

# Dinesh JAYARAMAN

Assistant Professor of Computer and Information Science, University of Pennsylvania

WEB:<https://www.seas.upenn.edu/~dineshj/> EMAIL:[dineshj@seas.upenn.edu](mailto:dineshj@seas.upenn.edu)

OFFICE:401 Levine Hall, 3330 Walnut St, Philadelphia PA-19104

CITIZENSHIP: Indian

## RESEARCH INTERESTS

LAST UPDATED: APRIL 15, 2025

My research group, PennPAL (Perception, Action, and Learning), works at the intersections of robotics, computer vision, and machine learning. Our recent research directions aim to improve vision-based robot learning by infusing “attentive abstractions” during observation, representation, policy mapping, and learning, to help the robot attend to task-relevant information on-the-fly. This enables efficiencies in pre-training, task specification, and resource allocation, paving the way for general-purpose robots of the future.

## EDUCATION

AUG 2011 - AUG 2017 | Doctor of Philosophy  
ELECTRICAL AND COMPUTER ENGINEERING, **UT Austin**

AUG 2007 - JUN 2011 | Bachelor of Technology  
ELECTRICAL ENGINEERING, **IIT Madras, India.**

## POSITIONS HELD

JAN 2020 - now | *Assistant Professor at Department of Computer and Information Science (CIS),  
University of Pennsylvania*

AUG 2019 - DEC 2019 | *Visiting Research Scientist at Facebook Artificial Intelligence Research, **Facebook***

SEP 2017 - AUG 2019 | *Postdoctoral Scholar at Berkeley Artificial Intelligence Research Laboratory, **UC Berkeley***

## FUNDED PROJECT GRANTS AND GIFTS

- OCT 2024 | NSF National Artificial Intelligence Research Resource (NAIRR) Pilot 240077. “Safe Recovery from Robotic Failures.” (USD 20k in compute)
- SEP 2024 | DARPA TIAMAT HR00112490421, as Penn PI (USD 700k). Automatic transfer of conceptual policies via simulator similarities and continua. Project led by GE Vernova Research.
- OCT 2023 | National Science Foundation (NSF) Safe Learning-Enabled Systems 2331783: SpecsRL: Specification-Guided Perception-Enabled Conformally-Safe Reinforcement Learning (co-PI, PI: Rajeev Alur), 1.5M USD over 10/01/2023-09/30/2027, shared equally among 4 PIs.
- MAY 2023 | National Science Foundation (NSF) CAREER 2239301: Flexible hierarchical abstractions for actionable visual perception (solo PI), 600k USD over 05/01/2023-04/30/2028.
- MAR 2023 | (Internal) Penn University Research Foundation, Immersive situated supervision for robot learning. (solo PI) 75k USD over 03/01/2023-02/29/2024.
- AUG 2022 | U.S. Office of Naval Research ONR N00014-22-1-2677: Uncertainty-based self-learning for perception (co-PI, PI: Kostas Daniilidis), 2.5M USD over 09/01/2022-08/31/2027.
- MAY 2022 | NEC Labs gift funding (solo PI): embodied visual recognition, 113k USD.
- MAR 2021 | Amazon Research Award gift funding (solo PI): Self-to-world modeling for robotic control, 100k USD.
- MAR 2021 | GE Research gift funding (PI, together with Pratik Chaudhari): Uncertainty-based diverse trajectory prediction, 25k USD.

## ACADEMIC HONORS AND AWARDS

- MAY 2024 Best Vision Paper Award Finalist, International Conference on Robotics and Automation, Yokohoma, Japan.
- MAY 2023 NSF CAREER Award.
- DEC 2022 Best Paper Award, Conference on Robot Learning, Auckland, New Zealand.
- MAR 2021 Amazon Research Award.
- FEB 2021 Selected for the AAAI New Faculty Highlights Program.
- MAY 2019 2018 Best Paper Award Runner-Up, IEEE Robotics and Automation-Letters (RA-L).
- MAY 2019 Work Featured on Cover Page of Science Robotics Issue.
- OCT 2017 University Nominee for ACM Best Dissertation Award, UT Austin 2017.
- NOV 2016 Best Application Paper Award at ACCV 2016, Taipei.
- JUL 2016 Graduate Dean's Prestigious Fellowship Supplement, UT Austin, 2016-17.
- JUN 2016 Outstanding Reviewer Award at CVPR 2016, Las Vegas, Nevada.
- APR 2016 Samsung PhD Fellowship, 2016-17.
- MAR 2016 Invited paper for IJCV Special Issue of Best Papers from ICCV 2015.
- OCT 2014 NVIDIA Academic grant award, 2014.
- AUG 2011 Microelectronics and Computer Development Fellowship, UT Austin, 2011-12.
- MAY 2011 Dr. Dilip Veeraraghavan Memorial Award, IIT Madras, 2011.
- MAY 2011 K. Srinivasan and Indira Srinivasan Prize, IIT Madras, 2011.
- MAY 2009 Rajalakshmi Krishnamurthy English Prize, IIT Madras, 2011.
- MAR 2009 OPJEMS National Award for academic and leadership excellence, O.P.Jindal group, 2009-10.
- FEB 2009 Winning team member, Robocon India robotics competition, 2009.
- MAY 2008 Certificate of Merit, Dept. of Electrical Engineering, IIT Madras.

