Viraj Prabhu

e-mail: prabhuviraj@gmail.com
webpage: www.virajprabhu.com

Interests

Developing reliable multimodal foundation models

Employment

Salesforce AI 2024-present
Research Scientist

Education

Georgia Tech 2019-2023
Ph.D. in Computer Science, advised by Judy Hoffman
Thesis: "Towards Reliable Computer Vision Systems"
Committee: Judy Hoffman, Dhruv Batra, Sanja Fidler, Zsolt Kira, James Hays

Georgia Tech 2017-2019
M.S. in Computer Science, advised by Devi Parikh
GPA: 4.0, Awarded M.S. Research Award

BITS Pilani 2011-2015
B.S. in Computer Science

Selected Past Experience

Georgia Tech, Atlanta Fall 2019-2023
Graduate Research Assistant, with Judy Hoffman
Data-efficient and resilient computer vision systems that can be deployed in the real world. Published at ICCV (2021, 2023), NeurIPS (2022, 2023), and BMVC (2021).

NVIDIA Research, Toronto Summer 2022
Research Intern, with Sanja Fidler, James Lucas, and David Acuna
Sim-to-real object detection adaptation for self-driving (published at TMLR 2023).

Salesforce Research, Palo Alto Summer 2021
Research Intern, with Nikhil Naik and Ramprasaath Selvaraju
Adapting visual classifiers to new geographies (published at L3D-IVU, CVPR 2022).

Curai, Palo Alto Summer 2018, 2019
Research Intern, with Anitha Kannan, David Sontag, and Xavier Amatriain
Few-shot learning for dermatological diagnosis (published at MLHC 2019). Open-set machine learning algorithms for diagnosis (published at ML4H at NeurIPS 2019).

Georgia Tech, Atlanta Fall 2017-Spring 2019
Graduate Research Assistant, with Devi Parikh
Human-in-the-loop evaluation of visual conversational agents, and of interpretability mechanisms proposed for such agents (published at HCOMP 2017, EMNLP 2018).

Virginia Tech, Blacksburg Fall 2016-Spring 2017
Visiting Scholar, with Dhruv Batra
Equipping VQA models with mechanisms for detecting the relevance of questions, and with better compositional reasoning (published at EMNLP 2017).

Adobe, Bangalore Summer 2014, Fall 2016-Spring 2017
Member of Technical Staff
Owner of the Android app for Adobe Captivate Prime through two release cycles. Developed and tech-transferred real-time background substitution algorithm for video.

Awards

Outstanding reviewer, NeurIPS 2021	2021
Outstanding reviewer, CVPR 2021	2021
M.S. Research Award, Georgia Tech Computing (1 student annually)	2018
Among top-30% reviewers, NeurIPS 2018	2018
Subfinalist, LDV Entrepreneurial Computer Vision Challenge	2017
Travel Scholarship, for Google Summer of Code Mentor summit	2016, 2017
1st, VTHacks, Virginia Tech's annual hackathon, (> 75 teams)	2017
1st, Google Hackathon at APOGEE 2014, (> 25 teams)	2014
2nd, Technical Project Competition at APOGEE	2013
Top-200 rank, BITSAT 2011 (>120k applicants)	2011
Amul Vidya Shree, awarded to top-100 in ICSE 2009 (>150k applicants)	2009

Publications (representative papers highlighted)

Preprints

28. **Trust but Verify: Programmatic VLM Evaluation in the Wild**
V. Prabhu, S. Purushwalkam, A. Yan, C. Xiong, R. Xu
Under review, 2024

Book Chapters

27. **Few-Shot Learning for Dermatological Disease Diagnosis**
V. Prabhu, A. Kannan, M. Ravuri, M. Chablani, D. Sontag, X. Amatriain
Meta Learning With Medical Imaging and Health Informatics Applications, Elsevier 2022 [Link]

Conference Publications

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

25. **Translating Labels to Solve Annotation Mismatches Across Object Detection Datasets**
A. Liao, D. Acuna, R. Mahmood, J. Lucas, V. Prabhu, S. Fidler
International Conference on Learning Representations (ICLR), 2024 [Paper]

24. **AUGCAL: Sim-to-Real Adaptation by Improving Uncertainty Calibration on Augmented Synthetic Images**
P. Chattopadhyay, B. Goyal, B. Ecsedi, V. Prabhu, J. Hoffman
International Conference on Learning Representations (ICLR), 2024 [Paper]

23. **LANCE: Stress-testing Visual Models by Generating Language-guided Counterfactual Images**
V. Prabhu, S. Yenamandra, P. Chattopadhyay, J. Hoffman
Neural Information Processing Systems (NeurIPS), 2023 [Paper]

22. **Battle of the Backbones: A Large-Scale Comparison of Pretrained Models across Computer Vision Tasks**
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
Neural Information Processing Systems (NeurIPS), Datasets & Benchmarks, 2023 [Paper]

21. **Bridging the Sim2Real gap with CARE: Supervised Detection Adaptation with Conditional Alignment and Reweighting**
V. Prabhu, D. Acuna, A. Liao, R. Mahmood, M. Law, J. Hoffman, S. Fidler, J.

