

Anand Bhattad

PIXELS · PERCEPTION · PHYSICS

Assistant Professor of Computer Science · DSAI & Institute of Robotics
Johns Hopkins University

✉ bhattad@jhu.edu 🌐 anandbhattad.github.io 📍 Baltimore, MD

— Research

Generative models look right without being right. My 3P Vision group builds the tools to tell the difference — to read out what they know and don't, and how to improve them — along three Ps:

Pixels — controllable generative models for relighting, compositing, and editing.

Perception — reading out the scene structure models implicitly know: intrinsics, geometry, and parts.

Physics — testing and enforcing physical and projective consistency, and carrying it to robots.

The group looks beyond generative models too: vision architectures, 3D perception, multisensory perception for embodied AI, and generative models for computational mechanics.

Keywords: Visual Generative Models; 3D Vision; Vision-based Robotics; Neural Rendering; Inverse Graphics; Intrinsic Images; Image-based Lighting; Physics Awareness in Large Vision Models.

— Education

	University of Illinois Urbana-Champaign (UIUC), USA	
	Ph.D., Computer Science	2019 – 2024
	Thesis: Exploring Knowledge in Generative Models	
	Thesis committee: David Forsyth (Advisor) , Alexei A. Efros, William T. Freeman, Derek Hoiem, Svetlana Lazebnik, Shenlong Wang	
	M.S., Computer Science	2017 – 2018
	M.S., Civil Engineering	2015 – 2017
	National Institute of Technology Karnataka (NITK) Surathkal, India	
	B.Tech., Civil Engineering	2011 – 2015

— Work Experience

	Johns Hopkins University , Assistant Professor of Computer Science	2025 – present
	Toyota Technological Institute at Chicago , Research Assistant Professor	2023 – 2025
	University of California Berkeley , Visiting Scholar; Host: Alexei A. Efros	2024
	Allen Institute for AI (AI2) , Research Intern with Aniruddha Kembhavi	2023
	Intel , Research Intern with Stephan R. Richter in Vladlen Koltun's team	2021
	NVIDIA , Research Intern with Andrew Tao and Bryan Catanzaro	2020
	Fyusion Inc , Research Intern with Abhishek Kar and Rodrigo Ortiz Cayon	2019
	Siemens Corporate Technology , Computer Vision & DL Intern with Jan Ernst	2017
	Michigan State University , Research Intern with Nizar Lajnef	2014
	IIT Hyderabad , Research Intern with Ramacharla Pradeep Kumar	2013

Selected Honors and Awards

Best Paper Award in Robot Learning, ICRA	2026
Outstanding Reviewer, ICCV	2023
Best Paper Finalist, CVPR	2022
Outstanding Emergency Reviewer, CVPR	2021
Excellent Teaching Assistant, Intro Computing: Engrg & Sci	2016

Selected Publications

Self-selected publications (not necessarily the most cited work).

- Generative Cinematographer: Composing Camera and Object Motion in 3D; [on arXiv](#)
- PowerSim: Differentiable Physics Simulation and Rendering with Power Diagrams; [on arXiv](#)
- Shadows Don't Lie and Lines Can't Bend! Generative Models Don't Know Projective Geometry...For Now; [CVPR 2024](#)
- StyleGAN Knows Normals, Depth, Albedo, and More; [NeurIPS 2023](#)
- PRINCIPIA: Relational Physics Tests for Video Models; [NeurIPS 2026 \(Evaluations & Datasets Track\)](#)
- Visual Jenga: Discovering Object Dependencies via Counterfactual Inpainting; [NeurIPS 2025](#)
- Latent Intrinsic Emerge from Training to Relight; [NeurIPS 2024 \(Spotlight\)](#)
- Generative Blocks World: Moving Things Around in Pictures; [ICLR 2026](#)
- Do You Know Where Your Camera Is? View-Invariant Policy Learning with Camera Conditioning; [ICRA 2026 \(Best Paper Award in Robot Learning\)](#)
- LumiNet: Latent Intrinsic Meets Diffusion Models for Indoor Scene Relighting; [CVPR 2025](#)

