

YILUN DU

Harvard Science and Engineering Complex
150 Western Ave, 6.138
Allston, MA 02134

ydu@seas.harvard.edu
<https://yilundu.github.io/>

Research Overview

My research is driven by the goal of developing intelligent decision-making agents. My research focuses on using generative AI to construct models of the world, using planning on top of the models as a way to determine the actions of the agent. A unique challenge of constructing world models is the lack of available decision-making data and the necessity to generalize well to previously unseen situations. My work addresses this by proposing compositional generative modeling, where simple generative models are learned and composed together to construct complex generative models that generalize beyond the narrow amount of available data. My work has illustrated how such compositionality solves various problems in constructing AI agents, ranging from complex scene understanding, reasoning, multimodal perception, and planning. My work has further shown how such developed techniques can be broadly applied to settings in both the sciences and engineering.

Research Areas: Generative AI, Machine Learning, Computer Vision, Robotics, AI for Science.

Education

Ph.D. EECS 2020-2024	Massachusetts Institute of Technology (MIT) Advisors: Leslie Kaelbling, Tomás Lozano-Pérez, and Joshua B. Tenenbaum Thesis: <i>Learning to Generate Compositionally</i>
M.S. EECS 2019-2020	Massachusetts Institute of Technology (MIT) Advisors: Leslie Kaelbling, Tomás Lozano-Pérez, and Joshua B. Tenenbaum Thesis: <i>Implicit Learning with Energy-Based Models</i>
B.S. 2015-2019	Massachusetts Institute of Technology (MIT) EECS, GPA: 5.0 / 5.0

Employment

Assistant Professor Summer 2025	Harvard University Kempner Institute and Computer Science
Senior Research Scientist 2024-2025	Google Deepmind Research Scientist on the Deepmind Robotics Team
Graduate Researcher 2019-2024	Computer Science and Artificial Intelligence Lab, MIT LIS Group (PIs: Tomás Lozano-Pérez and Leslie Pack Kaelbling)
Graduate Researcher 2019-2024	Brain and Cognitive Science Department, MIT Computational Cognitive Science Lab (PI: Joshua Tenenbaum)
Research Advisor 2023-2024	Pika Labs Text-to-Video Synthesis / Running Research Projects

Research Intern Summer 2023	Google Robotics Research Intern Working on Video Models for Robotics
Student Researcher 2022-2023	Google Brain Student Researcher Working on Video Models for Decision Making
Research Intern Summer 2022	Deepmind Research Intern Working on Compositional Generative Modeling
Visiting Researcher 2020-2022	Facebook AI Research Visiting Researcher on the FAIR Proteins (Generative Biology) Team
Research Intern Summer 2019	Facebook AI Research Research Intern on the FAIR Proteins (Generative Biology) Team
Research Fellow Summer/Fall 2018	OpenAI Research Fellow Working on Generative Models

Awards and Honors

Outstanding Paper Finalist at RSS	2026
Outstanding Paper Award at ICLR	2024
Best Paper Award at NeurIPS Instruction Following Workshop	2023
Outstanding Paper Award at ICRA RAP4Robots Workshop	2023
Outstanding Paper Award at NeurIPS Workshop on Controllable Generative Modeling	2021
National Science Foundation (NSF) Graduate Research Fellowship	2020-2024

Publications

The * sign denotes equal contribution.

Peer-Reviewed Papers

- C129** Haoyu Chen, Kaichen Zhou, Hang Hua, Kaile Zhang, Jingwen Qian, Wufei Ma, Haonan Chen, Chunjiang Liu, Yizhou Zhao, Xiaoyuan Wang, Weiyue Li, Alan Yuille, Paul Pu Liang, Yilun Du. “Memo-Bench: Benchmarking World Modeling in Dynamically Changing Environments” *European Conference on Computer Vision (ECCV)*, 2026.
- C128** Kaichen Zhou*, Yuzhen Chen*, Fangneng Zhan, Hang Hua, Grace Chen, Xinhai Chang, Ao Qu, Yilun Du, Zhuang Liu, Paul Pu Liang[†], Mengyu Wang[†]. “GEM-4D: Geometry-Enhanced Video World Models for Robot Manipulation” *European Conference on Computer Vision (ECCV)*, 2026.
- C127** Chaoqi Liu, Xiaoshen Han, Jiawei Gao, Yue Zhao, Haonan Chen, Yilun Du. “OAT: Ordered Action Tokenization” *Robotics: Science and Systems (RSS)*, 2026.
- C126** Zhiyi Li, Peilin Wu*, Xiaoshen Han*, Ruojin Cai, Yilun Du. “Structured 4D Latent Predictive Model for Robot Planning” *International Conference on Machine Learning (ICML)*, 2026.