Lucas
Transactions on Machine Learning Research (TMLR), 2023 [Paper]

20. **FACTS: First Amplify Correlations and Then Slice to Discover Bias**
S. Yenamandra, P. Ramesh, V. Prabhu, J. Hoffman
International Conference on Computer Vision (ICCV) 2023.

19. **Adapting Self-Supervised Vision Transformers by Probing Attention-Conditioned Masking Consistency**
V. Prabhu*, S. Yenamandra*, A. Singh, J. Hoffman
Neural Information Processing Systems (NeurIPS) 2022. [Paper]

18. **Mitigating Bias in Visual Transformers via Targeted Alignment**
S. Sudhakar, V. Prabhu, A. Krishnakumar, J. Hoffman
British Machine Vision Conference (BMVC) 2021. [Paper]

17. **Unsupervised Discovery of Bias in Deep Visual Recognition Models**
A. Krishnakumar, V. Prabhu, S. Sudhakar, J. Hoffman
British Machine Vision Conference (BMVC) 2021 [Paper]

16. **SENTRY: Selective Entropy Optimization via Committee Consistency for Unsupervised Domain Adaptation**
V. Prabhu, S. Khare, D. Kartik, J. Hoffman
International Conference on Computer Vision (ICCV) 2021 [Project Page]

15. **Active Domain Adaptation via Clustering Uncertainty-weighted Embeddings**
V. Prabhu, A. Chandrasekaran, K. Saenko, J. Hoffman
International Conference on Computer Vision (ICCV) 2021 [Project Page]

14. **Few-Shot Learning for Dermatological Disease Diagnosis**
V. Prabhu, A. Kannan, M. Ravuri, M. Chablani, D. Sontag, X. Amatriain
Machine Learning and Healthcare Conference, 2019 [Paper]

13. **Do Explanations make VQA Models more Predictable to a Human?**
A. Chandrasekaran*, V. Prabhu*, D. Yadav*, P. Chattopadhyay*, D. Parikh
Conference on Empirical Methods in Natural Language Processing (EMNLP) 2018 [Paper]

12. **The Promise of Premise: Harnessing Question Premises in VQA**
A. Mahendru*, V. Prabhu*, A. Mohapatra*, D. Batra, S. Lee
Conference on Empirical Methods in Natural Language Processing (EMNLP) 2017 [Paper]

11. **Evaluating Visual Conversational Agents via Cooperative Human-AI Games**
P. Chattopadhyay*, D. Yadav*, V. Prabhu, A. Chandrasekaran, A. Das, S. Lee, D. Batra, D. Parikh
AAAI Conference on Human Computation and Crowdsourcing (HCOMP) 2017 [Paper]

Workshop Publications

10. **xGen-MM (BLIP-3): A Family of Open Large Multimodal Models**
L. Xue†, M. Shu†, A. Awadalla*, J. Wang*, A. Yan*, S. Purushwalkam*, H. Zhou*, V. Prabhu*, Y. Dai*, M. Ryoo*, S. Kendre*, J. Zhang*, C. Qin, S. Zhang, C. Chen, N. Yu, J. Tan, T. Awalgaonkar, S. Heinecke, H. Wang, Y. Choi, L. Schmidt, Z. Chen, S. Savarese, J. Niebles, C. Xiong, R. Xu
Workshop on Emergent Visual Abilities and Limits of Foundation Models, ECCV 2024 [Paper]

9. **AUGCAL: Sim-to-Real Adaptation by Improving Uncertainty Calibration on Augmented Synthetic Images**
P. Chattopadhyay, B. Goyal, B. Ecsedi, V. Prabhu, J. Hoffman
Workshop on Uncertainty Quantification for Computer Vision, ICCV 2023

8. **ICON^2: Reliably Benchmarking Predictive Inequity in Object Detection by Identifying and Controlling for Confounders**
S. Sudhakar, V. Prabhu, O. Russakovsky, J. Hoffman
Workshop on Secure and Safe Autonomous Driving, CVPR 2023. [Paper]

7. **AUGCO: Augmentation Consistency-guided Self-training for Source-free Domain Adaptive Semantic Segmentation**
V. Prabhu*, S. Khare*, D. Kartik, J. Hoffman
Workshop on Distribution Shifts (DistShift), NeurIPS 2022. [Paper]

6. **Can domain adaptation make object recognition work for everyone?**
V. Prabhu, R. Selvaraju, J. Hoffman, N. Naik
Workshop on Learning with Limited Labeled Data, CVPR 2022 [Paper]