Papers Under Review


* denotes equal contribution, and † indicates equal advising



4. Generative Cinematographer: Composing Camera and Object Motion in 3D

Jiahao Zhang, Chaohao Yang, Namitha Guruprasad, Vivekkyoti Banerjee, Trong-Tung Nguyen, Alan Yuille, **Anand Bhattad**
[on arXiv](#)


[Paper](#) [Project](#)



3. PowerSim: Differentiable Physics Simulation and Rendering with Power Diagrams

Trong-Tung Nguyen, **Anand Bhattad**
[on arXiv](#)


[Paper](#) [Project](#) [Code](#)



2. FracGen: Learning How Objects Stretch and Tear with Physics-Informed Video Generation

Trong-Tung Nguyen, Jiahao Zhang, **Anand Bhattad**
[on arXiv](#)

[Paper](#) [Project](#)




1. ViT-5: Vision Transformers for The Mid-2020s

Feng Wang, Sucheng Ren, Tiezheng Zhang, Predrag Neskovic, **Anand Bhattad**, Cihang Xie, Alan Yuille
on arXiv

Paper Code

Refereed Conference Papers



32. PRINCIPIA: Relational Physics Tests for Video Models

Varun Varma Thozhiyoor*, Shivam Tripathi*, Venkatesh Babu Radhakrishnan, **Anand Bhattad**

NeurIPS 2026 (Evaluations & Datasets Track)

Paper Project Code



31. Thinking in Boxes: 3D Editing in Real Images Made Easy

Pradhaan S Bhat*, Naveen Chandra*, Rishubh Parihar, Vaibhav Vavilala, Venkatesh Babu, D.A. Forsyth, **Anand Bhattad**

NeurIPS 2026

Paper Project Code




30. SyncLight: Single-Edit Multi-View Relighting

David Serrano-Lozano, **Anand Bhattad**[†], Luis Herranz, Jean-François Lalonde[†], Javier Vazquez-Corral[†]

NeurIPS 2026

Paper Project Code




29. RATS! Patches Talk Through Registers: Emergent Parts in Register Attention Transformers

Timing Yang, Predrag Neskovic, Jansen Seheult, Wenchao Han, **Anand Bhattad**, Alan Yuille, Feng Wang

NeurIPS 2026

Paper Code



28. Context-Aware Autoregressive Image Generation for Emerging Reasoning Properties

Jixuan Ying, Haoyu Liu, Timing Yang, Tingyu Zhu, Predrag Neskovic, **Anand Bhattad**, Cihang Xie, Zeyu Zheng, Alan Yuille, Feng Wang

NeurIPS 2026



27. SyncFix: Fixing 3D Reconstructions via Multi-View Synchronization

Deming Li, Abhay Yadav, Cheng Peng, Rama Chellappa, **Anand Bhattad**

ECCV 2026

Paper Project Code




26. Relighting as a Probe of Visual Priors via Augmented Latent Intrinsic

Xiaoyan Xing, Xiao Zhang, Sezer Karaoglu, Theo Gevers, **Anand Bhattad**

ICML 2026

Paper Project Code



25. **Objects in Generated Videos Are Slower Than They Appear: Models Suffer Sub-Earth Gravity and Don't Know Galileo's Principle...for now**
Varun Varma Thozhiyoor, Shivam Tripathi, Venkatesh Babu Radhakrishnan, **Anand Bhattad**

CVPR 2026 (Findings)

Paper Project Code



24. **Name That Part: 3D Part Segmentation and Naming**

Soumava Paul, Prakhar Kaushik, Ankit Vaidya, **Anand Bhattad**, Alan Yuille

CVPR 2026 (Findings)


Paper Project Code



23. **Generalizable Sparse-View 3D Reconstruction from Unconstrained Images**
Vinayak Gupta, Chih-Hao Lin, Shenlong Wang, **Anand Bhattad**, Jia-Bin Huang

CVPR 2026

Paper Project Code




22. **Do You Know Where Your Camera Is? View-Invariant Policy Learning with Camera Conditioning**

Tianchong Jiang, Jingtian Ji, Xiangshan Tan, Jiading Fang, **Anand Bhattad**, Vitor Guizilini, Matthew R. Walter