## PEER-REVIEWED CONFERENCE PUBLICATIONS

---

1. Pratik Kunapuli, Jake Welde, Dinesh Jayaraman, and Vijay Kumar. Leveling the playing field: Carefully comparing classical and learned controllers for quadrotor trajectory tracking. *RSS*, 2025
2. Christopher Watson, Arjun Krishna, Rajeev Alur, and Dinesh Jayaraman. Illustrated landmark graphs for long-horizons policy learning. *TMLR*, 2025
3. Arjun Krishna, Edward S. Hu, and Dinesh Jayaraman. The value of sensory information to a robot. *ICLR*, 2025
4. Yecheng Jason Ma, Joey Hejna, Chuyuan Fu, Dhruv Shah, Jacky Liang, Zhuo Xu, Sean Kirmani, Peng Xu, Danny Driess, Ted Xiao, Osbert Bastani, Dinesh Jayaraman, Wenhao Yu, Tingnan Zhang, Dorsa Sadigh, and Fei Xia. Vision language models are in-context value learners. *ICLR*, 2025
5. Edward S. Hu, Kwangjun Ahn, Qinghua Liu, Haoran Xu, Manan Tomar, Ada Langford, Dinesh Jayaraman, Alex Lamb, and John Langford. Learning to achieve goals with belief state transformers. *ICLR*, 2025
6. Long Le, Jason Xie, William Liang, Hung-Ju Wang, Yue Yang, Yecheng Jason Ma, Kyle Vedder, Arjun Krishna, Dinesh Jayaraman, and Eric Eaton. Articulate-anything: Automatic modeling of articulated objects via a vision-language foundation model. *ICLR*, 2025
7. Kaustubh Sridhar, Souradeep Dutta, Dinesh Jayaraman, and Insup Lee. Regent: A retrieval-augmented generalist agent that can act in-context in new environments. *ICLR*, 2025
8. Tianyou Wang Ian Pedroza Amy Luo Jie Wang Jason Ma Dinesh Jayaraman Junyao Shi\*, Zhuolun Zhao\*. Zeromimic: Distilling robotic manipulation skills from web videos. *ICRA*, 2025
9. Jake Welde\*, Nishanth Rao\*, Pratik Kunapuli\*, Dinesh Jayaraman, and Vijay Kumar. Leveraging symmetry to accelerate learning of trajectory tracking controllers for free-flying robotic systems. *ICRA*, 2025

10. Jianing Qian, Bernadette Bucher, and Dinesh Jayaraman. Task-oriented hierarchical object decomposition for visuomotor control. *CORL*, 2024
11. William Liang, Sam Wang, Hungju Wang, Osbert Bastani, Dinesh Jayaraman\*, and Yecheng Jason Ma\*. Environment curriculum generation via large language models. *CORL*, 2024
12. Weilin Wan, Zhiyang Dou, Taku Komura, Wenping Wang, Dinesh Jayaraman, and Lingjie Liu. Tlcontrol: Trajectory and language control for human motion synthesis. *ECCV*, 2024
13. Yinsen Jia, Jingxi Xu, Dinesh Jayaraman, and Shuran Song. Learning a meta-controller for dynamic grasping. *CASE*, 2024
14. Jianing Qian, Anastasios Panagopoulos, and Dinesh Jayaraman. Recasting generic pretrained vision transformers as object-centric scene encoders for manipulation policies. *ICRA*, 2024
15. Junyao Shi\*, Jianing Qian\*, Yecheng Jason Ma, and Dinesh Jayaraman. Composing pre-trained object-centric representations for robotics from “what” and “where” foundation models. *ICRA*, 2024
16. Edward Hu, James Springer, Oleh Rybkin, and Dinesh Jayaraman. Privileged sensing scaffolds reinforcement learning. *ICLR*, 2024 [**Spotlight Presentation, 5% accept rate**]
17. Chuan Wen, Dinesh Jayaraman, and Yang Gao. Can transformers capture spatial relations between objects? *ICLR*, 2024
18. Kaustubh Sridhar, Souradeep Dutta, Dinesh Jayaraman, James Weimer, and Insup Lee. Memory-consistent neural networks for imitation learning. *ICLR*, 2024
19. Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. *ICLR*, 2024
20. Large collaboration. Open X-Embodiment: Robotic learning datasets and RT-X models. *ICRA*, 2024
21. Zichen Zhang, Yunshuang Li, Osbert Bastani, Abhishek Gupta, Dinesh Jayaraman, Yecheng Jason Ma, and Luca Weihs. Universal visual decomposer: Long-horizon manipulation made easy. *ICRA*, 2024 [**Best Paper, CORL '23 Workshop on Learning Effective Abstractions for Planning**]
22. Kyle Vedder, Neehar Peri, Nathaniel Chodosh, Ishan Khatri, Eric Eaton, Dinesh Jayaraman, Yang Liu, Deva Ramanan, and James Hays. Zeroflow: Fast zero label scene flow via distillation. *ICLR*, 2024
23. Sriram Narayanan, Dinesh Jayaraman, and Manmohan Chandraker. Long-hot: A modular hierarchical approach for long-horizon object transport. *ICRA*, 2024
24. Ashwin De Silva, Rahul Ramesh, Lyle Ungar, Marshall Hussain Shuler, Noah J. Cowan, Michael Platt, Chen Li, Leyla Isik, Seung-Eon Roh, Adam Charles, Archana Venkataraman, Brian Caffo, Javier J. How, Justus M Kebschull, John W. Krakauer, Maxim Bichuch, Kaleab Alemayehu Kinfu, Eva Yezerets, Dinesh Jayaraman, Jong M. Shin, Soledad Villar, Ian Phillips, Carey E. Priebe, Thomas Hartung, Michael I. Miller, Jayanta Dey, Ningyuan Huang, Eric Eaton, Ralph Etienne-Cummings, Elizabeth L. Ogburn, Randal Burns, Onyema Osuagwu, Brett Mensh, Alysson R. Muotri, Julia Brown, Chris White, Weiwei Yang, Andrei A. Rusu Timothy Verstynen, Konrad P. Kording, Pratik Chaudhari, and Joshua T. Vogelstein. Prospective learning: Principled extrapolation to the future. In *Proceedings of The 2nd Conference on Lifelong Learning Agents*. PMLR, 2023
25. Leon Kim, Yunshuang Li, Michael Posa, and Dinesh Jayaraman. Vision-based contact localization without touch or force sensing. *CORL*, 2023
26. Yecheng Jason Ma, Vikash Kumar, Amy Zhang, Osbert Bastani, and Dinesh Jayaraman. LIV: Language-image representations and rewards for robotic control. *ICML*, 2023
27. Yecheng Jason Ma, Kausik Sivakumar, Jason Yen, Osbert Bastani, and Dinesh Jayaraman. Learning policy-aware models for model-based reinforcement learning via transition occupancy matching. *L4DC*, 2023
28. Edward Hu, Richard Chang, Oleh Rybkin, and Dinesh Jayaraman. Planning goals for exploration. *ICLR (top 25 percent) and Best Workshop Paper at CORL 2022 Robot Adaptation Workshop*, 2023 [**Spotlight Presentation, 5% accept rate**] [**CoRL'22 Roboadapt Workshop, Best Paper Award**]
29. Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. VIP: Towards universal visual reward and representation via Value-Implicit Pre-Training. *ICLR (top 25 percent)*, 2023 [**Spotlight Presentation, 5% accept rate**]