- C125** Haldun Balim, Na Li[†], Yilun Du[†]. “Model-Based Diffusion Sampling for Predictive Control in Offline Decision Making” *International Conference on Machine Learning (ICML)*, 2026.
- C124** Hansen Jin Lillemark*, Benhao Huang*, Fangneng Zhan, Yilun Du, T. Anderson Keller. “Flow Equivariant World Models: Memory for Partially Observed Dynamic Environments” *International Conference on Machine Learning (ICML)*, 2026.
- C123** Evgenii Opryshko*, Junwei Quan*, Claas Voelcker, Yilun Du, Igor Gilitschenski. “Test-Time Graph Search for Goal-Conditioned Reinforcement Learning” *International Conference on Machine Learning (ICML)*, 2026.
- C122** Renhao Zhang, Haotian Fu, Mingxi Jia, George Konidaris, Yilun Du, Bruno Castro da Silva. “From Noise to Control: Parameterized Diffusion Policies” *International Conference on Machine Learning (ICML)*, 2026.
- C121** Maijunxian Wang, Ruisi Wang, Juyi Lin, Ran Ji, Thaddäus Wiedemer, Qingying Gao, Dezhi Luo, Yaoyao Qian, Lianyu Huang, Zelong Hong, Jiahui Ge, Qianli Ma, Hang He, Yifan Zhou, Lingzi Guo, Lantao Mei, Jiachen Li, Hanwen Xing, Tianqi Zhao, Fengyuan Yu, Weihang Xiao, Yizheng Jiao, Jianheng Hou, Danyang Zhang, Pengcheng Xu, Boyang Zhong, Zehong Zhao, Gaoyun Fang, John Kitaoka, Yile Xu, Hua Xu, Kenton Blacutt, Tin Nguyen, Siyuan Song, Haoran Sun, Shaoyue Wen, Linyang He, Runming Wang, Yanzhi Wang, Mengyue Yang, Ziqiao Ma, Raphaël Milliére, Freda Shi, Nuno Vasconcelos, Daniel Khashabi, Alan Yuille, Yilun Du, Ziming Liu, Bo Li, Dahua Lin, Ziwei Liu, Vikash Kumar, Yijiang Li, Lei Yang, Zhongang Cai, Hokin Deng. “A Very Big Video Reasoning Suite” *International Conference on Machine Learning (ICML)*, 2026.
- C120** Minda Zhao, Yilun Du, Mengyu Wang. “Large Language Models Are Bad Dice Players: LLMs Struggle to Generate Random Numbers from Statistical Distributions” *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026.
- C119** Zhiqiu Lin, Chancharik Mitra, Siyuan Cen, Isaac Li, Yuhan Huang, Yu Tong Tiffany Ling, Hewei Wang, Irene Pi, Shihang Zhu, Ryan Rao, George Liu, Jiayi Li, Ruojin Li, Yili Han, Yilun Du, Deva Ramanan. “CHAI: Building a Precise Video Language with Human-AI Oversight” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.
- C118** Haowen Liu*, Shaoxiong Yao*, Haonan Chen, Jiawei Gao, Jiayuan Mao, Jia-Bin Huang, Yilun Du. “SIMPACT: Simulation-Enabled Action Planning using Vision-Language Models” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.
- C117** Yifan Liu, Fangneng Zhan, Kaichen Zhou, Yilun Du, Paul Pu Liang, Hanspeter Pfister. “Abstract 3D Perception for Spatial Intelligence in Vision-Language Models” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.
- C116** Runhan Huang, Haldun Balim, Heng Yang, Yilun Du. “Flexible Locomotion Learning with Diffusion Model Predictive Control” *International Conference on Robotics and Automation (ICRA)*, 2026.
- C115** Haonan Chen, Jiaming Xu, Hongyu Chen, Kaiwen Hong, Binghao Huang, Chaoqi Liu, Jiayuan Mao, Yunzhu Li, Yilun Du[†], Katherine Driggs-Campbell[‡]. “Multi-Modal Manipulation via Multi-Modal Policy Consensus” *International Conference on Robotics and Automation (ICRA)*, 2026.
- C114** Aayush Karan, Yilun Du. “Reasoning with Sampling: Your Base Model is Smarter Than You Think” *International Conference on Learning Representations (ICLR)*, 2026.
- C113** Chenyu Wang*, Zishen Wan*, Hao Kang, Emma Chen, Zhiqiang Xie, Tushar Krishna, Vijay Reddi, Yilun Du. “SLM-MUX: Orchestrating Small Language Models for Reasoning” *International Conference on Learning Representations (ICLR)*, 2026.
- C112** Xiangcheng Zhang, Haowei Lin, Haotian Ye, James Zou, Jianzhu Ma, Yitao Liang, Yilun Du. “Inference-time Scaling of Diffusion Models through Classical Search” *International Conference on Learning Representations (ICLR)*, 2026.

- C111** Jen-Yuan Huang, Tong Ling, Yilun Du. “Long-Text-to-Image Generation via Compositional Prompt Decomposition” *International Conference on Learning Representations (ICLR)*, 2026.
- C110** Jiahao Zhang*, Muqing Jiang*, Nanru Dai, Taiming Lu, Arda Uzunoglu, Shunchi Zhang, Yana Wei, Jiahao Wang, Vishal M. Patel, Paul Pu Liang, Daniel Khashabi, Cheng Peng, Rama Chellappa, Tianmin Shu, Alan Yuille, Yilun Du, Jieneng Chen. “World-in-World: World Models in a Closed-Loop World” *International Conference on Learning Representations (ICLR)*, 2026.
- C109** Yunzhi Zhang, Carson Murtuza-Lanier, Zizhang Li, Yilun Du, Jiajun Wu. “Product of Experts for Visual Generation” *International Conference on Learning Representations (ICLR)*, 2026.
- C108** Calvin Luo*, Zilai Zeng*, Mingxi Jia, Yilun Du, Chen Sun. “Self-Improving Loops for Visual Robotic Planning” *International Conference on Learning Representations (ICLR)*, 2026.
- C107** Yiming Li, Nael Darwiche, Amirreza Razmjoo, Sichao Liu, Yilun Du, Auke Ijspeert, Sylvain Calinon. “Geometry-aware Policy Imitation” *International Conference on Learning Representations (ICLR)*, 2026.
- C106** Alexi Gladstone, Ganesh Nanduru, Md Mofijul Islam, Peixuan Han, Hyeonjeong Ha, Aman Chadha, Yilun Du, Heng Ji, Jundong Li, Tariq Iqbal. “Energy-Based Transformers are Scalable Learners and Thinkers” *International Conference on Learning Representations (ICLR)*, 2026.
- C105** Jaeyeon Kim*, Lee Cheuk-Kit*, Carles Domingo-Enrich, Yilun Du, Sham Kakade, Timothy Ngatiaoco, Sitan Chen, Michael S. Albergo. “Any-Order Flexible Length Masked Diffusion” *International Conference on Learning Representations (ICLR)*, 2026.
- C104** Zhiqiu Lin*, Siyuan Cen*, Daniel Jiang, Jay Karhade, Hewei Wang, Chancharik Mitra, Yo Tong Tiffany Ling, Yuhao Huang, Sifan Liu, Mingyu Chen, Rushikesh Zawar, Xue Bai, Yilun Du, Chuang Gan, Deva Ramanan. “Towards Understanding Camera Motion in Any Video” *Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*, 2025.
- C103** Zhenting Qi, Fan Nie, Alexandre Alahi, James Zou, Himabindu Lakkaraju, Yilun Du, Eric Xing, Sham Kakade, Hanlin Zhang. “EvoLM: In Search of Lost Language Model Training Dynamics” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C102** Alexandru Oarga, Yilun Du. “Generalizable Reasoning through Compositional Energy Minimization” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C101** Yunhao Luo, Utkarsh A. Mishra, Yilun Du[†], Danfei Xu. “Generative Trajectory Stitching through Diffusion Composition” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C100** Hao Kang, Qingru Zhang, Han Cai, Weiyan Xu, Tushar Krishna, Yilun Du, Tsachy Weissman. “Win Fast or Lose Slow: Balancing Speed and Accuracy in Latency-Sensitive Decisions of LLMs” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C99** Yuncong Yang*, Jiageng Liu*, Zheyuan Zhang, Siyuan Zhou, Reuben Tan, Jianwei Yang, Yilun Du, Chuang Gan. “MindJourney: Test-Time Scaling with World Models for Spatial Reasoning” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C98** Siyuan Zhou, Yilun Du, Yuncong Yang, Lei Han, Peihao Chen, Dit-Yan Yeung, Chuang Gan. “Learning 3D Persistent Embodied World Models” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C97** Louis Béthune, David Vigouroux, Yilun Du, Rufin VanRullen, Thomas Serre, Victor Boutin. “Follow the Energy, Find the Path: Riemannian Metrics from Energy-Based Models” *Neural Information Processing Systems (NeurIPS)*, 2025.
- C96** Yuncong Yang, Han Yang, Jiachen Zhou, Peihao Chen, Hongxin Zhang, Yilun Du, Chuang Gan. “3D-Mem: 3D Scene Memory for Embodied Exploration and Reasoning” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