ICRA 2026 (Oral; **Best Paper Award in Robot Learning**)

Paper Project Code



21. **Generative Blocks World: Moving Things Around in Pictures**

Vaibhav Vavilala, Seemandhar Jain, Rahul Vasanth, D.A. Forsyth, **Anand Bhattad**

ICLR 2026

Paper Project Code




20. **SPOTLIGHT: Shadow-Guided Object Relighting via Diffusion**

Frédéric Fortier-Chouinard, Zitian Zhang, Louis-Etienne Messier, Mathieu Garon, **Anand Bhattad**, Jean-François Lalonde

3DV 2026

Paper Project Code



19. **Improved Convex Decomposition with Ensembling and Boolean Primitives**

Vaibhav Vavilala, Florian Kluger, Seemandhar Jain, Bodo Rosenhahn, **Anand Bhattad**, D.A. Forsyth

3DV 2026

Paper




18. **VISUAL JENGA: Discovering Object Dependencies via Counterfactual Inpainting**

Anand Bhattad, Konpat Preechakul, Alexei A. Efros

NeurIPS 2025

Paper Project



17. **LUMINET: Latent Intrinsic Meets Diffusion Models for Indoor Scene Relighting**

Xiaoyan Xing, Konrad Groh, Sezer Karaoglu, Theo Gevers, **Anand Bhattad**


CVPR 2025

Paper Project Code




16. **ScribbleLight: Single Image Indoor Relighting with Scribbles**
Jun Myeong Choi, Annie N. Wang, Pieter Peers, **Anand Bhattad**, Roni Sengupta
CVPR 2025

Paper Project Code



15. **UrbanIR: Large-Scale Urban Scene Inverse Rendering from a Single Video**
Zhi-Hao Lin, Bohan Liu, Yi-Ting Chen, Kuan-Sheng Chen, D.A. Forsyth, Jia-Bin Huang, **Anand Bhattad**, Shenlong Wang
3DV 2025

Paper Project Code




14. **ZEROCOMP: Zero-shot Object Compositing from Image Intrinsic via Diffusion**
Zitian Zhang, Frédéric Fortier-Chouinard, Mathieu Garon, **Anand Bhattad**, Jean-François Lalonde
WACV 2025 (Oral presentation)

Paper Project Code



13. **Latent Intrinsic Emerge from Training to Relight**
Xiao Zhang, William Gao, Seemantdhhar Jain, Michael Maire, D.A. Forsyth, **Anand Bhattad**
NeurIPS 2024 (Spotlight presentation)

Paper Project Code



12. **From an Image to a Scene: Learning to Imagine the World from a Million 360° Videos**
Matthew Wallingford, **Anand Bhattad**, Aditya Kusupati, Vivek Ramanujan, Matt Deitke, Sham Kakade, Aniruddha Kembhavi, Roozbeh Mottaghi, Wei-Chiu Ma, Ali Farhadi
NeurIPS 2024

Paper Project Code



11. **Videoshop: Localized Semantic Video Editing with Noise-Extrapolated Diffusion Inversion**
Xiang Fan, **Anand Bhattad**[†], Ranjay Krishna[†]
ECCV 2024

Paper Project Code




10. **Shadows Don't Lie and Lines Can't Bend! Generative Models Don't Know Projective Geometry...For Now**
Ayush Sarkar*, Hanlin Mai*, Amitabh Mahapatra*, Svetlana Lazebnik, D.A. Forsyth, **Anand Bhattad**
CVPR 2024

Paper Project Code




9. **StyLitGAN: Image-based Relighting via Latent Control**
Anand Bhattad, James Soole, D.A. Forsyth
CVPR 2024

Paper Project Code



8. **StyleGAN Knows Normals, Depth, Albedo, and More**
Anand Bhattad, Daniel McKee, Derek Hoiem, D.A. Forsyth
NeurIPS 2023

Paper




7. OBJect 3DIT: Language-guided 3D-aware Image Editing

Oscar Michel, **Anand Bhattad**, Eli Vanderbilt, Ranjay Krishna, Aniruddha Kembhavi, Tanmay Gupta