30. Kun Huang, Edward Hu, and Dinesh Jayaraman. Training robots to evaluate robots: Example-based interactive reward functions for policy learning. *CORL*, 2022 [**Best Paper Award**]
31. Yecheng Jason Ma, Jason Yan, Dinesh Jayaraman, and Osbert Bastani. How far i'll go: Offline goal-conditioned reinforcement learning via  $f$ -advantage regression. *NeurIPS*, 2022 [**Nominated for Outstanding Paper**]
32. Jianing Qian, Anastasios Panagopoulos, and Dinesh Jayaraman. Discovering deformable keypoint pyramids. *ECCV*, 2022
33. Chuan Wen, Jianing Qian, Jierui Lin, Jiaye Teng, Dinesh Jayaraman, and Yang Gao. Fighting fire with fire: Avoiding dnn shortcuts through priming. *ICML*, 2022
34. Yecheng Jason Ma, Andrew Shen, Dinesh Jayaraman, and Osbert Bastani. Smodice: Versatile offline imitation learning via state occupancy matching. *ICML*, 2022
35. Edward S. Hu, Kun Huang, Oleh Rybkin, and Dinesh Jayaraman. Know thyself: Transferable visuomotor control through robot-awareness. *ICLR*, 2022
36. Yecheng Jason Ma, Andrew Shen, Osbert Bastani, and Dinesh Jayaraman. Conservative and adaptive penalty for model-based safe reinforcement learning. *AAAI*, 2022
37. Yecheng Jason Ma, Dinesh Jayaraman, and Osbert Bastani. Conservative offline distributional reinforcement learning. *NeurIPS*, 2021
38. Chuan Wen, Jierui Lin, Jianing Qian, Yang Gao, and Dinesh Jayaraman. Keyframe-focused visual imitation learning. *ICML*, 2021
39. Nikos Kolotouros, Georgios Pavlakos, Dinesh Jayaraman, and Kostas Daniilidis. Embracing the reconstruction uncertainty in 3d human pose estimation. *ICCV*, 2021
40. Yecheng Jason Ma, Jeevana Priya Inala, Dinesh Jayaraman, and Osbert Bastani. Likelihood-based diverse sampling for trajectory forecasting. *ICCV*, 2021
41. Jingxi Xu, Bruce Lee, Nikolai Matni, and Dinesh Jayaraman. How are learned perception-based controllers impacted by the limits of robust control? *L4DC*, 2021
42. Glen Berseth, Daniel Geng, Coline Devin, Chelsea Finn, Dinesh Jayaraman, and Sergey Levine. SMiRL: Surprise minimizing rl in dynamic environments. *ICLR*, 2021 (Oral Presentation)
43. Neha Das, Sarah Bechtle, Todor Davchev, Dinesh Jayaraman, Akshara Rai, and Franziska Meier. Model-based inverse reinforcement learning from visual demonstrations. *CORL*, 2020
44. Karl Pertsch, Oleh Rybkin, Frederik Ebert, Dinesh Jayaraman, Chelsea Finn, and Sergey Levine. Long-horizon visual planning with goal-conditioned hierarchical predictors. *NeurIPS*, 2020
45. Chuan Wen, Jierui Lin, Trevor Darrell, Dinesh Jayaraman, and Yang Gao. Fighting copycat agents in behavioral cloning from observation histories. *NeurIPS*, 2020
46. Jesse Zhang, Brian Cheung, Chelsea Finn, Sergey Levine, and Dinesh Jayaraman. Cautious adaptation for reinforcement learning in safety-critical settings. *ICML*, 2020
47. Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, Dinesh Jayaraman, and Roberto Calandra. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *ICRA and IEEE RA-L*, 2020
48. Brian Yang\*, Dinesh Jayaraman\*, Glen Berseth, Alexei Efros, and Sergey Levine. MAVRIC: Morphology-agnostic visual robotic control. *ICRA and IEEE RA-L*, 2020
49. Pim de Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. *NeurIPS*, 2019
50. Brian Yang, Dinesh Jayaraman, Jesse Zhang, and Sergey Levine. Replab: A reproducible low-cost arm benchmark for robotic learning. *ICRA*, 2019
51. Stephen Tian, Frederik Ebert, Dinesh Jayaraman, Mayur Mudigonda, Chelsea Finn, Roberto Calandra, and Sergey Levine. Manipulation by feel: Touch-based control with deep predictive models. *ICRA*, 2019
52. Dinesh Jayaraman, Frederik Ebert, Alexei A Efros, and Sergey Levine. Time-agnostic prediction: Predicting predictable video frames. *ICLR*, 2019
53. Dinesh Jayaraman, Ruohan Gao, and Kristen Grauman. Shapecodes: self-supervised feature learning by lifting views to viewgrids. *ECCV*, 2018