- C95** Shalev Lifshitz, Sheila A. McIlraith, Yilun Du. “Multi-Agent Verification: Scaling Test-Time Compute with Multiple Verifiers” *Conference on Language Modeling (CoLM)*, 2025.
- C94** Haoyu Zhen*, Qiao Sun*, Hongxin Zhang, Junyan Li, Siyuan Zhou, Yilun Du, Chuang Gan. “Tesser-Act: Learning 4D Embodied World Models” *International Conference on Computer Vision (ICCV)*, 2025.
- C93** Runhan Huang*, Shaoting Zhu*, Yilun Du, Hang Zhao. “MoELoco: Mixture of Experts for Multitask Locomotion” *International Conference on Intelligent Robots and Systems (IROS)*, 2025.
- C92** Yanbo Wang, Justin Dauwels, Yilun Du. “Compositional Scene Understanding through Inverse Generative Modeling” *International Conference on Machine Learning (ICML)*, 2025.
- C91** Kiwhan Song*, Boyuan Chen*, Max Simchowitz, Yilun Du, Russ Tedrake, Vincent Sitzmann. “History-Guided Video Diffusion” *International Conference on Machine Learning (ICML)*, 2025.
- C90** Shenyan Gao, Siyuan Zhou, Yilun Du, Jun Zhang, Chuang Gan. “Learning Adaptable World Models with Latent Actions” *International Conference on Machine Learning (ICML)*, 2025.
- C89** Zishen Wan, Jiayi Qian, Yuhang Du, Jason Jabbour, Yilun Du, Yang (Katie) Zhao, Arijit Raychowdhury, Tushar Krishna, Vijay Reddi. “Generative AI in Embodied Systems: System-Level Analysis of Performance, Efficiency and Scalability” *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2025.
- C88** Sigmund H. Hoeg, Yilun Du, Olav Egeland. “Streaming Diffusion Policy: Fast Policy Synthesis with Variable Noise Diffusion Models” *International Conference on Robotics and Automation (ICRA)*, 2025.
- C87** Yanwei Wang, Lirui Wang, Yilun Du, Balakumar Sundaralingam, Xuning Yang, Yu-Wei Chao, Claudia Perez-D’Arpino, Dieter Fox, Julie Shah. “Inference-Time Policy Steering through Human Interactions” *International Conference on Robotics and Automation (ICRA)*, 2025.
- C86** Siddharth Ancha, Sunshine Jiang, Travis Manderson, Laura Brandt, Yilun Du, Phillip Osteen, Nicholas Roy. “Anomalies-by-Synthesis: Anomaly Detection using Generative Diffusion Models for Off-Road Navigation” *International Conference on Robotics and Automation (ICRA)*, 2025.
- C85** Vighnesh Subramaniam*, Yilun Du*, Joshua B. Tenenbaum, Antonio Torralba, Shuang Li, Igor Mordatch. “Multiagent Finetuning: Self-Improvement with Diverse Reasoning Chains” *International Conference on Learning Representations (ICLR)*, 2025.
- C84** Yunhao Luo, Yilun Du. “Grounding Video Models to Actions through Goal Conditioned Exploration” *International Conference on Learning Representations (ICLR)*, 2025.
- C83** Calvin Luo, Zilai Zeng, Yilun Du, Chen Sun. “Solving New Tasks by Adapting Internet Video Knowledge” *International Conference on Learning Representations (ICLR)*, 2025.
- C82** Ying Fan, Yilun Du, Kannan Ramchandran, Kangwook Lee. “Looped Transformers for Length Generalization” *International Conference on Learning Representations (ICLR)*, 2025.
- C81** Wenhao Chai*, Enxin Song*, Yilun Du, Chenlin Meng, Vashisht Madhavan, Omer Bar-Tel, Jenq-Neng Hwang, Saining Xie, Christopher D. Manning. “AuroraCap: Efficient, Performant Video Detailed Captioning and a New Benchmark” *International Conference on Learning Representations (ICLR)*, 2025.
- C80** Hongxin Zhang*, Zeyuan Wang*, Qiushi Lyu*, Zheyuan Zhang, Sunli Chen, Tianmin Shu, Yilun Du, Chuang Gan. “COMBO: Compositional World Models for Embodied Multi-Agent Cooperation” *International Conference on Learning Representations (ICLR)*, 2025.
- C79** Julia Briden, Yilun Du, Enrico Zucchelli, Richard Linares. “Compositional Diffusion Models for Powered Descent Trajectory Generation with Flexible Constraints” *IEEE Aerospace Conference*, 2025.