NeurIPS 2023

Paper Project Code



6. Improving Equivariance in State-of-the-Art Supervised Depth and Normal Predictors

Yuanyi Zhong, **Anand Bhattad**, Yuxiong Wang, D.A. Forsyth

ICCV 2023

Paper Code



5. Cut-and-Paste Object Insertion by Enabling Deep Image Prior for Reshading

Anand Bhattad, D.A. Forsyth

3DV 2022 (**Spotlight presentation**)

Paper Project Code




4. DiVeR: Real-time and Accurate Neural Radiance Fields with Deterministic Integration for Volume Rendering

Liwen Wu, Jae Yong Lee, **Anand Bhattad**, Yuxiong Wang, D.A. Forsyth

CVPR 2022 (**Best Paper Finalist; Oral presentation**)

Paper Project Code



3. View Generalization for Single Image Textured 3D Models

Anand Bhattad, Aysegul Dundar, Guilin Liu, Andrew Tao, Bryan Catanzaro

CVPR 2021

Paper Project




2. Unrestricted Adversarial Perturbations via Semantic Manipulation

Anand Bhattad*, Min-Jin Chong*, Kaizhao Liang, Bo Li, D.A. Forsyth

ICLR 2020

Paper Code




1. Improved Style Transfer with Calibrated Metrics

Mao-Chuang Yeh*, Shuai Tang*, **Anand Bhattad**, Chuhan Zou, D.A. Forsyth

WACV 2020

Paper

Refereed Workshop Papers




5. Intrinsic LoRA: A Generalist Approach for Discovering Knowledge in Generative Models

Xiaodan Du, Nicholas Kolkin, Greg Shakhnarovich, **Anand Bhattad**

CVPR-W 2024 (**Oral presentation**)

workshops: Generative Models for CV, AI for 3D Gen, SyntaGen, SynthData4CV, Dataset Distillation

Paper Project Code



4. **UrbanIR: Large-Scale Urban Scene Inverse Rendering from a Single Video**
Zhi-Hao Lin, Bohan Liu, Yi-Ting. Chen, D.A. Forsyth, Jia-Bin Huang, **Anand Bhattad**, Shenlong. Wang

CVPR-W 2024

workshop: SynthData4CV

Paper Project Code


3. **MIMIC: Masked Image Modeling with Image Correspondences**

Kalyani Marathe, Mahtab Bigverdi, Nishat Khan, Tuhin Kundu, Patrick Howe, Sharan Ranjit S, **Anand Bhattad**, Aniruddha Kembhavi, Linda G. Shapiro, Ranjay Krishna

CVPR-W 2024

workshop: 3D with Multi-View Supervision (Archival long paper)

Paper Code



2. **Big but Imperceptible Adversarial Perturbations via Semantic Manipulation**

Anand Bhattad*, Min-Jin Chong*, Kaizhao Liang, Bo Li, D.A. Forsyth

CVPR-W 2019 (**Oral presentation**)

workshop: Adversarial Machine Learning in Real-World Computer Vision Systems

Paper Code



1. **Detecting Anomalous Faces with “No Peeking” Autoencoders**


Anand Bhattad, Jason Rock, D.A. Forsyth

CVPR-W 2018 (**Oral presentation**)

workshop: Vision with Biased and Scarce Data

Paper

— Technical Reports




4. **Generative Models: What Do They Know? Do They Know Things? Let's Find Out!**

Xiaodan Du, Nicholas Kolkin, Greg Shakhnarovich, **Anand Bhattad**

Technical Report, arXiv 2024

Paper Project Code




3. **Blocks2World: Controlling Realistic Scenes with Editable Primitives**

Vaibhav Vavilala, Seemadhar Jain, Rahul Vasanth, **Anand Bhattad**, D.A. Forsyth

Technical Report, arXiv 2023

Paper



2. **Make It So: Steering StyleGAN for Any Image Inversion and Editing**

Anand Bhattad, Viraj Shah, Derek Hoiem, D.A. Forsyth

Technical Report, arXiv 2023

Paper Project

1. **SIRfyN: Single Image Relighting from your Neighbors**

D.A. Forsyth, **Anand Bhattad**, Pranav Asthana, Yuanyi Zhong, Yuxiong Wang

Technical Report, arXiv 2021

Paper

— Patents

1. Intrinsic-ControlNet: Zero-shot Object Compositing from Image Intrinsic

Zitian Zhang, Frédéric Fortier-Chouinard, Mathieu Garon, **Anand Bhattad**, Jean-François Lalonde
Pending

