Abhishek Sinha

abhishek.sinha94@gmail.com | 650.709.5501 | a7b23.github.io | Google Scholar | LinkedIn

INTERESTS

MultiModal Models, Generative Models, Representation Learning, Few-shot Learning

EDUCATION

Stanford University
MS in Computer Science
Sept. 2019 - June 2021
Cum. GPA: 4.13/4.0

IIT Kharagpur
BTech in E & ECE
2013-2017 | Kharagpur, India
Minor in Computer Science
Cum. GPA: 9.63 / 10.0
Minor Cum. GPA: 9.8 / 10.0

COURSEWORK

CS 221, CS 231N, CS 236, CS 234, CS 224N, CS 271

ACHIEVEMENTS

Cheetah Award
Won the Cheetah Award for my work on latency and compute optimization at Waymo.

Young Engineer Award
Won the Outstanding Young Engineers Award at Adobe Inc.

Adobe MAX
My work on image synthesis was showcased on stage at Adobe MAX.

Winner OF AI HACKATHON
Winner of the Microsoft AI Hackathon competition held at IIT Kharagpur.

SKILLS

Python • C • C++
TensorFlow • PyTorch • Caffe
OpenCV • Scikit-learn • Numpy

POSITIONS

Reviewer for NeurIPS 2021,2022, ICLR 2022,2023
Course Assistant, CS 330

EXPERIENCE

Google DeepMind | Software Engineer
Feb 2024 - Present | Sunnyvale, USA
- Post-training co-lead and pre-training core-contributor for nano banana pro.
- Core contributor for the nano banana and Imagen 4 launch. I led the integration of nano banana with Gemini and worked on post-training of it.
- Before image generation, I was a core contributor for Gemini image understanding post-training.

Waymo LLC | Software Engineer
June 2021 - Feb 2024 | Mountain View, USA
- Led the fine-tuning of LVLMs for driving domain and achieved large data efficiency improvements for several Perception models.

Stanford Univ. | Research Assistant Under Stefano Ermon
January 2020 – June 2021 | Stanford, USA
- Developed the first latent diffusion model and worked on improving generative models and representation learning models by designing novel loss functions.

Waymo LLC | Perception Research and Development Intern
June 2020 – September 2020 | Mountain View, USA
- Showed the efficacy of different active learning algorithms for 3D detection of vehicles and pedestrians over the Waymo Open Dataset.

Adobe | Software Engineer
June 2017 – August 2019 | Noida, India
- Developed a deep learning-based visual search product for apparel.

SELECTED PUBLICATIONS

Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities | Core Contributor | Paper
- Worked on native image generation integration with Gemini.

Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context | Core Contributor | Paper
- Worked on Multi-Modal Post-training of Gemini.

Comparing Distributions by Measuring Differences that Affect Decision Making | Best Paper Award at ICLR, 2022 | Joint First Author | Paper
- Proposed a new divergence metric using H-entropy computed from log-likelihood of generative models.
- Our approach outperformed the FID metric for evaluating image quality.

D2C: Diffusion-Denoising Models for Few-shot Conditional Generation | NeurIPS, 2021 | Joint First Author | Paper | Project
- Improved the representation learning and generation abilities of VAE via contrastive loss and strong diffusion prior respectively.
- Our diffusion model was SOTA for few-shot conditional generation.

Negative Data Augmentation | Joint First Author | ICLR, 2020 | Paper
- Proposed a new training objective for GAN and contrastive learning approaches using negative data augmentation.

Introspection: Accelerating Neural Network Training By Learning Weight Evolution | First Author | ICLR, 2017 | Paper
- Developed an algorithm to speed up training of deep neural networks.