- C78** Aviv Netanyahu, Yilun Du, Antonia Bronars, Jyotish Pari, Joshua B. Tenenbaum, Tianmin Shu, Pulkit Agrawal. “Few-Shot Task Learning through Inverse Generative Modeling” *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- C77** Boyuan Chen, Diego Marti Monso, Yilun Du, Max Simchowitz, Russ Tedrake, Vincent Sitzmann. “Diffusion Forcing: Next-token Prediction Meets Full-Sequence Diffusion” *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- C76** Yinan Zhang, Eric Tzeng, Yilun Du, Dmitry Kislyuk. “Large-scale Reinforcement Learning for Diffusion Models” *International Conference on Computer Vision (ICCV)*, 2024.
- C75** Jie Ye, Andrew Luo, Yilun Du, Anoop Cherian, Tim K. Marks, Jonathan Le Roux, Chuang Gan. “Disentangled Acoustic Fields For Multimodal Physical Scene Understanding” *International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- C74** Yiqing Xu, Jiayuan Mao, Yilun Du, Tomás Lozano-Pérez, Leslie Kaelbling, David Hsu. “‘Set It Up!’: Functional Object Arrangement with Compositional Generative Models” *Robotics: Science and Systems (RSS)*, 2024.
- C73** Lirui Wang, Alan Zhao, Yilun Du, Edward Adelson, Russ Tedrake. “PoCo: Policy Composition from and for Heterogeneous Robot Learning” *Robotics: Science and Systems (RSS)*, 2024.
- C72** Yilun Du, Leslie Kaelbling. “Compositional Generative Modeling: A Single Model is Not All You Need” *International Conference on Machine Learning (ICML)*, 2024.
- C71** Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, Igor Mordatch. “Improving Factuality and Reasoning in Language Models through Multiagent Debate” *International Conference on Machine Learning (ICML)*, 2024.
- C70** Yilun Du*, Jiayuan Mao*, Joshua B. Tenenbaum. “Learning Iterative Reasoning through Energy Diffusion” *International Conference on Machine Learning (ICML)*, 2024.
- C69** Yunhao Luo, Chen Sun, Joshua B. Tenenbaum, Yilun Du. “Potential Based Diffusion Motion Planning” *International Conference on Machine Learning (ICML)*, 2024.
- C68** Jocelin Su*, Nan Liu*, Yanbo Wang*, Joshua B. Tenenbaum, Yilun Du. “Compositional Image Decomposition with Diffusion Models” *International Conference on Machine Learning (ICML)*, 2024.
- C67** Siyuan Zhou, Yilun Du, Jiaben Chen, Yandong Li, Dit-Yan Yeung, Chuang Gan. “RoboDreamer: Learning Compositional World Models for Robot Imagination” *International Conference on Machine Learning (ICML)*, 2024.
- C66** Sherry Yang, Jacob Walker, Jack Parker-Holder, Yilun Du, Jake Bruce, Andre Barreto, Pieter Abbeel, Dale Schuurmans. “Video as the New Language for Real-World Decision Making” *International Conference on Machine Learning (ICML)*, 2024.
- C65** Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, Chuang Gan. “3D-VLA: A 3D Vision-Language-Action Generative World Model” *International Conference on Machine Learning (ICML)*, 2024.
- C64** Yilun Du, Sherry Yang, Pete Florence, Fei Xia, Ayzaan Wahid, Brian Ichter, Pierre Sermanet, Tianhe Yu, Pieter Abbeel, Joshua B. Tenenbaum, Leslie Kaelbling, Andy Zeng, Jonathan Tompson. “Video Language Planning” *International Conference on Learning Representations (ICLR)*, 2024.
- C63** Kevin Black*, Michael Janner*, Yilun Du, Ilya Kostrikov, Sergey Levine. “Training Diffusion Models with Reinforcement Learning” *International Conference on Learning Representations (ICLR)*, 2024.
- C62** Tailin Wu*, Takashi Maruyama*, Long Wei*, Tao Zhang*, Yilun Du*, Gianluca Iaccarino, Jure Leskovec. “Compositional Generative Inverse Design” *International Conference on Learning Representations (ICLR)*, 2024.

- C61** Mengjiao Yang*, Yilun Du*, Bo Dai, Dale Schuurmans, Joshua B. Tenenbaum, Pieter Abbeel. “Probabilistic Adaptation of Text-to-Video Models” *International Conference on Learning Representations (ICLR)*, 2024.
- C60** Mengjiao Yang, Yilun Du, Kamyar Ghasemipour, Jonathan Tompson, Dale Schuurmans, Pieter Abbeel. “Learning Interactive Real-World Simulators” *International Conference on Learning Representations (ICLR)*, 2024.
- C59** Po-Chen Ko, Jiayuan Mao, Yilun Du, Shao-Hua Sun, Joshua B. Tenenbaum. “Learning to Act from Actionless Video through Dense Correspondences” *International Conference on Learning Representations (ICLR)*, 2024.
- C58** Yichen Li, Yilun Du, Chao Liu, Francis Williams, Michael Foshey, Benjamin Eckart, Jan Kautz, Joshua B. Tenenbaum, Antonio Torralba, Wojciech Matusik. “Learning to Jointly Understand Visual and Tactile Signals” *International Conference on Learning Representations (ICLR)*, 2024.
- C57** Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B. Tenenbaum, Tianmin Shu, Chuang Gan. “Building Cooperative Embodied Agents Modularly with Large Language Models” *International Conference on Learning Representations (ICLR)*, 2024.
- C56** Qinhong Zhou, Sunli Chen, Yisong Wang, Haozhe Xu, Weihua Du, Hongxin Zhang, Yilun Du, Joshua B. Tenenbaum, Chuang Gan. “BHAZARD Challenge: Embodied Decision Making in Dynamically Changing Environments” *International Conference on Learning Representations (ICLR)*, 2024.
- C55** Anurag Ajay*, Seungwook Han*, Yilun Du*, Shuang Li, Abhi Gupta, Tommi Jaakkola, Joshua B. Tenenbaum, Leslie Kaelbling, Akash Srivastava, Pulkit Agrawal. “Compositional Foundation Models for Hierarchical Planning” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C54** Yilun Du*, Mengjiao Yang*, Bo Dai, Hanjun Dai, Ofir Nachum, Joshua B. Tenenbaum, Dale Schuurmans, Pieter Abbeel. “Learning Universal Policies via Text-Guided Video Generation” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C53** Siyuan Zhou, Yilun Du, Shun Zhang, Mengdi Xu, Yikang Shen, Wei Xiao, Dit-Yan Yeung, Chuang Gan. “Adaptive Online Replanning with Diffusion Models” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C52** Cameron Smith, Yilun Du, Ayush Tewari, Vincent Sitzmann. “FlowCam: Training Generalizable 3D Radiance Fields without Camera Poses via Pixel-Aligned Scene Flow” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C51** Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, Chuang Gan. “3D-LLM: Injecting the 3D World into Large Language Models” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C50** Tsun-Hsuan Wang, Juntian Zheng, Pingchuan Ma, Yilun Du, Byungchul Kim, Andrew Spielberg, Joshua B. Tenenbaum, Chuang Gan, Daniela Rus. “DiffuseBot: Breeding Soft Robots With Physics-Augmented Generative Diffusion Models” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C49** Shengzhuang Chen, Long-Kai Huang, Jonathan Richard Schwarz, Yilun Du, Ying Wei. “Secure Out-of-Distribution Task Generalization with Energy-Based Models” *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- C48** Zhutian Yang, Jiayuan Mao, Yilun Du, Jiajun Wu, Joshua B. Tenenbaum, Tomás Lozano-Pérez, Leslie Kaelbling. “Compositional Diffusion-Based Continuous Constraint Solvers” *Conference on Robot Learning (CoRL)*, 2023.
- C47** Nan Liu*, Yilun Du*, Shuang Li*, Joshua B. Tenenbaum, Antonio Torralba. “Unsupervised Compositional Concepts Discovery with Text-to-Image Generative Models” *International Conference on Computer Vision (ICCV)*, 2023.