— Advising

Postdoctoral Researchers

- **Jamila Taaki** (JHU, primary advisor Laurent Pueyo) 2026 – now

Doctoral Students

- **Deming (Remus) Li** (JHU, co-advised with Rama Chellappa) 2026 – now
- **Jiahao Zhang** (JHU) 2026 – now
- **Trong-Tung Nguyen** (JHU) 2026 – now

Master's Students

- **Chaohao Yang** (JHU) 2026 – now
- **Joanna Cheng** (JHU, co-advised with Homanga Bharadhwaj) 2026 – now
- **Jiafeng Gu** (JHU) 2026 – now
- **Rayhaneh Behraves** (JHU) 2026 – now
- **Yixin (Tracy) Zhu** (UChicago); Project on Improving Projective Geometry 2024 – 2025
- **Feiran Wang** (UIUC); Project on NeRF + Latent Illumination 2023
- **Kexuan (Klaus) Zou** (UIUC | Now: Software Engineer at NVIDIA); Project on 2D Meshes 2019 – 2020

Undergraduate Students

- **Jin Hong (Jayden) Moon** (JHU) 2026 – now
- **Jiuzhou (Anthony) Cheng** (JHU) 2026 – now
- **Anudeep Chimata** (BITS Goa, India); Project on Representation Learning & Generative Models 2025 – now
- **Joshua Ahn** (UChicago); Project on Wavelets in Neural Radiance Fields 2024
- **Zhiyan (Alex) Wang** (UChicago); Project on Improving Projective Geometry 2024 – 2025
- **Kuan-Sheng Chen** (UIUC); Project on Improving Generative Models with Intrinsic Images 2023
- **Brian Chen** (UIUC | REU at CMU's RI | Next: Hive AI); Project on Lighting Correction 2020 – 2022
- **Anchu Zhu** (UIUC | Next: MS in CS at USC); Project on Anomaly Detection 2018 – 2019

— Student Collaborators

Students whom I actively co-mentor or collaborate with (home institution in parentheses).

Doctoral-Level

- **Feng Wang** (JHU); Papers in NeurIPS 2026 & under review 2026 – now
- **Namitha Guruprasad** (JHU) 2026 – now
- **Timing Yang** (JHU); Papers in NeurIPS 2026 2026 – now
- **Vivekjiyoti Banerjee** (JHU); Paper under Review 2026 – now
- **Zhaoliang Zhang** (JHU) 2026 – now
- **Ozgur Kara** (UIUC) 2026 – now
- **David Serrano-Lozano** (UAB, Spain); Paper in NeurIPS 2026 2025 – now
- **Soumava Paul** (JHU); Paper in CVPR 2026 2025 – now
- **Jun Myeong Choi** (UNC Chapel Hill); Paper in CVPR 2025 2024 – 2025
- **William Gao** (UChicago); Paper in NeurIPS 2024 2024 – now
- **Xiaoyan Xing** (University of Amsterdam); Papers in CVPR 2025 & ICML 2026 2024 – now

• Ayush Sarkar (UIUC); Paper in CVPR 2024	2023 – 2024
• Kalyani Marathe (University of Washington); Paper in CVPR-W 2024	2023 – 2024
• Matthew Wallingford (University of Washington); Paper in NeurIPS 2024	2023 – 2024
• Vaibhav Vavilala (UIUC); Papers in ICLR 2026, 3DV 2026 & NeurIPS 2026	2023 – 2025
• Xiang Fan (University of Washington); Paper in ECCV 2024	2023 – 2024
• Xiao Zhang (UChicago); Papers in NeurIPS 2024 & ICML 2026	2023 – 2025
• Xiaodan Du (TTIC); Paper in CVPR-W 2024	2023 – 2024
• Zhi-Hao Lin (UIUC); Paper in CVPR-W 2024 & Paper in 3DV 2025	2023 – 2024
• Zitian Zhang (Université Laval); Papers in WACV 2025 & 3DV 2026	2023 – 2024

