

Arpit Bansal

COLLEGE PARK MD, USA

✉ bansal01@umd.edu | 🏠 arpitbansal297.github.io | 📄 arpitbansal297 | 🌐 Arpit Bansal

Research Interests

My research primarily revolves around the empirical analysis and optimization of Deep Neural Network architectures with the aim of improving their training and inference mechanisms. I am particularly interested in a range of multi-modal vision tasks, including text-to-image generation, image editing, and image reasoning (VLMs). Additionally, I have investigated algorithmic synthesis via neural networks to enable logical reasoning and developed robust watermarking techniques to maintain the security and integrity of neural network models.

Education

University of Maryland, College Park

PHD IN COMPUTER SCIENCE

- Advisor: Prof. Tom Goldstein
- Dean's Fellowship

College Park
Jan 2021 - Jan 2025

Indian Institute of Technology, Kharagpur

BACHELORS + MASTERS IN ELECTRICAL ENGINEERING

- Minor in Computer Science
- Advisor: Prof. Rajiv Ranjan Sahay

Kharagpur, India
Aug 2014 - May 2019

Industry Experience

March 2025
- Present

Google AI Innovation & Research

Research Scientist

As a Research Scientist, I advance image-to-image generation by developing vision encoders useful for generative tasks and improving algorithms for RLHF

May 2024 -
Dec 2024

Meta GenAI @ Llama

AI Research Scientist Intern

In collaboration with the Llama team, developing multi-modal Generative Models with a focus on image reasoning and image editing thus working on VLMs and autoregressive image generation.

June 2023 -
Aug 2023

Amazon AWS AI labs

Applied Science Intern

Examined the challenges of existing Text-to-Image Diffusion models' low fidelity and devised a novel approach to produce high-fidelity images conditioned on text.

July 2019 -
Dec 2020

Visa Incorporated, India

Software Developer

Designed an NLP engine for intuitive English queries and an impactful Recommendation System, earning a Visa Trade Secret. Introduced a Merchant Rating system based on transaction nature, recognized with another Trade Secret for its Research and Development.

Selected Publications

Universal guidance for diffusion models

A. Bansal*, H. Chu*, A. Schwarzschild, S. Sengupta, M. Goldblum, J. Geiping, T. Goldstein
International Conference on Learning Representations (ICLR) 2024

Cold diffusion: Inverting arbitrary image transforms without noise

A. Bansal, E. Borgnia, H. Chu, J Li, H. Kazemi, F. Huang, M. Goldblum, J. Geiping, T. Goldstein
Conference on Neural Information Processing Systems (Neurips) 2023

Transformers Can Do Arithmetic with the Right Embeddings

S. McLeish*, **A. Bansal***, A. Stein, N. Jain, J. Kirchenbauer, B. Kailkhura, A. Bhatele, J. Geiping, A. Schwarzschild, T. Goldstein
Conference on Neural Information Processing Systems (Neurips) 2024

End-to-end Algorithm Synthesis with Recurrent Networks: Logical Extrapolation Without Overthinking

A. Bansal*, A. Schwarzschild*, E. Borgnia, Z. Emam, F. Huang, M. Goldblum, T. Goldstein
Conference on Neural Information Processing Systems (Neurips) 2022

Certified Neural Network Watermarks with Randomized Smoothing

A. Bansal, P. Yeh Chiang, M. Curry, R. Jain, C. Wightington, V. Manjunatha, J. P Dickerson, T. Goldstein
International Conference on Machine Learning (ICML - Spotlight) 2022

Can You Learn the Same Model Twice? Investigating Reproducibility and Double Descent from the Decision Boundary Perspective

G. Somepalli, L. Fowl, **A. Bansal**, P. Yeh Chiang, Y. Dar, R. Baraniuk, M. GoldBlum, T. Goldstein
Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR - Oral) 2022

Transfer Learning with Deep Tabular Models

R. Levin*, V. Cherepanova*, A. Schwarzschild[†], **A. Bansal[†]**, C Bayan Bruss, T. Goldstein, A. G. Wilson, M. Goldblum
International Conference on Learning Representations (ICLR) 2023

Canary in a Coalmine: Better Membership Inference with Ensembled Adversarial Queries

Y. Wen, **A. Bansal**, H. Kazemi, E. Borgnia, M. Goldblum, J. Geiping, T. Goldstein
International Conference on Learning Representations (ICLR - Spotlight) 2023

Loss Landscapes are All You Need: Neural Network Generalization Can Be Explained Without the Implicit Bias of Gradient Descent

P. Yeh Chiang, R. Ni, D. Yu Miller, **A. Bansal**, J. Geiping, M. Goldblum, T. Goldstein
International Conference on Learning Representations (ICLR - Spotlight) 2023

Teaching Experience

Fall 2021 **Control Systems**, Teaching Assistant
Spring 2021 **Operating Systems**, Teaching Assistant

University of Maryland, College Park
University of Maryland, College Park

Relevant Course-work

Machine Learning Information Retrieval (IIT), Machine Learning (IIT), Speech and NLP (IIT), Deep Learning (UMD), Algorithms in Machine Learning (UMD)

Signal Processing Digital Signal Processing (IIT), Statistical Signal Processing (IIT), Probability and Stochastic Processes (IIT), Random Processes (UMD), Information Theory (UMD), Numerical Analysis I (UMD), Advanced Numerical Optimization (UMD)