindi text is a major issue in building a powerful Language Model
- Developed a robust transliteration system to utilize large amounts of Roman text data from the web.

Publications

Pre-print/Workshop

P5. RLSbench: Investigating Domain Adaptation Methods Under Relaxed Label Shift

Saurabh Garg, Nick Erickson, James Sharpnack, Alex Smola, Sivaraman Balakrishnan, Zachary C. Lipton

Short version at NeurIPS Workshop on Distribution Shifts (DistShift), 2022

Longer version under submission

P4. CHiLS: Zero-shot Image Classification with Hierarchical Label Sets

Zachary Novack*, Saurabh Garg*, Zachary C. Lipton

Under Submission

P3. Downstream Datasets Make Surprisingly Good Upstream Corpora

Kundan Krishna, Saurabh Garg, Jefferey Bigham, Zachary C. Lipton

Short version at NeurIPS Workshop on Transfer Learning for NLP, 2022

Longer version under submission

P2. Disentangling the Mechanisms Behind Implicit Regularization in SGD

Zachary Novack, Simran Kaur, Tanya Marwah, Saurabh Garg, Zachary Lipton

Spotlight at NeurIPS Workshop on The Benefits of Higher-Order Optimization in Machine Learning, 2022

Longer version under submission

P1. Deconstructing Distributions: A Pointwise Framework of Learning Performance

Gal Kaplun*, Nikhil Ghosh*, Saurabh Garg, Boaz Barak, Preetum Nakkiran

Short version at NeurIPS Workshop on Distribution Shifts (DistShift), 2022

Longer version under submission

Conference

C11. Domain Adaptation under Open Set Label Shift

Saurabh Garg, Sivaraman Balakrishnan, Zachary Lipton

Advances in Neural Information Processing (NeurIPS), 2022

ICML Workshop on Spurious Correlations, Invariance, and Stability (SCIS), 2022

C10. Unsupervised Learning under Latent Label Shift

Manley Roberts*, Pranav Mani*, **Saurabh Garg**, Zachary C. Lipton
Advances in Neural Information Processing (NeurIPS), 2022
ICML Workshop on Spurious Correlations, Invariance, and Stability (SCIS), 2022

- C9. **Characterizing Datapoints via Second-Split Forgetting**
Pratyush Maini, **Saurabh Garg**, Zachary Lipton, Zico Kolter
Advances in Neural Information Processing (NeurIPS), 2022
Spotlight at ICML Workshop on Spurious Correlations, Invariance, and Stability (SCIS), 2022
- C8. **Leveraging Unlabeled Data to Predict Out-of-Distribution Performance**
Saurabh Garg, Sivaraman Balakrishnan, Zachary Lipton, Behnam Neyshabur, Hanie Sedghi [\[Paper\]](#)
International Conference on Learning Representations (ICLR), 2022
NeurIPS Workshop on Distribution Shift (DistShift), 2021
- C7. **Mixture Proportion Estimation and PU Learning: A Modern Approach**
Saurabh Garg, Yifan Wu, Alex Smola, Sivaraman Balakrishnan, Zachary Lipton [\[Paper\]](#)
Spotlight at Advances in Neural Information Processing (NeurIPS), 2021
ICML Workshop on Uncertainty & Robustness in Deep Learning (UDL), 2021
- C6. **RATT: Leveraging Unlabeled Data to obtain Generalization Guarantees**
Saurabh Garg, Zico Kolter, Sivaraman Balakrishnan, Zachary Lipton [\[Paper\]](#)
Long Talk at International Conference of Machine Learning (ICML), 2021
ICLR Workshop on Robust Machine Learning (RobustML), 2021
- C5. **On Proximal Policy Optimization's Heavy-Tailed Gradients**
Saurabh Garg, Joshua Zhanson, Emilio Parisotto, Adarsh Prasad, Zico Kolter, Zachary Lipton, Sivaraman Balakrishnan, Ruslan Salakhutdinov, Pradeep Ravikumar [\[Paper\]](#)
International Conference of Machine Learning (ICML), 2021
ICLR Workshop on Science and Engineering of Deep Learning (SEDL), 2021
- C4. **A Unified View of Label Shift Estimation**
Saurabh Garg, Yifan Wu, Sivaraman Balakrishnan, Zachary Lipton [\[Paper\]](#)
Advances in Neural Information Processing Systems (NeurIPS) 2020
Contributed Talk at ICML Workshop on Uncertainty & Robustness in Deep Learning (UDL), 2020
- C3. **Code-Switched Language models using Dual RNNs and Same-Source Pretraining**
Saurabh Garg*, Tanmay Parekh*, Preethi Jyothi [\[Paper\]](#) (* joint first authors)
Empirical Methods in Natural Language Processing (EMNLP), 2018
- C2. **Uncertainty Estimation in Segmentation with Perfect MCMC Sampling in Bayesian MRFs**
Saurabh Garg, Suyash Awate [\[Paper\]](#)
Medical Image Computing & Computer Assisted Intervention (MICCAI), 2018
- C1. **Dual Language Models for Code Mixed Speech Recognition**
Saurabh Garg, Tanmay Parekh, Preethi Jyothi [\[Paper\]](#)
Interspeech 2018 (19th Annual Conference of ISCA)

Journal.....