54. Roberto Calandra, Andrew Owens, Dinesh Jayaraman, Justin Lin, Wenzhen Yuan, Jitendra Malik, Edward H Adelson, and Sergey Levine. More than a feeling: Learning to grasp and regrasp using vision and touch. *IROS and IEEE RA-L*, 2018
55. Dinesh Jayaraman and Kristen Grauman. Learning to look around: Intelligently exploring unseen environments for unknown tasks. *CVPR*, 2018
56. Yu-Chuan Su, Dinesh Jayaraman, and Kristen Grauman. Pano2vid: Automatic cinematography for watching 360-degree videos. *ACCV*, 2016
57. Ruohan Gao, Dinesh Jayaraman, and Kristen Grauman. Object-centric representation learning from unlabeled videos. *ACCV*, 2016
58. Dinesh Jayaraman and Kristen Grauman. Look-ahead before you leap: end-to-end active recognition by forecasting the effect of motion. *ECCV*, 2016 (Oral Presentation)
59. Dinesh Jayaraman and Kristen Grauman. Slow and steady feature analysis: higher order temporal coherence in video. *CVPR*, 2016 (Oral Spotlight Presentation)
60. Dinesh Jayaraman and Kristen Grauman. Learning image representations tied to ego-motion. *ICCV*, 2015 (Oral Presentation)
61. Dinesh Jayaraman and Kristen Grauman. Zero-shot recognition with unreliable attributes. *NeurIPS*, 2014
62. Dinesh Jayaraman, Fei Sha, and Kristen Grauman. Decorrelating semantic visual attributes by resisting the urge to share. *CVPR*, 2014
63. Dinesh Jayaraman, Anish Mittal, Anush K Moorthy, and Alan C Bovik. Objective quality assessment of multiply distorted images. *ASILOMAR Signals, Systems and Computers*, 2012

## JOURNAL PUBLICATIONS

---

1. Menachem Stern, Sam Dillavou, Dinesh Jayaraman, Douglas J Durian, and Andrea J Liu. Training self-learning circuits for power-efficient solutions. *Applied Physics Letters (APL) Machine Learning*, 2024
2. Hui Chen, Zhao Li, Sheng Feng, Anni Wang, Melissa Richard-Greenblatt, Emily Hutson, Stefen Andrianus, Laurel J. Glaser, Kyle G. Rodino, Jianing Qian, Dinesh Jayaraman, Ronald G. Collman, Abigail Glascock, Fred-eric D. Bushman, Jae Seung Lee, Sara Cherry, Alejandra Fausto, Susan R. Weiss, Hyun Koo, Patricia M. Corby, Una O'Doherty, Alfred L. Garfall, Dan T. Vogl, Edward A. Stadtmauer, and Ping Wang. Femtomolar sars-cov-2 antigen detection using the microbubbling digital assay with smartphone readout enables antigen burden quantitation and dynamics tracking. Cold Spring Harbor Laboratory Press, 2021
3. Santhosh K Ramakrishnan, Dinesh Jayaraman, and Kristen Grauman. An exploration of embodied visual exploration. *IJCV*, 2021
4. Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, Dinesh Jayaraman, and Roberto Calandra. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *ICRA and IEEE RA-L*, 2020
5. Brian Yang\*, Dinesh Jayaraman\*, Glen Berseth, Alexei Efros, and Sergey Levine. MAVRIC: Morphology-agnostic visual robotic control. *ICRA and IEEE RA-L*, 2020
6. Santhosh K Ramakrishnan\*, Dinesh Jayaraman\*, and Kristen Grauman. Emergence of exploratory look-around behaviors through active observation completion. *Science Robotics*, 2019
7. Roberto Calandra, Andrew Owens, Dinesh Jayaraman, Justin Lin, Wenzhen Yuan, Jitendra Malik, Edward H Adelson, and Sergey Levine. More than a feeling: Learning to grasp and regrasp using vision and touch. *IROS and IEEE RA-L*, 2018 (Best Paper Award Runner-Up)
8. Dinesh Jayaraman and Kristen Grauman. End-to-end policy learning for active visual categorization. *IEEE TPAMI*, 2018
9. Dinesh Jayaraman and Kristen Grauman. Learning image representations tied to egomotion from unlabeled video. *IJCV Special Issue of Best Papers from ICCV 2015*, 2017

## BOOK CHAPTERS

---

1. Dinesh Jayaraman, Chao-Yeh Chen, Fei Sha, and Kristen Grauman. Divide, share, and conquer: Multi-task attribute learning with selective sharing. In *Visual attributes*, pages 49–85. Springer, Cham, 2017