Master's-Level

• Naveen Chandra R (IISc, India); Paper in NeurIPS 2026	2026 – now
• Shivam Tripathi (IISc, India); Papers in CVPR 2026 & NeurIPS 2026	2025 – now
• Varun Varma Thozhiyoor (IISc, India); Papers in CVPR 2026 & NeurIPS 2026	2025 – now
• Frédéric Fortier-Chouinard (Université Laval); Papers in WACV 2025 & 3DV 2026	2023 – 2025
• Hanlin (Asher) Mai (UIUC); Paper in CVPR 2024	2023 – 2025
• James Soole (UIUC Now: Research Engineer at MATLAB); Paper in CVPR 2024	2023
• Seemantihar Jain (UIUC Now: PhD at UCSD); Papers in ICLR 2026, 3DV 2026 & NeurIPS 2024	2023 – 2024
• Pranav Asthana (UIUC Now: PhD at UMD); Project on Relighting from Neighbors	2021 – 2022

Undergraduate-Level

• Furkan Horoz (METU, Turkey / UIUC)	2026 – now
• Pradhaan S Bhat (PES University / IISc, India); Paper in NeurIPS 2026	2026 – now
• Amitabh Mahapatra (UIUC Now: PhD at UIUC); Paper in CVPR 2024	2023 – 2025
• Oscar Michel (AI2 Now: PhD at NYU); Paper in NeurIPS 2023	2023
• Liwen Wu (UIUC Now: PhD at UCSD); Paper in CVPR 2022 (Best Paper Finalist)	2022

Doctoral Thesis Committees

• William Gao, “Neural Methods for 3D Mesh Editing” (UChicago)	2025
• Xiao Zhang, “Representation Learning from and for Generative Models” (UChicago)	2025
• Xin Yuan, “Interpretable Unsupervised Generative Learning via Factorized Architectures and Structured Bottlenecks” (UChicago)	2024

— Courses Designed

Generative Vision: The Art and Science of Visual Synthesis, Johns Hopkins University; *Course Instructor* Fall 2026

- Developed a new course on visual synthesis with generative models

Learning-based Vision, Johns Hopkins University; *Course Instructor* Spring 2026

- Developed a new course on modern approaches in computer vision
- Taught a class of 35 students from Johns Hopkins University

Past Meets Present: A Tale of Two Visions, TTIC; *Course Instructor* Spring 2024

- Developed a new course from scratch on Computer Vision, combining historical and modern approaches
- Taught a class of 20+ students from TTIC and UChicago

— Teaching Assistant

Computer Science, UIUC; *Graduate Teaching Assistant*

2016 – 2018

- Applied Machine Learning (CS498 AML), Fall 2018
- Data Structures (CS 225), Spring 2017
- Intro Computing: Engrg & Sci (CS 101), Spring 2016 & Fall 2016

— Invited Talks

Principia: What Video Models Know about High-School Physics

- CVPR AC Workshop (Lightning Talk) Jun 2026

The Three P's of Modern Computer Vision: Pixels, Perception, and Physics

- Indian Institute of Science, Bangalore, Kotak-IISC AI/ML Seminar Jan 2026
- Johns Hopkins University, Guest Lecture in Rama Chellappa's Machine Perception Course Nov 2026
- Johns Hopkins University, Physics of Learning Seminar, invited by Brice Ménard Oct 2026
- Johns Hopkins University, ECE Department Seminar Sept 2025

Emergent Latent Intrinsic Representations for Scene Relighting

- UChicago, Great Lakes Graphics Workshop Apr 2025

What Generative Image Models Understand (and Don't) about the Physical World

- Univ of Wisconsin at Madison Apr 2025
- Johns Hopkins University, CS/LCSR Seminar Mar 2025
- UC Irvine, CS Seminar Mar 2025
- Stony Brook University, CS Seminar Mar 2025
- UIUC: External Speaker Series Feb 2025
- ULaval: Vision Seminar Feb 2025
- UCSD: Pixel Cafe Seminar Jan 2025

Are Generative Image Models Physically Grounded?

