PRITHVIJIT CHATTOPADHYAY

prithvijitchattopadhyay@gmail.com www.prithv1.xyz

RESEARCH INTERESTS

World Models, Synthetic Data, Robust & Reliable Machine Learning, Sim2Real Transfer, Embodied AI

EDUCATION

Ph.D. in Computer Science, School of Interactive Computing, Georgia Tech 2019-2024
Thesis: Harnessing Synthetic Data for Robust and Reliable Vision
Advisor: Prof. Judy Hoffman
Committee: Prof. Dhruv Batra, Prof. James Hays, Prof. Animesh Garg, Dr. Roozbeh Mottaghi
Award: Rising Star Doctoral Student Research Award

M.S. in Computer Science, College of Computing, Georgia Tech 2017-2019
Thesis: Evaluating Visual Conversational Agents via Cooperative Human-AI Games
Advisor: Prof. Devi Parikh
Committee: Prof. Dhruv Batra, Prof. Stefan Lee
Award: M.S. Research Award

B.Tech. in Electrical Engineering, Delhi Technological University (Formerly DCE) 2012-2016

SELECTED RESEARCH EXPERIENCE

Research Scientist, NVIDIA Research 2024-Present
Manager Dr. Ming-Yu Liu Santa Clara, CA
Core Contributor to Cosmos World Foundation Models:

- Cosmos-Reason2: VLM for Physical AI Blog
- Cosmos-Predict2.5: World-Model [Project Page]
- Cosmos-Predict2: World-Model [Project Page]
- Cosmos-Embed1: Video-Text Embedding Model [Project Page]
- Cosmos-Reason1: VLM for Physical AI Reasoning [Project Page]
- Cosmos-Predict1: World Model [Project Page]

Research Assistant, Hoffman Group, Georgia Tech 2019-2024
Advised by Prof. Judy Hoffman Atlanta, GA
Getting vision models to work across changing visual distributions.

- Synthetic Aerial Imagery Benchmark [ECCV24]
- Calibration in Sim2Real Adaptation [ICLR24]
- Sim2Real Generalization [ICCV23]
- Embodied Robustness Benchmark [ICCV21]
- Interpreting Adversarial Robustness [ECCVW20]
- Multi-source Domain Generalization [ECCV20]
- Low-Shot Robustness [ICCV23]
- Language-Guided Counterfactuals [NeurIPS23]
- Backbone Benchmark [NeurIPS23]

Research Intern, PRIOR, Allen Institute for AI Summer 2022
Mentored by Ani Kembhavi, Roozbeh Mottaghi and Judy Hoffman Seattle, WA
Learning representations of environments from house tours to improve sample efficiency and generalization for embodied agents across tasks and simulators

Research Intern, PRIOR, Allen Institute for AI Summer 2020

Mentored by Ani Kembhavi, Roozbeh Mottaghi and Judy Hoffman Atlanta, GA
Benchmark to assess robustness of embodied navigation agents [Project Page][ICCV21]

Research Intern, Deep Learning Group, Microsoft Research AI Summer 2018
Mentored by Hamid Palangi Redmond, WA
Improving goal-driven visually grounded dialog under the presence
of an adversarial utterance evaluator

Research Assistant, Visual Intelligence Lab, Georgia Tech 2017-2019
Mentored by Prof. Devi Parikh and Prof. Dhruv Batra Atlanta, GA
Worked on problems at the intersection of computer vision and natural language processing

- Zero-shot Learning [ECCV18]
- Cooperative Human-AI Games [HCOMP18]
- (Diverse) Generative Visual Dialog [EMNLP19]
- Sub-goals in RL [IJCAI20]
- Evaluating Explanations via Human-AI Teams [EMNLP18]
- AI Challenge Evaluation Framework [SOSPW19]

Research Assistant, CVMLP Lab, Virginia Tech 2016-2017
Mentored by Prof. Devi Parikh and Prof. Dhruv Batra Blacksburg, VA
Worked on problems at the intersection of computer vision and natural language processing

- Counting Objects in Everyday Scenes [CVPR17]
- Human-AI Teams [CVPRW17]