## PEER-REVIEWED WORKSHOP PAPERS AND ABSTRACTS

---

1. Jake Welde\*, Nishanth Rao\*, Pratik Kunapuli\*, Dinesh Jayaraman, and Vijay Kumar. Leveraging symmetry to accelerate learning of trajectory tracking controllers for free-flying robotic systems. *ICRA*, 2025 [Best Contribution - Neuroscience & Interpretability, at Workshop on Symmetry and Geometry in Neural Representations (NeurReps) at NeurIPS 2024]
2. Jianing Qian and Dinesh Jayaraman. Object representations guided by optical flow. *NeurIPS 4th Robot Learning Workshop: Self-Supervised and Lifelong Learning*, 2021

## PATENTS

---

1. Dinesh Jayaraman, Oscar Nestares, and Kalpana Seshadrinathan. "Techniques for improved focusing of camera arrays." U.S. Patent 9,743,016, issued August 22, 2017.
2. Dinesh Jayaraman, Tao Ma, Wei Sun, Oscar Nestares, and Kalpana Seshadrinathan. "Techniques for rectification of camera arrays." U.S. Patent 9,875,543, issued January 23, 2018.

## THESES

---

- Dinesh Jayaraman, "Embodied Learning for Visual Recognition". PhD Thesis. Supervisor: Prof. Kristen Grauman, UT Austin, 2017.
- Dinesh Jayaraman, "Modeling Natural and Distorted Image Statistics". Bachelors Thesis. Supervisor: Prof. R. Aravind, IIT Madras, 2011.

## WIDER PRESS AND TECHNICAL BLOG COVERAGE

---

- The Economist's Babbage Podcast: "How AI can make robots fit for a human world." <https://www.economist.com/podcasts/2024/06/12/how-ai-can-make-robots-fit-for-a-human-world> (Interview) (Jun 12, 2024)
- Excitement and Money Flows to Intersection of AI and Robotics <https://www.pymnts.com/artificial-intelligence-2024/excitement-and-money-flows-to-intersection-of-ai-and-robotics/>, (Interview) (Feb 27, 2024)
- Inside UPenn's new Artificial Intelligence degree program debuting fall 2024 <https://www.fox29.com/news/university-of-pennsylvania-school-of-engineering-to-offer-new-artificial-intelligence-ai-degree> (Interview) (Feb 23, 2024)
- GRASP Faculty Members Awarded Two Inaugural NSF SLES Grants: "Specification-guided Perception-enabled Conformal Safe Reinforcement Learning" <https://www.grasp.upenn.edu/news/nsf-sles-jayaraman/>, (Interview) (Dec 2023)
- NVIDIA Top 10 projects of 2023 [https://www.youtube.com/watch?v=d7RM-3dQaJ8&ab\\_channel=NVIDIADeveloper](https://www.youtube.com/watch?v=d7RM-3dQaJ8&ab_channel=NVIDIADeveloper) (Work covered) (Dec 2023)
- Eureka: With GPT-4 overseeing training, robots can learn much faster <https://arstechnica.com/information-technology/2023/10/eureka-uses-gpt-4-and-massively-parallel-simulations-to-accelerate-robot-training/> (Work covered) (Oct 2023)
- Engadget: NVIDIA's latest AI model helps robots perform pen spinning tricks as well as humans <https://news.yahoo.com/nvidias-latest-ai-model-helps-robots-perform-pen-spinning-tricks-as-well-as-humans-13.html> (Oct 2023)
- TechCrunch, Embodied AI spins a pen and helps clean the living room in new research. <https://techcrunch.com/2023/10/20/embodied-ai-spins-a-pen-and-helps-clean-the-living-room-in-new-research/> (Oct, 2023)
- VentureBeat, New Nvidia AI agent, powered by GPT-4, can train robots. <https://venturebeat.com/ai/new-nvidia-ai-agent-powered-by-gpt-4-can-train-robots/>. (Oct, 2023)

- Nature: Teaching robots to touch. <https://www.nature.com/articles/d41586-022-01401-y>, (May 26, 2022)
- VentureBeat: Facebook battles the challenges of tactile sensing. <https://venturebeat.com/2021/11/01/facebook-battles-the-challenges-of-tactile-sensing/>
- NSF.gov: “NSF advances artificial intelligence research with new nationwide institutes”, [https://www.nsf.gov/news/special\\_reports/announcements/082620.jsp](https://www.nsf.gov/news/special_reports/announcements/082620.jsp), (Aug 26, 2020).
- TechXplore: “DIGIT: A high-resolution tactile sensor to enhance robot in-hand manipulation skills.” <https://techxplore.com/news/2020-07-digit-high-resolution-tactile-sensor-robot.html>, July 2020.
- FreeNews: “Scientists have made a high-resolution tactile sensor: the robot will work more efficiently.” <https://freenews.live/scientists-have-made-a-high-resolution-tactile-sensor-the-robot-will-work-more> July 2020.
- VentureBeat: “Facebook’s Digit is a low-cost tactile sensor for robotic hands.” <https://venturebeat.com/2020/06/01/facebooks-digit-is-a-low-cost-tactile-sensor-for-robotic-hands/>, June 2020.
- NVIDIA developer news center: “NVAAIL Partners present AI Research at ICLR.” <https://news.developer.nvidia.com/nvail-partners-present-ai-research-at-iclr/>, May 2020.
- TechXplore: “REPLAB: a low-cost benchmark platform for robotic learning.” <https://techxplore.com/news/2019-05-replab-low-cost-benchmark-platform-robotic.html>
- Science Daily: “New AI sees like a human, filling in the blanks”. <https://www.sciencedaily.com/releases/2019/05/190515144017.htm>. May 2019.
- Psychology Today: “Scientists Create Human-like AI Computer Vision”. <https://www.psychologytoday.com/us/blog/the-future-brain/201905/scientists-create-human-ai-computer-vision>, (May 29, 2019).
- Science Robotics Front Page: “Developing a good eye” <https://robotics.sciencemag.org/content/4/30>, May 2019.
- Interesting Engineering, “New Artificial Intelligence Sees Like a Human, Bringing Us Closer to Skynet”. May 2019. <https://interestingengineering.com/new-artificial-intelligence-sees-like-a-human-bringing-us-closer-to-skynet>
- UT Austin News: “New AI Sees Like a Human, Filling in the Blanks”, <https://news.utexas.edu/2019/05/15/new-ai-sees-like-a-human-filling-in-the-blanks/>, May 2019.
- Samsung PhD Fellowship Program Recognizes Best and Brightest Student Innovators, <https://news.samsung.com/global/samsung-phd-fellowship-program-recognizes-best-and-brightest-student-innovators/1000>, May 2016