- C46 Yilun Du, Conor Durkan, Robin Strudel, Joshua B. Tenenbaum, Sander Dieleman, Rob Fergus, Jascha Sohl-Dickstein, Arnaud Doucet, Will Grathwohl. “Reduce, Reuse, Recycle: Compositional Generation with Energy-Based Diffusion Models and MCMC” *International Conference on Machine Learning (ICML)*, 2023.
- C45 Armand Comas, Yilun Du, Christian Fernandez, Sandesh Ghimire, Mario Sznaiier, Joshua B. Tenenbaum, Octavia Camps. “Inferring Relational Potentials in Interacting Systems” *International Conference on Machine Learning (ICML)*, 2023.
- C44 Jiahui Fu, Yilun Du, Kurran Singh, Joshua B. Tenenbaum, John J. Leonard. “NeuSE: Neural SE(3)-Equivariant Embedding for Consistent Spatial Understanding with Objects” *Robotics: Science and Systems (RSS)*, 2023.
- C43 Weiyu Liu, Yilun Du, Tucker Hermans, Sonia Chernova, Chris Paxton. “StructDiffusion: Language-Guided Creation of Physically-Valid Structures using Unseen Objects” *Robotics: Science and Systems (RSS)*, 2023.
- C42 Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin Burchfiel, Shuran Song. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion” *Robotics: Science and Systems (RSS)*, 2023.
- C41 Yilun Du, Cameron Smith, Ayush Tewari, Vincent Sitzmann. “Learning to Render Novel Views from Wide-Baseline Stereo Pairs” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- C40 Yining Hong, Chunru Lin, Yilun Du, Zhenfang Chen, Joshua B. Tenenbaum, Chuang Gan. “3D Concept Learning and Reasoning from Multi-View Images” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- C39 Anurag Ajay*, Yilun Du*, Abhi Gupta*, Joshua B. Tenenbaum, Tommi S. Jaakkola, Pulkit Agrawal. “Is Conditional Generative Modeling all You Need for Decision-Making?” *International Conference on Learning Representations (ICLR)*, 2023.
- C38 Shuang Li*, Yilun Du*, Joshua B. Tenenbaum, Antonio Torralba, Igor Mordatch. “Composing Ensembles of Pre-trained Models via Iterative Consensus” *International Conference on Learning Representations (ICLR)*, 2023.
- C37 Hongyi Chen*, Yilun Du*, Yiye Chen*, Joshua B. Tenenbaum, Patricio Antonio Vela. “Planning with Sequence Models through Iterative Energy Minimization” *International Conference on Learning Representations (ICLR)*, 2023.
- C36 Prafull Sharma, Ayush Tewari, Yilun Du, Sergey Zakharov, Rares Ambrus, Adrien Gaidon, William T. Freeman, Fredo Durand, Joshua B. Tenenbaum, Vincent Sitzmann. “Seeing 3D Objects in a Single Image via Self-Supervised Static-Dynamic Disentanglement” *International Conference on Learning Representations (ICLR)*, 2023.
- C35 Ethan Chun, Yilun Du, Anthony Simeonov, Tomás Lozano-Pérez, Leslie Kaelbling. “Local Neural Descriptor Fields: Locally Conditioned Object Representations for Manipulation” *International Conference on Robotics and Automation (ICRA)*, 2023.
- C34 Jose Iturralde*, Aiden Curtis*, Yilun Du, Leslie Kaelbling, Tomás Lozano-Pérez. “Visibility-Aware Navigation Among Movable Objects” *International Conference on Robotics and Automation (ICRA)*, 2023.
- C33 Anthony Simeonov*, Yilun Du*, Yen-Chen Lin, Alberto Rodriguez, Leslie Kaelbling, Tomás Lozano-Pérez, Pulkit Agrawal. “SE(3)-Equivariant Relational Rearrangement with Neural Descriptor Fields” *Conference on Robot Learning (CoRL)*, 2022.
- C32 Yen-Chen Lin, Pete Florence, Andy Zeng, Jonathan T. Barron, Yilun Du, Wei-Chiu Ma, Anthony Simeonov, Alberto Rodriguez, Phillip Isola. “MIRA: Mental Imagery for Robotic Affordances” *Conference on Robot Learning (CoRL)*, 2022.