AWARDS AND RECOGNITION

2025 **Outstanding Reviewer** for ICCV
2023 **Awarded ICCV Doctoral Consortium**
2023 **Outstanding Reviewer** for CVPR
2022 **Outstanding Reviewer** for CVPR
2022 **Highlighted Reviewer** for ICLR
2021 **Outstanding Reviewer** for CVPR
2021 **Outstanding Reviewer** for MLRC
2020 **Among Top 33% Reviewers** for ICML
2020 **NVIDIA Best Runner Up Paper Award** at AROW, ECCV
2020 **Rising Star Doctoral Student Award**, School of Interactive Computing, Georgia Tech
2019 **One of the best reviewers** for NeurIPS
2019 **Outstanding Reviewer** for ICLR
2018 **IC Student Travel Grant** to attend NeurIPS
2018 **Among Top 30% Reviewers** for NeurIPS
2018 **MS Research Award**, College of Computing, Georgia Tech
2017 **Subfinalist**, LDV Enterpreneurial Computer Vision Challenge
2017 **Winner**, VTHacks (MLH event at Virginia Tech)
2013 **Semi-Finalists** out of 30 participating teams at ROBOSUB-AUVSI
2013 **Finalists** out of 27 participating teams at NIOT-SAVe
2014 **Merit Scholarships** for Academic Performance 2012-2014
2013 **National Top 1%:** Indian National Physics Olympiad (InPhO)
2013 Cleared **Indian Statistical Institute (ISI)** entrance exam (36 students selected across the country)
2012 **KVPY and INSPIRE Fellowships**

MEDIA & PRESS COVERAGE

Cosmos World Foundation Models (NVIDIA)

- NVIDIA Cosmos launch at CES 2025 – Best of AI & Best of CES [NVIDIA Newsroom]
- Major release of Cosmos world foundation models and Physical AI data tools [NVIDIA Newsroom]
- Cosmos world foundation models platform [NVIDIA Blog]
- NVIDIA releases its own brand of world models [TechCrunch]
- NVIDIA unveils new Cosmos world models [TechCrunch]
- NVIDIA wants to be the Android of generalist robotics [TechCrunch]
- NVIDIA Cosmos AI helps robots and self-driving cars [WIRED]
- NVIDIA signs Toyota to use its self-driving tech [ZDNet]
- NVIDIA launches Cosmos world foundation model [Robotics 24/7]
- Cosmos Reason 2 for physical AI reasoning [VentureBeat]
- Cosmos Predict 2.5 world models [NVIDIA Investor Relations]

Research Coverage

- LANCE: Stress-test method for object recognition shortcuts [TechXplore] [Georgia Tech]
- SkyScenes dataset for autonomous flying vehicles [Georgia Tech] [Mirage News] [GT News]
- Counting objects in everyday scenes featured in CVPR 2017 coverage [GVU Center]
- Work cited in ML for Science textbook [ml-science-book.com]
- Explanations for interpretability discussed in AI Magazine Vol. 45 [AI Magazine]
- Domain generalization work discussed in textbook [Springer]