## PROFESSIONAL SERVICE

---

### Service within the University of Pennsylvania:

- Organizer, Departmental Colloquium, Department of Computer and Information Science, '22, '23.
- Recruitment Ambassadors for outreach to URM applicants, Spring 2023 - Present.
- Organizer, Industry Day, GRASP lab. Fall 2023.
- Faculty Representative, GRASP Compute Cluster Administration, Summer 2024 - Present.
- UPenn Apple Scholarship Committee, 2024.
- Organizer, GRASP Lab Summit: What’s Next in Robotics, Spring '25.
- IDEAS faculty recruitment committee, Spring '25
- SEAS Strategic Planning Committee, Spring '25

### Organizer:

- RSS Workshop on Robotic Tasks and How to Specify Them? – Task Specification for General-Purpose Intelligent Robots, RSS '24
- Workshop on Tactile Sensing for Robotics, Facebook Menlo Park, November '21.
- RSS Workshop on Visual Learning and Reasoning for Robotics, RSS '20, '21

- ICLR Workshop on Embodied Multimodal Learning, ICLR '21
- Workshop on Benchmarking in Robotics, Pittsburgh, August '19.
- ECCV Workshop on Action and Anticipation for Visual Learning, ECCV '16

#### **Proposal Review:**

- National Science Foundation Proposals Review Panel, '20, '22, '23, '24
- Air Force Office of Scientific Research Proposals Reviewer, '22
- Army Research Office Proposals Reviewer Panel, '21
- National Science Foundation Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Program Review Panel, '20.
- Hong Kong University Grants Commission Review Panel, '20

#### **Senior Area Chair:**

- Association for the Advancement of Artificial Intelligence (AAAI) Conference on Artificial Intelligence, '22, '23.

#### **Area Chair:**

- Robotics Science and Systems (RSS), '25.
- Computer Vision and Pattern Recognition (CVPR), '24
- International Conference on Computer Vision (ICCV), '23.
- International Conference on Machine Learning (ICML), '23, '25.
- International Conference on Learning Representations (ICLR), '21, '22, '23.
- Neural and Information Processing Systems (NeurIPS), '18, '19, '20, '21, '22.
- Association for the Advancement of Artificial Intelligence (AAAI) Conference on Artificial Intelligence, '20, '21.

#### **Associate Editor:**

- International Conference on Robotics and Automation (ICRA) '20, '21, '22.

#### **Session Chair:**

- Conference on Robot Learning (CORL), '22.
- International Conference on Robotics and Automation (ICRA), '22.
- International Conference on Learning Representations (ICLR), '21.
- Neural and Information Processing Systems (NeurIPS), '18, '20.

#### **Local Arrangements Chair:**

- Koenderink Symposium on the Art and Science of Vision, UC Berkeley, October '18.

#### **Conference Program Committee Member / Reviewer:**

- International Conference on Learning Representations '22
- Robotics: Science and Systems (RSS) '18, '21, '22
- Computer Vision and Pattern Recognition (CVPR) '15, '16, '17, '18, '20
- International Conference on Robotics and Automation (ICRA) '18, '19, '23
- Conference on Robot Learning (CORL) '18, '19
- European Conference on Computer Vision (ECCV) '16, '18
- Asian Conference on Computer Vision (ACCV) '16, '18
- International Conference on Computer Vision (ICCV) '15, '17
- International Conference on Image Processing Theory, Tools and Applications (IPTA) '17

- Neural and Information Processing Systems (NeurIPS) '15
- International Conference on Machine Learning (ICML) '15
- Workshop on Egocentric Perception, Interaction and Computing, ICCV '19
- Workshop on Multitask and Lifelong Reinforcement Learning, ICML '19
- Egocentric Vision Workshop at CVPR '16

**Journal Action Editor:**

- Transactions on Machine Learning Research (TMLR), '22, '23, '24.

**Journal Reviewer:**

- Transactions on Pattern Analysis and Machine Intelligence (TPAMI) '16,'17,'18,'19,'21,'22.
- IEEE Robotics and Automation-Letters (RA-L) '18,'19
- Journal of Machine Learning Research (JMLR) '19
- International Journal of Computer Vision (IJCV) '17,'18,'19.
- Machine Vision and Applications (MVAP) '17
- Transactions in Image Processing (TIP) '16

**Award Committee:**

- ACM India Doctoral Dissertation Award Committee. '23.

