## Yash Sharma


Contact Information
ysharma1126@gmail.com
https://www.yash-sharma.com
Google Scholar
Kaggle


Research Interests
Compositional Generalization, Representation Learning, Adversarial Robustness

Education

**Max Planck Institute for Intelligent Systems (IMPRS-IS)** May 2019 - September 2024
Tübingen, Germany
PhD, Computer Science
Advised by Wieland Brendel & Matthias Bethge

**Cooper Union for the Advancement of Science and Art** September 2014 - May 2018
New York, NY, USA
*B.Eng and M.Eng, Computer Engineering*
Thesis Advisor: Sam Keene

Experience

**Entrepreneur First** October 2024 - December 2024
Paris, France
*Entrepreneur in Residence*
Unlocking the promise of computational discovery by identifying what's needed for generalization.

**Flagship Pioneering** June 2023 - August 2023
Cambridge, MA, USA
*AI Fellow*
Worked on formulating and testing promising venture hypotheses in the life sciences.

**Google Brain** February 2023 - June 2023
Mountain View, CA, USA
*Student Researcher*
Worked on predicting model performance from the training set.

**Meta AI (FAIR)** August 2022 - February 2023
New York, NY, USA
*Research Scientist Intern*
Worked with the core learning group on out-of-distribution generalization.

**Amazon (AWS AI)** October 2021 - April 2022
Tübingen, Germany
*Applied Science Intern*
Worked with the causality lab on self-supervised learning from video.

**Borealis AI** September 2018 - February 2019
Toronto, ON, Canada
*ML Researcher*
Worked on understanding the effectiveness of and robustifying models to adversarial examples.

**IBM Research** June 2017 - August 2017
Yorktown Heights, NY, USA
*Research Intern*
Worked with the AI group on generating adversarial examples in limited access settings.

## Honors And Awards

**Keynote Speaker**, MICCAI Medical Applications with Disentanglement (MAD) Workshop. 2022
**Outstanding Reviewer**, International Conference on Machine Learning (ICML). 2022
**Finalist**, NVIDIA Graduate Fellowship. 2021
**Nominee**, Google PhD Fellowship. 2021
**Reviewer Award**, International Conference on Learning Representations (ICLR). 2021
**Gold Medal** in Abstraction and Reasoning Challenge. 2020
**Full Financial Support** for doctoral studies. 2019-2024
**CAAD Overall Winner**; Prize: $**38,000**. 2018
**DEFCON 26 Presenter** on practical adversarial attacks in challenging environments. 2018
**Kaggle Competitions Master** achieved; Highest Rank: **325**. 2018
**One Gold & Two Silver Medals** in NeurIPS Competition Track. 2017
**Blockchain NYC Hackathon Winner**, IBM. 2016
**CodeSuisse Winner**, Credit Suisse. 2016
**HackRU Prize**, Rutgers University. 2015
**Half-Tuition Merit Scholarship** for undergraduate studies. 2014-2018

## Research

[1] **Pretraining Frequency Predicts Compositional Generalization of CLIP on Real-World Tasks**
*Compositional Learning, NeurIPS 2024*
Thaddäus Wiedemer*, ***Yash Sharma****, Ameya Prabhu, Wieland Brendel, Matthias Bethge (*equal contribution)

[2] **No "Zero-Shot" Without Exponential Data: Pretraining Concept Frequency Determines Multimodal Model Performance**
*Neural Information Processing Systems (NeurIPS) 2024*
Vishaal Udandarao*, Ameya Prabhu*, Adhiraj Ghosh, ***Yash Sharma***, Philip H.S. Torr, Adel Bibi, Samuel Albanie, Matthias Bethge (*equal contribution)

[3] **Attribute Diversity Determines the Systematicity Gap in VQA**
*Conference on Empirical Methods in Natural Language Processing (EMNLP) 2024*
Ian Berlot-Attwell, A. Michael Carrell, Kumar Krishna Agrawal, ***Yash Sharma***†, Naomi Saphra† (†senior author)

[4] **On Transfer of Adversarial Robustness from Pretraining to Downstream Tasks**
*Neural Information Processing Systems (NeurIPS) 2023*
Laura Fee Nern, Harsh Raj, Maurice Georgi, ***Yash Sharma***† (†senior author)
Also at *Adversarial Learning Methods for Machine Learning and Data Mining, KDD 2022*

[5] **Provably Learning Object-Centric Representations**
*International Conference on Machine Learning (ICML) 2023* (**Oral**)
Jack Brady*, Roland Zimmermann*, ***Yash Sharma***, Bernhard Schölkopf, Julius von Kügelgen, Wieland Brendel (*equal contribution)

[6] **Jacobian-based Causal Discovery with Nonlinear ICA**
*Transactions on Machine Learning Research (TMLR) 2023*
Patrik Reizinger, ***Yash Sharma***, Matthias Bethge, Bernhard Schölkopf, Ferenc Huszár, Wieland Brendel
Also at *Causal Representation Learning, UAI 2022* (**Oral**)

[7] **Pixel-level Correspondence for Self-Supervised Learning from Video**
*Pre-training: Perspectives, Pitfalls, and Paths Forward, ICML 2022*
***Yash Sharma***, Yi Zhu, Chris Russell, Thomas Brox

[8] **Disentanglement via Mechanism Sparsity Regularization: A New Principle for Nonlinear ICA**
*Causal Learning and Reasoning (CLeaR) 2022*
Sebastien Lachapelle, Pau Rodriguez Lopez, ***Yash Sharma***, Katie Everett, Remi Le Priol, Alexandre Lacoste, Simon Lacoste-Julien