PREPRINTS

1. H. Ye, K. Zheng, J. Xu, P. Li, H. Chen, J. Han, S. Liu, Q. Zhang, H. Mao, Z. Hao, **P. Chattopadhyay**, D. Yang, L. Feng, M. Liao, J. Bai, M.-Y. Liu, J. Zou, S. Ermon. "Data-regularized Reinforcement Learning for Diffusion Models at Scale." *ArXiv 2025*
2. NVIDIA$^{\Psi}$: A. Ali, J. Bai, M. Bala, Y. Balaji, A. Blakeman, T. Cai, J. Cao, T. Cao, E. Cha, Y.-W. Chao, **P. Chattopadhyay**, M. Chen, Y. Chen, Y. Chen, S. Cheng, Y. Cui, J. Diamond, Y. Ding, J. Fan, L. Fan, L. Feng, F. Ferroni, S. Fidler, X. Fu, R. Gao, Y. Ge, J. Gu, A. Gupta, S. Gururani, I. El Hanafi, A. Hassani, Z. Hao, J. Huffman, J. Jang, P. Jannaty, J. Kautz, G. Lam, X. Li, Z. Li, M. Liao, C.-H. Lin, T.-Y. Lin, Y.-C. Lin, H. Ling, M.-Y. Liu, X. Liu, Y. Lu, A. Luo, Q. Ma, H. Mao, K. Mo, S. Nah, Y. Narang, A. Panaskar, L. Pavao, T. Pham, M. Ramezanali, F. Reda, S. Reed, X. Ren, H. Shao, Y. Shen, S. Shi, S. Song, B. Stefaniak, S. Sun, S. Tang, S. Tasmeen, L. Tchapmi, W.-C. Tseng, J. Varghese, A. Z. Wang, H. Wang, H. Wang, H. Wang, T.-C. Wang, F. Wei, J. Xu, D. Yang, X. Yang, H. Ye, S. Ye, X. Zeng, J. Zhang, Q. Zhang, K. Zheng, A. Zhu, Y. Zhu. "World Simulation with Video Foundation Models for Physical AI." *ArXiv 2025* ($^{\Psi}$ Authors sorted alphabetically)
3. NVIDIA$^{\Psi}$: A. Azzolini, H. Brandon, **P. Chattopadhyay**, H. Chen, J. Chu, Y. Cui, J. Diamond, Y. Ding, F. Ferroni, R. Govindaraju, J. Gu, S. Gururani, I.E. Hanafi, Z. Hao, J. Huffman, J. Jin, B. Johnson, R. Khan, G. Kurian, E. Lantz, N. Lee, Z. Li, X. Li, TY Lin, YC Lin, MY Liu, A. Mathau, Y. Ni, L. Pavao, W. Ping, D. W. Romero, M. Smelyanskiy, S. Song, L. Tchapmi, A. Z. Wang, B. Wang, H. Wang, F. Wei, J. Xu, Y. Xu, X. Yang, Z. Yang, X. Zeng, Z. Zhang. "Cosmos-Reason1: From Physical Common Sense To Embodied Reasoning." *ArXiv 2025* ($^{\Psi}$ Authors sorted alphabetically)
4. NVIDIA$^{\Psi}$: N. Agarwal, A. Ali, M. Bala, Y. Balaji, E. Barker, T. Cai, **P. Chattopadhyay**, Y. Chen, Y. Cui, Y. Ding, D. Dworakowski, J. Fan, M. Fenzi, F. Ferroni, S. Fidler, D. Fox, S. Ge, Y. Ge, J. Gu, S. Gururani, E. He, J. Huang, J. Huffman, P. Jannaty, J. Jin, S. Kim, G. Klár, G. Lam, S. Lan, L. Leal-Taixe, A. Li, Z. Li, C. Lin, T. Lin, H. Ling, M. Liu, X. Liu, A. Luo, Q. Ma, H. Mao, K. Mo, A. Mousavian, S. Nah, S. Niverty, D. Page, D. Paschalidou, Z. Patel, L. Pavao, M. Ramezanali, F. Reda, X. Ren, V. Sabavat, E. Schmerling, S.

Shi, B. Stefaniak, S. Tang, L. Tchapmi, P. Tredak, W. Tseng, J. Varghese, H. Wang, H. Wang, H. Wang, T. Wang, F. Wei, X. Wei, J. Wu, J. Xu, W. Yang, L. Yen-Chen, X. Zeng, Y. Zeng, J. Zhang, Q. Zhang, Y. Zhang, Q. Zhao, A. Zolkowski. "Cosmos World Foundation Model Platform for Physical AI." *ArXiv 2025*
($^{\Psi}$ Authors sorted alphabetically)

5. A. Chandrasekaran*, D.Yadav*, **P. Chattopadhyay***, V. Prabhu*, D. Parikh. "It Takes Two to Tango: Towards Theory of AI's Mind." *ArXiv 2017*
([Talk] *at Chalearn Looking at People Workshop, CVPR 2017)*