**Thesis Committees:**

- Soroush Seifi, (external), PhD thesis: Visual attention in partially observable environments, KU Leuven, '22. (PhD Advisor: Tinne Tuytelaars)
- Elijah Lee, CIS PhD Thesis Committee Chair, UPenn '22.
- Karl Schmeckpeper, CIS PhD Thesis Committee Chair, UPenn, '22.
- Chao Qu, CIS PhD Thesis: Deep Basis Fitting for Depth Completion, UPenn, '22.
- Oleh Rybkin, CIS PhD Thesis Committee Chair: Learning to Action from Diverse Data Sources via World Models, '23.
- Nikolaos Kolotouros, CIS PhD Thesis: Reconstructing 3D Humans from Images, '22.
- Wenxin Liu, CIS PhD Thesis: Visual-Inertial State Estimation with Information Deficiency, '22.
- Georgios Pavlakos, CIS PhD thesis: Learning to Reconstruct 3D Humans, University of Pennsylvania, '20.
- Jyh-Jing Hwang, CIS PhD thesis: Learning Image Segmentation with Relation-Centric Loss and Representation, University of Pennsylvania, '20.
- Wenxin Liu, WPE II Committee, Uncertainty Estimation for Regression Using Neural Networks, '20.
- Ian Miller, ESE Provisional Doctoral Committee, '21.
- Daniel Feshbach, WPE II Committee, '22.
- Rahul Ramesh, WPE II Committee, '22.
- Nikolaos Kolotouros, WPE II Committee, '21.
- Stephen Mell, WPE II Committee, '21.
- Elijah Lee, WPE II Committee, '21.
- Kyle Vedder, WPE II Committee, '21.
- Tony Liu, WPE II Committee Chair, '21.

**Research Advising (current PhD students):**

- Aurora Qian
- Edward Hu

- Jason Ma (co-advised with Osbert Bastani)
- Leon Kim (co-advised with Michael Posa)
- Pratik Kunapuli (co-advised with Vijay Kumar)
- Junyao Shi
- Arjun Krishna
- Chris Watson (co-advised with Rajeev Alur)

**Research Advising (current Masters students):**

- Hungju Wang
- Tony Jie Wang
- Tianyou Wang
- George Gao
- Oliver Dong
- Xingfang Yuan
- Arjun Arasappan

**Research Advising (current undergraduate students):**

- Sam Wang
- Lilian Muyao Li
- Fiona Luo
- Josh Smith
- Ethan Yu
- Rujia Yang (visiting from Tsinghua)
- Ian Pedroza
- Amy Luo

**Other Student Collaborators (PhD students and postdocs only):**

- Kaustubh Sridhar
- Kyle Vedder (graduated 2024, now at Dyna Robotics)
- Chuan Wen (graduated Tsinghua University 2019-2024, moved to Assistant Professor at SJTU)
- Weilin Wan (University of Hong Kong)
- Oleg Rybkin (UC Berkeley)

**Research Advising (graduated students):**

- James Springer
- (Masters Student) Yunshuang Li (SEAS outstanding Masters student researcher award)
- (Masters Student) Kun Huang (SEAS outstanding Masters student researcher award)
- (Masters Student) Kausik Sivakumar (GRASP outstanding Masters student researcher award)
- (Undergraduate student) Andrew Shen
- (Undergraduate student) Anastasios Panagopoulos
- Will Liang

## TEACHING

---

- CIS 4190/5190 Applied Machine Learning, University of Pennsylvania – Spring '20, '21, '22, '23 & Fall '21, '24.
- Winter School Course on Introduction to Machine Learning, Infosys Campus, Mysuru, India, December '22.
- CIS 7000 Vision-Based Robot Learning – Fall '23.
- CIS 5800 Machine Perception – Fall '22
- CIS 5220 Deep Learning: Reinforcement Learning Module, University of Pennsylvania – Spring '21
- CIS 7000 Data-Driven Robotic Perception and Control, University of Pennsylvania – Fall '20