- J2. **Estimating Uncertainty in MRF-based Image Segmentation: An Exact-MCMC Approach**
Suyash Awate*, **Saurabh Garg***, Rohit Jena* [\[Paper\]](#) (*alphabetic ordering)
Medical Image Analysis (MedIA) Journal, 2019
- J1. **Neural Architecture for Question Answering Using a Knowledge Graph and Web Corpus**
Uma Sawant, **Saurabh Garg**, Soumen Chakrabarti, Ganesh Ramakrishnan [\[Paper\]](#)
Information Retrieval Journal, 2019
Invited Oral Talk at European Conference on Information Retrieval (ECIR), 2020

Invited Talks

Domain Adaptation under Structural Distribution Shift

· ML Theory seminar at Princeton	May' 22
· ML Seminar at IIT Bombay	July' 22
Leveraging Unlabeled Data to Predict Out-of-Distribution Performance	
· Google Brain Deep Phenomena Group	Nov '21
· Carnegie Mellon University	Nov '21
Mixture Proportion Estimation and PU Learning: A Modern Approach	
· Advances in Neural Information Processing Systems	Dec '21
· Carnegie Mellon University	Sept '21
RATT: Leveraging Unlabeled Data to obtain Generalization Guarantees	
· IIT Bombay	Oct '21
· International Conference on Machine Learning 2021	July '21
· Google Brain Deep Phenomena Group	June '21
· Carnegie Mellon University (Andrej's Reading Group)	June '21
On Proximal Policy Optimization's Heavy-Tailed Gradients	
· ICLR Workshop on Science and Engineering of Deep Learning (SEDL)	April '21
· Carnegie Mellon University (Zico's Reading Group)	June '21
Unified View of Label Shift Estimation	
· ICML Workshop on Uncertainty and Deep Learning Workshop (UDL) 2020	July '20
Neural Architecture for Question Answering using KG and Corpus	
· European Conference on Information Retrieval (ECIR) 2020	April '20
Uncertainty Estimation with Perfect MCMC Sampling	
· IIT Bombay Seminar	April '18
Code-Switched Language models	
· IIT Bombay Seminar	April '18
· Microsoft Research Labs, India	Dec '17
Approximation algorithms for weighted b-Matching	
· Purdue University	July '16

Mentorship

Bachelors in Computer Science, CMU student: Zachary Novack	2021 – ongoing
CHiLS: Zero-shot Image Classification with Hierarchical Label Sets (<i>under submission</i>)	
Understanding properties of stochastic gradient noise in deep learning (<i>Accepted at NeurIPS Workshop on Higher-order Optimization, 2022</i>)	
MS in Machine Learning, CMU, Students: Pranav Mani and Manley Roberts	2022 – ongoing
Unsupervised Learning under Latent Label Shift (<i>Accepted at NeurIPS 2022</i>)	
Ph.D. in Machine Learning, CMU, Student: Pratyush Maini	2021–2022
Characterizing Datapoints via Second-Split Forgetting (<i>Accepted at NeurIPS 2022</i>)	

Academic Service

Workshop Organizer Principles of Distribution Shift (PODS) Workshop at ICML 2022.

Reviewer. NeurIPS (2021, 2022), ICML (2021, 2022), ICLR (2022, 2023), EMNLP (2019, 2020), ACL (2020, 2021), NACL (2021), TMLR (2022).

Ph.D. Admission's Committee. Machine Learning Department, CMU, 2023-21

Teaching

Graduate Teaching Assistant, Carnegie Mellon University

- Advanced Introduction to Machine Learning, Prof. Nihar Shah Fall 2021
- Theory of Machine Learning, Prof. Pradeep Ravikumar Spring 2022

Undergraduate Teaching Assistant, IIT Bombay

- Introduction to Machine Learning, Prof. Preethi Jyothi Spring 2018
- Data Analysis and Interpretation, Prof. Suyash Awate Autumn 2017
- Computer Programming and Utilisation, Prof. Sunita Sarawagi Spring 2017
- Computer Programming and Utilisation, Prof. Benard Menezes Autumn 2016

Selected Coursework

Carnegie Mellon University: Advanced Introduction to Machine Learning (A+), Intermediate Statistics (A+), Advanced Statistical Theory 1 (A+), Convex Opt. (A+), Advanced Machine Learning Theory (A)

IIT Bombay: Web Search and Mining (AA), Organization of Web Information (AA), Optimization (AA), Artificial Intelligence (AA), Automatic Speech Recognition (AA), Linear Algebra (AA), Numerical Analysis (AA), Operating Systems (AA), Compilers (AP), Automata theory and logic (AA)