[9] **Unsupervised Learning of Compositional Energy Concepts**
*Neural Information Processing Systems (NeurIPS) 2021*
Yilun Du, Shuang Li, ***Yash Sharma***, Joshua B. Tenenbaum, Igor Mordatch

[10] **Self-Supervised Learning with Data Augmentations Provably Isolates Content from Style**
*Neural Information Processing Systems (NeurIPS) 2021*
Julius von Kügelgen*, ***Yash Sharma****, Luigi Gresele*, Wieland Brendel, Bernhard Schölkopf, Michel Besserve, Francesco Locatello (*equal contribution)
Also at *Self-Supervised Learning for Reasoning and Perception, ICML 2021*

[11] **Contrastive Learning Inverts the Data Generating Process**
*International Conference on Machine Learning (ICML) 2021*
Roland Zimmermann*, ***Yash Sharma****, Steffen Schneider*, Matthias Bethge, Wieland Brendel (*equal contribution)
Also at *Self-Supervised Learning: Theory and Practice, NeurIPS 2020*

[12] **Towards Nonlinear Disentanglement in Natural Data with Temporal Sparse Coding**
*International Conference on Learning Representations (ICLR) 2021* (**Oral; 53/2997**)
David Klindt*, Lukas Schott*, ***Yash Sharma****, Ivan Ustyuzhaninov, Wieland Brendel, Matthias Bethge, Dylan Paiton (*equal contribution)

[13] **Benchmarking Unsupervised Object Representations for Video Sequences**
*Journal of Machine Learning Research (JMLR) 2021*
Marissa A. Weis, Kashyap Chitta, ***Yash Sharma***, Wieland Brendel, Matthias Bethge, Andreas Geiger, Alexander S. Ecker

[14] **Spatially Structured Recurrent Modules**
*International Conference on Learning Representations (ICLR) 2021*
Nasim Rahaman, Anirudh Goyal, Muhammad Waleed Gondal, Manuel Wuthrich, Stefan Bauer, ***Yash Sharma***, Yoshua Bengio, Bernhard Schölkopf
Also at *Inductive Biases, Invariances and Generalization in Reinforcement Learning, ICML 2020*

[15] **MMA Training: Direct Input Space Margin Maximization through Adversarial Training**
*International Conference on Learning Representations (ICLR) 2020*
Gavin Weiguang Ding, ***Yash Sharma***, Kry Yik Chau Liu, Ruitong Huang
Also at *Safe Machine Learning: Specification, Robustness, and Assurance, ICLR 2019*

[16] **On the Effectiveness of Low Frequency Perturbations**
*International Joint Conference on Artificial Intelligence (IJCAI) 2019*
***Yash Sharma***, Gavin Weiguang Ding, Marcus Brubaker

[17] **Are Generative Classifiers More Robust to Adversarial Attacks?**
*International Conference on Machine Learning (ICML) 2019*
Yingzhen Li, John Bradshaw, ***Yash Sharma***
Also at *Theoretical Foundations and Applications of Deep Generative Models, ICML 2018*

[18] **GenAttack: Practical Black-box Attacks with Gradient-Free Optimization**
*Genetic and Evolutionary Computation Conference (GECCO) 2019*
Moustafa Alzantot, ***Yash Sharma***, Supriyo Chakraborty, Huan Zhang, Cho-Jui Hsieh, Mani Srivastava

[19] **CAAD 2018: Generating Transferable Adversarial Examples**
*In arXiv:1810.01268, 2018*
***Yash Sharma***, Tien-Dung Le, Moustafa Alzantot

[20] **Generating Natural Language Adversarial Examples**
*Conference on Empirical Methods in Natural Language Processing (EMNLP) 2018*
Moustafa Alzantot*, ***Yash Sharma****, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, Kai-Wei

Chang (*equal contribution)
Also at *Security in Machine Learning, NeurIPS 2018 (Encore Track)*

[21] **Technical Report on the CleverHans v2.1.0 Adversarial Examples Library**
*In arXiv:1610.00768, 2018*
Nicolas Papernot, Fartash Faghri, Nicholas Carlini, Ian Goodfellow, Reuben Feinman, Alexey Kurakin, Cihang Xie, ***Yash Sharma*** et al.

[22] **Bypassing Feature Squeezing by Increasing Adversary Strength**
*In arXiv:1803.09868, 2018*
***Yash Sharma***, Pin-Yu Chen

[23] **Attacking the Madry Defense Model with L1-based Adversarial Examples**
*Workshop Track, ICLR 2018*
***Yash Sharma***, Pin-Yu Chen

[24] **EAD: Elastic-Net Attacks to Deep Neural Networks via Adversarial Examples**
*AAAI Conference on Artificial Intelligence (AAAI) 2018* (**Oral**)
Pin-Yu Chen*, ***Yash Sharma****, Huan Zhang, Jinfeng Yi, Cho-Jui Hsieh (*equal contribution)

[25] **ZOO: Zeroth Order Optimization based Black-box Attacks to Deep Neural Networks without Training Substitute Models**
*ACM Workshop on Artificial Intelligence and Security (AISec) 2017*
**Best Paper Award Finalist**
Pin-Yu Chen*, Huan Zhang*, ***Yash Sharma***, Jinfeng Yi, Cho-Jui Hsieh (*equal contribution)