5. **Open Set Medical Diagnosis**
V. Prabhu, A. Kannan, G. Tso, N. Katariya, M. Chablani, D. Sontag, X. Amatriain
ML for Health Workshop, NeurIPS 2019 [Paper]

4. **Fabrik: An Online Collaborative Neural Network Editor**
U. Garg, V. Prabhu, D. Yadav, R. Ramrakhya, H. Agarwal, D. Batra
Workshop on AI Systems, SOSP 2019 [Paper]

3. **Few-Shot Learning for Dermatological Disease Diagnosis**
V. Prabhu, A. Kannan, M. Ravuri, M. Chablani, D. Sontag, X. Amatriain
ML for Health Workshop, NeurIPS 2018

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

Patents

1. **Systems and methods for responding to healthcare inquiries**
A. Kannan, M. Ravuri, V. Rodrigues, V. Venkataraman, T. Geoffrey, N. Khosla, N. Hunt, X. Amatriain, M. Chablani, D. Sontag, V. Prabhu
US Patent 10,847,265 [Paper]

Talks

Towards Reliable Computer Vision Systems March 2024
Computer Vision Group, California Institute of Technology

Towards Reliable Computer Vision Systems December 2023
AWS Responsible AI

Towards Reliable Computer Vision Systems September 2023
VASC Seminar, Carnegie Mellon University[Link]

Towards Reliable Computer Vision Systems August 2023
UC Berkeley [Video]

Reliable Computer Vision for a Changing World Jan 2023
Google Research Zurich, with Judy Hoffman and Prithivijit Chattopadhyay

Responsible CV: How do models fail and what can we do about it? 2022
Human-Centered AI tutorial at CVPR 2022, with Judy Hoffman [Website]

Introduction to Reinforcement Learning November 2019
Guest lecturer for Deep Learning (Course instructor: Dhruv Batra) [Slides]

Professional Activities

Reviewing

Transactions on Machine Learning Research (TMLR)	2023
International Conference on Computer Vision (ICCV)	2023
Neural Information Processing Systems (NeurIPS)	2018-2023
Conference on Computer Vision and Pattern Recognition (CVPR)	2018, 2021-2022
Winter Conference on Applications of Computer Vision (WACV)	2022
International Conference on Learning Representations (ICLR)	2018, 2020, 2024
Association for Computational Liguistics (ACL)	2019
European Conference on Computer Vision and Pattern Recognition (ECCV)	2018

Mentoring

Simar Kareer, PhD student, Georgia Tech	2022-2023
Sriram Yenamandra, Master's student, Georgia Tech	2022-2023
Aaditya Singh, Master's student, Georgia Tech	2022
Sruthi Sudhakar, Bachelor's student, Georgia Tech	2021-2022
Shivam Khare, Master's student, Georgia Tech	2020-2021
Deeksha Kartik, Master's student, Georgia Tech	2020-2021
Arvind Krishnakumar, Master's student, Georgia Tech	2021
Utsav Garg, Bachelor's student, NUS (Google Summer of Code 2017)	2017
Gaurav Gupta, Bachelor's student, IIT BHU (Google Summer of Code 2016)	2016

Workshop Organization

Learning from Limited and Imperfect Data (L2ID), ECCV 2022	2022

Other Projects

Fabrik, an Online Collaborative Neural Network Editor Summer 2016-2017
Lead mentor and maintainer of Fabrik, an open-source web platform to collaboratively build, visualize, and design neural networks in the browser (1000+ GitHub stars).

Learning Cooperative Visual Dialog Agents via Deep RL Fall 2017
PyTorch code for Das & Kottur *et al.*, ICCV '17. Used for the 2018 Visual Dialog challenge. [Code] (160+ GitHub stars)

Learning Active Learning Policies for Visual Recognition Spring 2019
Learning active learning policies for visual recognition via RL. [Report]

Cooperative Visual Dialog Models with Mental Models Fall 2017
Explored self-play strategies based on dialog rollouts to develop cooperative visual dialog agents.[Poster]

Exploring Generative Models for Semantic Segmentation Spring 2018
Semantic segmentation via deep probabilistic generative models. [Report]

Relevant Coursework

Machine Learning, Deep Learning, Computer Vision, Advanced ML
Adaptive Control and Reinforcement Learning, Probabilistic Graphical Models
Computability & Algorithms, High-dimensional Data Analytics
Information Retrieval, Parallel Computing, Advanced Algorithms

Teaching Experience	**Introduction to Computer Vision**, Georgia Tech Head teaching assistant with Judy Hoffman	Spring 2021
	Deep Learning, Georgia Tech Teaching assistant with Dhruv Batra	Fall 2019
	Introduction to Machine Learning, Virginia Tech Teaching assistant with Stefan Lee	Fall 2016