## INVITED AND CONFERENCE TALKS

---

- MAR 2025 The Pretrain-Finetune-Adapt Paradigm for Robots, TIAMAT PI Meeting, Berkeley, CA.  
FEB 2025 Engineering Better Robot Learners: Exploitation and Exploration, Keynote at AAAI Doctoral Consortium, Philadelphia, PA.  
FEB 2025 How to Train Your Robot, AI Research Seminar Series, UPenn.  
OCT 2024 Exploring and Exploiting Robot Learning, Cornell Robotics Seminar, Cornell University.  
OCT 2024 Exploring and Exploiting Robot Learning, Michigan AI Symposium, University of Michigan.  
SEP 2024 How To Train Your Robot, Nemertes [Next] Conference, Baltimore, Maryland.  
SEP 2024 How To Train Your Robot, ASSET & Warren Center Research Mixer Event, UPenn.  
AUG 2024 Attention and Inattention in Minimalist Robot Learners, OSU ECE Colloquium, The Ohio State University.  
JUL 2024 How to Train Your Robot, Penn Rising Scholars Success Academy Outreach Talk, 2024.  
JUL 2024 Attention and Inattention in Minimalist Robot Learners, Boston Dynamics AI Institute, Cambridge.  
MAY 2024 How to Train Your Robot, AI for Executives: Daylong Immersion at Penn Engineering, UPenn.  
APR 2024 The Sensory Needs of Robot Learners, Lehigh University.  
MAR 2024 Attention and Inattention in Robot Learners, University of Michigan.  
DEC 2023 The Sensory Needs of Robot Learners, NSF Convergence Workshop, Washington DC.  
OCT 2023 Attentive Abstractions for Flexible Vision-Based Robot Learners, Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA), Philadelphia.  
OCT 2023 Attentive Abstractions For Flexible Vision-Based Robot Learners, UC Berkeley.  
OCT 2023 Attentive Abstractions For Flexible Vision-Based Robot Learners, Stanford University.  
MAY 2023 Vision-Language Representations for Robots, 2023 Embedded Vision Summit, Santa Clara.  
MAR 2023 Polyglot Robots: Goal-Based Task Specification in Vision, Language, and More, ASSET Seminar Series.  
MAR 2023 Polyglot Robots: Goal-Based Task Specification in Vision, Language, and More, MIT Embodied Intelligence Seminar.  
FEB 2023 Robot Attention, Rachleff Scholars Guest Lecture, University of Pennsylvania.  
JAN 2023 Training Scrupulous and Safe Robots, Microsoft Research Trustworthy AI Workshop, Bengaluru, India.  
SEP 2022 Scaling Visual Robot Learning By Attending to the Right Things, ECE Colloquium Series, University of Delaware.  
SEP 2022 Scaling Visual Robot Learning By Attending to the Right Things, GE Research Edge Symposium.  
AUG 2022 Spatial and Temporal Abstractions for Scaling Visual Robot Learning, IIT Bombay.  
AUG 2022 Spatial and Temporal Abstractions for Scaling Visual Robot Learning, IIT Madras.  
AUG 2022 Self-Supervised Abstract Scene Representations for Vision-Based Robot Learning, Office of Naval Research Science of Autonomy Meeting.  
APR 2022 How Human-Machine Miscommunication Breaks Machine Learning, Virtual Quakers Day talk to incoming undergraduate class at UPenn.  
FEB 2022 Towards Scalable Visual Reinforcement Learning for Robotics, Computational Sensorimotor Learning Seminar Series, Massachusetts Institute of Technology.  
OCT 2021 Learning Visual Dynamics Models for Robotic Control Policies, Guest Lecture at MEAM 520 Introduction to Robotics, University of Pennsylvania.  
AUG 2021 Tutorial on Reinforcement Learning: Machine Learning for Decision Making and Control, 5th Summer School on Artificial Intelligence, International Institute of Information Technology, Hyderabad, India.

- FEB 2021 AAAI New Faculty Highlight Talk, 35th AAAI Conference on Artificial Intelligence (Virtual), 2021.
- SEP 2020 Intelligent Sensing Summer School, Queen Mary University of London, United Kingdom.
- DEC 2019 "Causal Confusion in Imitation Learning", Invited Talk at Center for Human-Compatible AI, Berkeley, CA.
- NOV 2019 "Understanding Embodied Visual Intelligence", Invited Talk at Samsung Strategy and Innovation Center Forum, San Jose, CA.
- OCT 2019 "REPLAB: Challenges and Opportunities of Low-Cost Robots", Invited Tutorial at CORL, Osaka, Japan.
- JUL 2019 "Towards Embodied Visual Intelligence", Facebook AI Research, Menlo Park, California, USA.
- JUL 2019 "Vision and Action: Visual Approaches to Robotic Control", Invited Talk at the VISUM Summer School on Computer Vision and Machine Learning, Porto, Portugal.
- MAY 2019 "Visual Learning for Tactile Manipulation", Invited Talk at ViTac Workshop on Integrating Vision and Touch for Multimodal and Cross-modal Perception, ICRA 2019, Montreal, Canada.
- JUL 2018 "Embodied learning for vision", Invited talk at ICML workshop on "Learning with Limited Labels: Invariance, Equivariance, and Beyond", Stockholm, Sweden.
- MAR 2017 "Embodied learning for visual recognition", Robotics Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.
- FEB 2017 "Embodied learning for visual recognition", Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA.
- FEB 2017 "Embodied learning for visual recognition", Toyota Technological Institute, Chicago, Illinois, USA.
- FEB 2017 "Embodied learning for visual recognition", Stanford Computer Vision Lab, Stanford University, Palo Alto, California, USA.
- FEB 2017 "Embodied learning for visual recognition", Berkeley Artificial Intelligence Research Laboratory, The University of California at Berkeley, Berkeley, California, USA.
- DEC 2016 "Embodied learning for visual recognition", Image Processing and Computer Vision Lab, IIT Madras, Chennai, India.
- OCT 2016 "Look-ahead before you leap: end-to-end active recognition by forecasting the effect of motion", Oral presentation at European Conference on Computer Vision (ECCV), Amsterdam, the Netherlands.
- JUN 2016 "Slow and steady feature analysis", Oral spotlight presentation at Computer Vision and Pattern Recognition (CVPR), Las Vegas, Nevada, USA.
- DEC 2015 "Learning image representations tied to ego-motion", Oral presentation at International Conference on Computer Vision (ICCV), Santiago, Chile.
- DEC 2015 "Embodied learning of image representations from video", Image Processing and Computer Vision Lab, IIT Madras, Chennai, India.
- DEC 2014 "Zero-shot recognition and cross-view geolocation", Computer Vision group at Cornell Tech, New York City, New York, USA.
- SEP 2014 "Zero-shot recognition with unreliable attributes", Computer Vision group at UC Berkeley, Berkeley, California, USA.
- JUN 2014 "Decorrelating semantic visual attributes", Oral presentation at Computer Vision and Pattern Recognition (CVPR), Columbus, Ohio, USA.