- UPenn: GRASP Seminar Dec 2024
- NYU: Guest Lecture in Saining Xie's Computer Vision Course Dec 2024

Generative Models Inside Out

- Midwest Computer Vision Workshop, Indiana University Sep 2024

What Do Generative Image Models Know?

- IIIT Hyderabad, India; Vision Seminar Jan 2024
- Exactech, Inc.; Tech Talk Oct 2023

Exploring Knowledge in Generative Models

- Stanford University in Jiajun Wu's group Jun 2023
- University of Tübingen, Autonomous Vision Group May 2023

What Do Generative Models Know about Geometry and Illumination?

- UC Berkeley: Vision Seminar Apr 2023
- NVIDIA Research Apr 2023
- MIT: Vision and Graphics Seminar Apr 2023
- CMU: VASC Seminar Mar 2023
- UW: RAIVN Vision Seminar Mar 2023

Towards a Productive and Fun PhD Experience

- UIUC: Computer Vision Workshop Apr 2023
- Learning about Light Without Labeled Data**
- UMD: Vision Seminar Mar 2023
 - UCSD: Pixel Cafe Seminar Feb 2023
 - TTIC: Research Talk Feb 2023
- Drag-and-Drop Rendering: Towards In the Wild Image Editing**
- CMU: Misc-Read Vision Reading Group Aug 2021

— Services & Leadership

Reviewing

- **Area Chair**, ICLR 2027, CVPR 2027, NeurIPS 2026, CVPR 2026, ICLR 2026, 3DV 2026, NeurIPS 2024 – now 2025, ICCV 2025, CVPR 2025, WACV 2025
- **Reviewer**, CVPR, NeurIPS, ECCV, ICCV, ICLR, SIGGRAPH, SIGGRAPH Asia and PAMI 2018 – now

University Service

- **Faculty Hiring Committee**, Department of Computer Science, Johns Hopkins University 2025 – 2026

Community Building Workshops

- **Lead Organizer**, “Bitter Lessons” workshop at CVPR Jun 2026
- **Lead Organizer**, “How to Stand Out in the Crowd?” workshop at CVPR Jun 2025
- **Lead Organizer**, “CV 20/20: A Retrospective Vision” workshop at CVPR Jun 2024
- **Lead Organizer**, “Scholars & Big Models: How Can Academics Adapt?” workshop at CVPR Jun 2023

Technical Workshops

- **Organizer**, Findings, a new track at CVPR Jun 2026
- **Organizer**, “Humans of Generative Models (HuGs)” at CVPR Jun 2026
- **Lead Organizer**, “UniLight Workshop” at ICCV Oct 2025
- **Co-Organizer**, “Ind3D: Enforcing geometric, physical, topological, and functional inductive bias in 3D generation” workshop at CVPR Jun 2025
- **Lead Organizer**, “Knowledge in Generative Models” workshop at ECCV Sep 2024
- **Co-lead Organizer**, “Multimodal Artificial Intelligence ” workshop at TTIC Aug 2024
- **Lead Organizer**, Computer Vision workshop at Allerton, UIUC Apr 2023
- **Committee**, “Adversarial ML in Real-World Computer Vision Systems”, workshop at CVPR Jun 2019
- **Committee**, “Security and Privacy in Machine Learning” workshop at ICML Jun 2019

Seminars and Reading Groups

- **Organizer**, UIUC Vision Seminar, a new weekly seminars inviting speakers outside UIUC 2021 – 2023
- **Organizer**, Vision Lunch (computer vision reading group) 2021 – 2023
- **Organizer**, UIUC Summer Vision Coffee 2018

Co-Founder, SchoolEngg

- Forum for high-school students; providing firsthand knowledge of all engineering disciplines 2013 – 2015
- Renamed as *PrepLift*, an education counselling start-up in India

Co-Founder, American Society of Civil Engineers (ASCE) NITK Chapter

- Co-Founded the *first* ASCE student chapter of India 2013 – 2015
- Served as treasurer for one year and as a senior advisor for the second year
- Coordinated with board members, faculty, and outside organizations to facilitate 10+ events