PEER-REVIEWED CONFERENCE PAPERS

1. S. Khose*, A. Pal*, A. Agarwal*, D. Deepanshi*, J. Hoffman, **P. Chattopadhyay**. "SkyScenes: A Synthetic Dataset for Aerial Scene Understanding." *European Conference on Computer Vision (ECCV) 2024*
2. **P. Chattopadhyay**, B. Goyal, B. Ecsedi, V. Prabhu, J. Hoffman. "AugCal: Improving Sim2Real Adaptation by Uncertainty Calibration on Augmented Synthetic Images." *International Conference on Learning Representations (ICLR) 2024*
(Also presented at Workshop on Uncertainty Quantification for Computer Vision (UNCV), ICCV 2023)
3. M. Goldblum, H. Souri, R. Ni, M. Shu, V. Prabhu, G. Somepalli, **P. Chattopadhyay**, A. Bardes, M. Ibrahim, J. Hoffman, R. Chellappa, A. Wilson, T. Goldstein. "Battle of the Backbones: A Large-Scale Comparison of Pretrained Models across Computer Vision Tasks." *Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks 2023*
4. V. Prabhu, S. Yenanmandra, **P. Chattopadhyay**, J. Hoffman. "LANCE: Stress-testing Visual Models by Generating Language-guided Counterfactual Images" *Neural Information Processing Systems (NeurIPS) 2023*
5. **P. Chattopadhyay***, K. Sarangmath*, V. Vijaykumar, J. Hoffman. "PASTA: Proportional Amplitude Training Spectrum Augmentation for Syn-to-Real Domain Generalization." *International Conference on Computer Vision (ICCV) 2023*
6. A. Singh, K. Sarangmath, **P. Chattopadhyay**, J. Hoffman. "Benchmarking Low-Shot Robustness to Natural Distribution Shifts." *International Conference on Computer Vision (ICCV) 2023*
7. **P. Chattopadhyay**, J. Hoffman, R. Mottaghi, A. Kembhavi. "RobustNav: Towards Benchmarking Robustness in Embodied Navigation." *International Conference on Computer Vision (ICCV) 2021* [Oral]
(Also presented at Embodied AI Workshop, CVPR 2021)
8. **P. Chattopadhyay**, Y. Balaji, J. Hoffman. "Learning to Balance Specificity and Invariance for In and Out of Domain Generalization." *European Conference on Computer Vision (ECCV) 2020*
(Also presented at Visual Learning with Limited Labels (LwLL), CVPR 2020)
9. N. Modhe, **P. Chattopadhyay**, M. Sharma, A. Das, D. Parikh, D. Batra, R. Vedantam. "IR-VIC: Unsupervised Discovery of Sub-goals for Transfer in RL." *International Joint Conference on Artificial Intelligence (IJCAI) 2020*
10. V. Murahari, **P. Chattopadhyay**, D. Batra, D. Parikh, A. Das. "Improving Generative Visual Dialog by Answering Diverse Questions." *Empirical Methods in Natural Language Processing (EMNLP) 2019*
(Also presented at Visual Question Answering and Dialog Workshop, CVPR 2019)
11. R. Selvaraju*, **P. Chattopadhyay***, M. Elhoseiny, T. Sharma, D. Batra, D. Parikh, S. Lee. "Choose Your Neuron: Incorporating Domain Knowledge Through Neuron-Importance." *European Conference on Computer Vision (ECCV) 2018*
(Also presented at Continual Learning Workshop, NeurIPS 2018)
(Also presented at Visually Grounded Interaction and Language (ViGIL) Workshop, NeurIPS 2018)
12. A. Chandrasekaran*, V. Prabhu*, D.Yadav*, **P. Chattopadhyay***, D. Parikh. "Do Explanations make VQA models more predictable to a human?" *Empirical Methods in Natural Language Processing (EMNLP) 2018*

13. **P. Chattopadhyay***, D.Yadav*, V. Prabhu, A. Chandrasekaran, A. Das, S. Lee, D. Batra, D. Parikh. "Evaluating Visual Conversational Agents via Cooperative Human-AI Games." *AAAI Conference on Human Computation and Crowdsourcing (HCOMP) 2017* [Oral]
14. **P.Chattopadhyay***, R.Vedantam*, R. Selvaraju, D. Batra, D. Parikh. "Counting Everyday Objects in Everyday Scenes." *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017* [Spotlight]

WORKSHOP PAPERS

1. F. Lin, R. Mittapali, **P. Chattopadhyay**, D. Bolya, J. Hoffman. "Likelihood Landscapes: A Unifying Principle Behind Many Adversarial Defenses." *Adversarial Robustness in the Real World (AROW), ECCV 2020* [Talk]
   NVIDIA Best Paper Runner Up 🏆
2. N. Modhe, **P. Chattopadhyay**, M. Sharma, A. Das, D. Parikh, D. Batra, R. Vedantam. "DS-VIC: Unsupervised Discovery of Decision States for Transfer in RL." *Task-Agnostic Reinforcement Learning (TARL) Workshop, ICLR 2019* [Talk]
3. D. Yadav, R. Jain, H. Agrawal, **P. Chattopadhyay**, T. Singh, A. Jain, S. Singh, S. Lee, D. Batra. "EvalAI: Towards Better Evaluation Systems for AI Agents." *Workshop on AI Systems, SOSP 2019*

JOURNAL PAPERS

1. S. Kareer, V. Vijaykumar, H.Maheshwari, **P. Chattopadhyay**, J. Hoffman, V. Prabhu. "We're Not Using Videos Effectively: An Updated Domain Adaptive Video Segmentation Baseline." *Transactions on Machine Learning Research (TMLR) 2024*

TALKS

- **"Harnessing Synthetic Data for Training Robust and Reliable Vision Models"** at NC A&T — April 2024
- **"Reducing Visual Distribution Sensitivity"** at CODA AI Synapse, Georgia Tech — Feb 2024
- **"Reliable Vision for a Changing World"** at DRDO, India — Jan 2024
- **"Reliable Vision for a Changing World"** at Machine Perception, Google — Jan 2023
  (with Viraj Prabhu and Judy Hoffman)

PROFESSIONAL SERVICES

Manuscript Reviewer (🏆indicates reviewer awards)

IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 🏆x3	2018-2024
Neural Information Processing Systems (NeurIPS) 🏆x2	2018-2025
Association for Computational Linguistics (ACL)	2019
International Conference on Learning Representations (ICLR) 🏆x2	2019-2022
IEEE International Conference on Robotics and Automation (ICRA)	2021-2022
International Conference on Machine Learning (ICML) 🏆	2019-2020
International Conference on Computer Vision (ICCV) 🏆x2	2023, 2025
European Conference on Computer Vision (ECCV)	2018
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	2021-2022
Workshop on Uncertainty Quantification for Computer Vision (UNCV), ICCV	2023
Workshop on Distribution Shifts (DistShift), NeurIPS	2021-2022
Machine Learning Reproducibility Challenge (MLRC) 🏆	2021-2022
Workshop on Robustness in Sequence Modeling (RobustSeq), NeurIPS	2022
Learning from Limited and Imperfect Data (L2ID), ECCV	2022
Challenge Organization	
Visual Dialog Challenge	CVPR 2020
Workshop Organization	
EMACS: Experimental Model Auditing with Controllable Synthesis	CVPR 2025
WorldModelBench: The 1st Workshop on Benchmarking World Models	CVPR 2025
Panel Discussion	

Workshop on Reliable and Interactable World Models: Geometry, Physics, Interactivity and Real-World Generalization — ICCV 2025

TEACHING EXPERIENCE

CS 8803: Machine Learning with Limited Supervision — Atlanta, GA
Graduate Teaching Assistant — Fall 2022

CS 4476: Introduction to Computer Vision — Atlanta, GA
Graduate Teaching Assistant — Spring 2021

MENTORING

Sahil Khose, Master's, Georgia Tech — 2023-2024
Now: Ph.D. student in CS at UC Irvine (Advisor: Prof. Judy Hoffman)

Anisha Pal, Master's, Georgia Tech — 2023-2024
Now: Scientific ML Engineer at Planette AI (AI-based climate forecasting)

Vivek Vijaykumar, Bachelor's, Georgia Tech — 2022-2024
Now: Software Engineer at Databricks (AI Data Quality); BS with highest honors

Aaditya Singh, Master's, Georgia Tech — 2022-2023
Now: Research Engineer at AWS AI Rekognition

Aayushi Agarwal, Master's, Georgia Tech — 2021-2023
Now: Senior Software Engineer at Mesh; contributed to SkyScenes dataset

Deepanshi Deepanshi, Master's, Georgia Tech — 2021-2023
Now: Software Engineer (Signal Processing & ML); contributed to SkyScenes

Kartik Sarangmath, Master's, Georgia Tech — 2021-2022
Now: Co-Founder at Omnara (YC S25); prev. ThirdAI, Meta AI

Rohit Mittapalli, Bachelor's, Georgia Tech — 2020-2021
Now: Inference Engineer at Anthropic

Fu Lin, Master's, Georgia Tech — 2020-2021

PROJECTS

Investigating Visual Dialog Models for Goal-Driven Self-Talk [PDF]
As a project for CS 7001: Grad. Studies Computing, Fall 2019

Exploring Weak-Supervision and Generative Models for Semantic Segmentation [PDF]
As a project for CS 8803: Probabilistic Graphical Models, Spring 2018

DTU AUV: Autonomous Underwater Vehicle [PDF]
As a part of DTU-AUV (undergraduate research) team

SELECTED COURSEWORK

Deep Learning, Machine Learning, Machine Learning Theory, Advanced Machine Learning, Probabilistic Graphical Models, Adaptive Control and Reinforcement Learning, Numerical Linear Algebra High Dimensional Data Analytics, Computability and Algorithms

OTHER RESEARCH EXPERIENCE

Research Intern, Robotics Research Lab, IIIT Hyderabad — Winter 2014
Mentored by Prof. K Madhava Krishna — Hyderabad, India
Robotics: Implemented an efficient strategy for a robot to discover, recognize and navigate to a selected few objects among some scattered in an environment

Research Intern, IACS, Kolkata — Summer 2014

Mentored by Prof. Soumitra Sengupta Kolkata, India

Theoretical Physics: Worked on finding Charged Rotating Black Hole solutions in Einstein-Gauss-Bonnet dilaton coupled gravity

Undergraduate Researcher, Autonomous Underwater Vehicle Team, DTU 2012-2016

Mentored by Prof. R K Sinha Delhi, India

Underwater Acoustics: Developed and implemented range estimation algorithms for Passive Source Localization from Time Difference of Arrival (TDOA) values