- C31** Andrew Luo, Yilun Du, Michael J. Tarr, Joshua B. Tenenbaum, Antonio Torralba, Chuang Gan. “Learning Neural Acoustic Fields” *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- C30** Yining Hong, Yilun Du, Chunru Lin, Joshua B. Tenenbaum, Chuang Gan. “3D Concept Grounding on Neural Fields” *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- C29** Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyurek, Anima Anandkumar, Jacob Andreas, Igor Mordatch, Antonio Torralba, Yuke Zhu. “Pre-Trained Language Models for Interactive Decision-Making” *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- C28** Nan Liu*, Shuang Li*, Yilun Du*, Antonio Torralba, Joshua B. Tenenbaum. “Compositional Visual Generation with Composable Diffusion Models” *European Conference on Computer Vision (ECCV)*, 2022.
- C27** Yilun Du, Tomás Lozano-Pérez, Leslie Kaelbling. “Learning Object Based State Estimators for Household Autonomy” *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- C26** Jiahui Fu, Yilun Du, Kurran Singh, Joshua B. Tenenbaum, John J. Leonard. “Robust Change Detection Based on Neural Descriptor Fields” *International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- C25** Michael Janner*, Yilun Du*, Joshua B. Tenenbaum, Sergey Levine. “Planning with Diffusion for Flexible Behavior Synthesis” *International Conference on Machine Learning (ICML)*, 2022.
- C24** Yilun Du, Shuang Li, Joshua B. Tenenbaum, Igor Mordatch. “Learning Iterative Reasoning through Energy Minimization” *International Conference on Machine Learning (ICML)*, 2022.
- C23** Rylan Schaeffer, Yilun Du, Gabrielle Liu, Ila Fiete. “Streaming Inference for Infinite Feature Models” *International Conference on Machine Learning (ICML)*, 2022.
- C22** Shuang Li, Yilun Du, Guido M. van de Ven, Antonio Torralba, Igor Mordatch. “Energy-Based Models for Continual Learning” *Conference on Lifelong Learning Agents (CoLLA)*, 2022.
- C21** Rylan Schaeffer, Gabrielle Liu, Yilun Du, Scott Linderman, Ila Fiete. “Streaming Inference for Infinite Non-Stationary Clustering” *Conference on Lifelong Learning Agents (CoLLA)*, 2022.
- C20** Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J. Fleet, Dan Gnanaprasasam, Florian Golemo, Charles Herrmann, Thomas Kipf, Abhijit Kundu, Dmitry Lagun, Issam Laradji, Derek Liu, Henning Meyer, Yishu Miao, Derek Nowrouzezahrai, Cengiz Oztireli, Etienne Pot, Noha Radwan, Daniel Rebain, Sara Sabour, Mehdi Sajjadi, Matan Sela, Vincent Sitzmann, Austin Stone, Deqing Sun, Suhani Vora, Ziyu Wang, Tianhao Wu, Kwang Moo Yi, Fangcheng Zhong, Andrea Tagliasacchi. “Kubric: A Scalable Dataset Generator” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- C19** Anthony Simeonov*, Yilun Du*, Andrea Tagliasacchi, Joshua B. Tenenbaum, Alberto Rodriguez, Pulkit Agrawal, Vincent Sitzmann. “Neural Descriptor Fields: SE(3)-Equivariant Object Representations for Manipulation” *International Conference on Robotics and Automation (ICRA)*, 2022.
- C18** Yilun Du, Shuang Li, Yash Sharma, Joshua B. Tenenbaum, Igor Mordatch. “Unsupervised Learning of Compositional Energy Concepts” *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- C17** Nan Liu*, Shuang Li*, Yilun Du*, Joshua B. Tenenbaum, Antonio Torralba. “Learning to Compose Visual Relations” *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- C16** Yilun Du, Katie Collins, Joshua B. Tenenbaum, Vincent Sitzmann. “Learning Signal-Agnostic Manifolds of Neural Fields” *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

- C15 Joseph Suarez, Yilun Du, Clare Zhu, Igor Mordatch, Phillip Isola. “The Neural MMO Platform for Massively Multiagent Research” *Advances in Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS)*, 2021.
- C14 Linqi Zhou, Yilun Du, Jiajun Wu. “3D Shape Generation and Completion through Point-Voxel Diffusion” *International Conference on Computer Vision (ICCV)*, 2021.
- C13 Yilun Du, Yinan Zhang, Hong-Xing Yu, Joshua B. Tenenbaum, Jiajun Wu. “Neural Radiance Flow for 4D View Synthesis and Video Processing” *International Conference on Computer Vision (ICCV)*, 2021.
- C12 Yilun Du, Chuang Gan, Phillip Isola. “Curious Representation Learning for Embodied Intelligence” *International Conference on Computer Vision (ICCV)*, 2021.
- C11 Shuang Li, Yilun Du, Antonio Torralba, Josef Sivic, Bryan Russell. “Weakly Supervised Human-Object Interaction Detection in Video via Contrastive Spatiotemporal Regions” *International Conference on Computer Vision (ICCV)*, 2021.
- C10 Yilun Du, Shuang Li, Joshua B. Tenenbaum, Igor Mordatch. “Improved Contrastive Divergence Training of Energy Based Models” *International Conference on Machine Learning (ICML)*, 2021.
- C9 Yilun Du, Kevin Smith, Tomer Ullman, Joshua B. Tenenbaum, Jiajun Wu. “Unsupervised Discovery of 3D Physical Objects from Video” *International Conference on Learning Representations (ICLR)*, 2021.
- C8 Yilun Du, Shuang Li, Igor Mordatch. “Compositional Visual Generation with Energy Based Models” *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- C7 Anthony Simeonov, Yilun Du, Beomjoon Kim, Francois Hogan, Joshua B. Tenenbaum, Pulkit Agrawal, Alberto Rodriguez. “A Long Horizon Planning Framework for Manipulating Rigid Pointcloud Objects” *Conference on Robot Learning (CoRL)*, 2020.
- C6 Yilun Du, Joshua Meier, Jerry Ma, Rob Fergus, Alexander Rives. “Energy-based models for atomic-resolution protein conformations” *International Conference on Learning Representations (ICLR)*, 2020.
- C5 Xingyou Song, Yiding Jiang, Yilun Du, Behnam Neyshabur. “Observational Overfitting in Reinforcement Learning” *International Conference on Learning Representations (ICLR)*, 2020.
- C4 Yilun Du, Toru Lin, Igor Mordatch. “Model Based Planning with Energy Based Models” *Conference on Robot Learning (CoRL)*, 2019.
- C3 Yilun Du, Igor Mordatch. “Implicit Generation and Generalization with Energy Based Models” *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- C2 Yilun Du, Karthik Narasimhan. “Task-Agnostic Dynamics Priors for Deep Reinforcement Learning” *International Conference on Machine Learning (ICML)*, 2019.
- C1 Yilun Du, Zhijian Liu, Hector Basevi, Ales Leonardis, William T. Freeman, Joshua B. Tenenbaum, Jiajun Wu. “Learning to Exploit Stability for 3D Scene Parsing” *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Journal

- J10 Han Qi, Haocheng Yin, Aris Zhu, Yilun Du, Heng Yang. “Strengthening Generative Robot Policies through Predictive World Modeling” *IEEE Robotics and Automation Letters (RA-L)*, 2026.
- J9 Chaoqi Liu, Haonan Chen, Sigmund H. Høeg, Shaoxiong Yao, Yunzhu Li, Kris Hauser, Yilun Du. “Flexible Multitask Learning with Factorized Diffusion Policy” *IEEE Robotics and Automation Letters (RA-L)*, 2026.

- J8** Sigmund H. Høeg, Aksel Vaaler, Chaoqi Liu, Olav Egeland, [Yilun Du](#). “Hybrid Diffusion for Simultaneous Symbolic and Continuous Planning” *IEEE Robotics and Automation Letters (RA-L)*, 2026.
- J7** Robert McCarthy, Daniel C.H. Tan, Dominik Schmidt, Fernando Acero, Nathan Herr, [Yilun Du](#), Thomas Thuruthel, Zhibin Li. “Towards Generalist Robot Learning from Internet Video: A Survey” *Journal of Artificial Intelligence Research (JAIR)*, 2025.
- J6** Julen Urain, Ajay Mandlekar, [Yilun Du](#), Mahi Shafiullah, Danfei Xu, Katerina Fragkiadaki, Georgia Chalvatzaki, Jan Peters. “Deep Generative Models in Robotics: A Survey on Learning from Multimodal Demonstrations” *IEEE Transactions on Robotics (T-RO)*, 2025.
- J5** Yiqing Xu, Jiayuan Mao, Linfeng Li, [Yilun Du](#), Tomas Lozano-Pérez, Leslie Pack Kaelbling, David Hsu. “"Set It Up": Functional Object Arrangement with Compositional Generative Models” *International Journal on Robotics Research (IJRR)* 2025
- J4** Robert McCarthy, Daniel CH Tan, Dominik Schmidt, Fernando Acero, Nathan Herr, [Yilun Du](#), Thomas G Thuruthel, Zhibin Li. “Towards generalist robot learning from internet video: A survey” *Journal on Artificial Intelligence Research (IJRR)* 2025
- J3** Julen Urain, Ajay Mandlekar, [Yilun Du](#), Mahi Shafiullah, Danfei Xu, Katerina Fragkiadaki, Georgia Chalvatzaki, Jan Peters. “Deep generative models in robotics: A survey on learning from multimodal demonstrations” *IEEE Transactions on Robotics (TRO)* 2025
- J2** Jiahui Fu, [Yilun Du](#), Kurran Singh, Joshua B. Tenenbaum, John J. Leonard. “NeuSE: Neural SE(3)-Equivariant Embedding for Consistent Spatial Understanding with Objects” *International Journal on Robotics Research (IJRR)* 2025
- J1** Cheng Chi*, Zhenjia Xu*, Siyuan Feng, Eric Cousineau, [Yilun Du](#), Benjamin Burchfiel, Russ Tedrake, Shuran Song. “Diffusion Policy: Visuomotor Policy Learning via Action Diffusion” *International Journal on Robotics Research (IJRR)* 2024

Technical Reports

- P5** Gemini Team. “Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities” *arXiv preprint arXiv:2507.06261*, 2025. [Blog Paper](#).
- P4** Gemini Robotics Team. “Gemini Robotics: Bringing AI into the Physical World” *arXiv preprint arXiv:2503.20020*, 2025. [Blog Paper](#).
- P3** Gemini Robotics Team. “Gemini Robotics 1.5: Pushing the Frontier of Generalist Robots with Advanced Embodied Reasoning, Thinking, and Motion Transfer” *arXiv preprint arXiv:2510.03342*, 2025. [Blog Paper](#).
- P2** Gemini Robotics Team. “Evaluating Gemini Robotics Policies in a Veo World Simulator” *arXiv preprint arXiv:2512.10675*, 2025. [Website Paper](#).
- P1** Robert Verkuil*, Ori Kabeli*, [Yilun Du](#), Basile Wicky, Lukas Milles, Justas Dauparas, David Baker, Sergey Ovchinnikov, Tom Sercu, Alexander Rives. “Language Models Generalize Beyond Natural Proteins” *bioRxiv preprint*, 2022. [Paper](#).

Theses

- T1** [Yilun Du](#) 2024. “Learning Generalizable Systems by Learning Composable Energy Landscapes”. PhD Thesis, Massachusetts Institute of Technology.

Teaching

Experience

6.867 Machine Learning (Graduate) Teaching Assistant, Massachusetts Institute of Technology	Fall 2021
6.036/6.862 Machine Learning (Undergraduate/Graduate) Teaching Assistant, Massachusetts Institute of Technology	Spring 2018
6.036 Machine Learning (Undergraduate) Teaching Assistant, Massachusetts Institute of Technology	Fall 2017

Invited Talks

RSS Workshop on Diffusion Models for Robot Learning Model Based Decision Making with Diffusion Models, July 2026	
RSS Workshop on Foundation Models and Planning Planning with World Models, July 2026	
CVPR Workshop on Embodied Reasoning in Action Embodied Reasoning with World Models, June 2026	
CVPR Workshop on Embodied Intelligence Robot Manipulation through World Models, June 2026	
CVPR Workshop on Generative Models in Computer Vision Scene Understanding through Inverse Generative Modeling, June 2026	
CVPR Workshop on Visual Concepts Compositional Energy-Based Concepts, June 2026	
CVPR Workshop on GeoFreeNVS: Geometry-Free Novel View Synthesis and Controllable Video Models A Unified Generative-Predictive Modeling for 4D Scene Understanding, June 2026	
CVPR Workshop on Emerging Directions in Data for Multimodal Foundation Models A Very Big Video Reasoning Suite, June 2026	
CVPR Workshop on Video World Models Robot Manipulation with World Models, June 2026	
CVPR Workshop on From Labs to Life: Embodied Intelligence in the Wild Embodied Intelligence with World Models, June 2026	
Brown University Workshop on Agentic Scientific Computing and Scientific Machine Learning Multiagent Scientific Reasoning Systems, May 2026	

Oregon State University

Building Embodied Intelligence with Compositional World Models, January 2026

NeurIPS Workshop on Embodied World Models

Building Intelligent Robots with World Models, December 2025

NeurIPS Tutorial on Beyond Autoregressive Models

Autoregressive World Models, December 2025

MIT FutureTech

AI Automation with Generative World Models, November 2025

Amazon Robotics

Robot Manipulation with World Models, November 2025

Tufts ECE Seminar

Learning Compositional Models of the World, September 2025

Nvidia

Multiagent Reasoning with Language Models, August 2025

Toyota Research Institute ML and Robotics Seminar

Flexible Decision Making with Compositional Constraints, August 2025

Meta MRS AI Seminar

Multiagent Reasoning with Language Models, July 2025

Microsoft Research Seminar of Sampling and Generative Models

Inference Time Reasoning with Generative Models, July 2025

ICML Workshop on Multiagent Systems

Multiagent Language Model Reasoning Systems, July 2025

CVPR Tutorial on Continual Robot Learning from Humans

Continual Robot Learning with Compositional Generative Models, June 2025

RSS Workshop on Leveraging Foundation Models for Interactive Robot learning

Scaling Robot Learning with Video Models, June 2025

CVPR Tutorial on Robotics 101: An Odyssey from A Vision Perspective

Learning Compositional Models of the World, June 2025

CVPR Workshop on 3D-LLM/VLA

Build 4D Models of the World, June 2025

CVPR Workshop on Foundation Models Meet Embodied Agents

Learning Compositional Models of the World, June 2025

ICRA Workshop on Foundation Models and Neuro-Symbolic AI for Robotics

Generalizable Foundation Models through Compositional Generation, May 2025

New England NLP Symposium

Constructing Generalizable Multimodal Models through Compositional Generation, April 2025

Weizmann AI and Vision Seminar

Learning Compositional Models of the World, March 2025

Shanghai AI Lab

Learning Compositional Models of the World, Nov 2024

ICML Workshop on Multimodal Foundation Models in Embodied AI

Compositional Foundation Models, July 2024

RSS Workshop on Semantic Reasoning and Goal Understanding in Robotics

Panelist, July 2024

RSS Workshop on GenAI and HRI

Build Customizable Generative Models through Compositional Generation, July 2024

Generalizing Outside the Training Distribution through Composition Generation

UCLA, November 2024

KAIST, October 2024

Vanderbilt Machine Learning, September 2024

NUS Tea Time, July 2024

CMU MLD, April 2024

UMich CSE, April 2024

UT Austin, April 2024

UBC, April 2024

UT Austin, April 2024

Princeton, April 2024

Brown, April 2024

NYU, March 2024

Columbia, March 2024

UPenn, March 2024

University of Washington, March 2024

Georgia Tech, March 2024

University of Toronto, March 2024

Harvard, March 2024

UCSD, Feb 2024

Boston Dynamics AI Institute, Feb 2024

ASU Frontiers in GenAI Seminar, Jan 2024

Genesis Therapeutics, Jan 2024

JHU, Nov 2023

KAIST, Nov 2023

CMU, AI Seminar, Nov 2023

Brown University, Robotics Seminar, Oct 2023

Cornell University, Robotics Seminar, Oct 2023

UC Berkeley, Oct 2023

UT Austin, Oct 2023

Stanford, Vision Seminar, July 2023

NeurIPS Tutorial on Language Models Meet World Models

Panelist, Dec 2023

Learning a Visual Imagination

BuzzRobot, Dec 2023

Towards Multi-AI Intelligent Systems

University of Tokyo, IRCN Keynote Speech, Oct 2023

Eye on AI Podcast

AI Self-Improvement from Multiagent Debate, Oct 2023

Riding the Wave of Generative Modeling for Decision Making

University of Montreal, Robot Learning Seminar, Oct 2023

Planning with Diffusion Models

Donders Institute, Feb 2023

Deepmind, Aug 2022

Implicit Learning with Energy Based Models

Nuro, October 2022

Tencent, September 2022

CMU, August 2022

Generally Intelligent Podcast

Modularity and Energy Based Models, Nov 2021

Professional Service

Conference Area Chair

Robotics Science and Systems (RSS)	2026
Neural Information Processing Systems (NeurIPS)	2026
International Conference in Machine Learning (ICML)	2026
International Conference in Learning Representations (ICLR)	2025-26
Conference on Robot Learning (CoRL)	2025-26
Conference on Language Modeling (CoLM)	2025-26
International Conference in Computer Vision (ICCV)	2025

Conference Refereeing

Neural Information Processing Systems (NeurIPS)	2021-25
International Conference in Machine Learning (ICML)	2021-25
International Conference in Learning Representations (ICLR)	2021-24
Computer Vision and Pattern Recognition (CVPR)	2021-23
International Conference in Computer Vision (ICCV)	2021,23
European Conference in Computer Vision (ECCV)	2022
Special Interest Group in Graphics (SIGGRAPH)	2022-24
SIGGRAPH Asia	2024
Conference on Robot Learning (CoRL)	2023-24
Robotics: Science and Systems (RSS)	2023
International Conference on Robotics and Automation (ICRA)	2022-23
International Conference on Intelligent Robots and Systems (IROS)	2022-23

Workshop Organization

Workshop on Memory in Foundation Models at CoRL	December 2026
Workshop on Compositional Learning at CoRL	December 2026
Workshop on Embodied Spatial Reasoning at NeurIPS	December 2026
Workshop on Compositional Learning in the Era of Scaling in Robotics at IROS	October 2026
Workshop on Tactile Sensing for Robotic Foundation Models at RSS	July 2026
Workshop on Compositional Learning at ICML 2026	July 2026
Workshop on 3D-LLM/VLA: Bridging Language, Vision and Action in 3D Environments at CVPR 2026	June 2026
Workshop on World Models Meet Active Sensing and Closed-Loop Planning at CVPR 2026	June 2026
Workshop on Multi-Agent Robotic Systems: Scaling with Compositional Intelligence at CVPR 2026	June 2026
Workshop on Robotics World Modeling at CoRL 2025	September 2025
Workshop on Robot Planning in the Era of Foundation Models at RSS 2025	June 2025
Workshop on World Models in Physical World at ICML 2025	June 2025
Workshop on Multi-Agent AI in the Real World at AAAI 2025	March 2025
Workshop on Compositional Learning at NeurIPS 2024	Dec 2024
Workshop on Safe Generative AI at NeurIPS 2024	Dec 2024
Workshop on Geometric and Algebraic Structure in Robot Learning at RSS 2024	July 2024
Generative Models for Decision Making Workshop at ICLR 2023	May 2024
Foundation Models for Decision Making Workshop at NeurIPS 2023	Dec 2023
Foundation Models for Decision Making Workshop at NeurIPS 2022	Dec 2022

Seminar Organization

MIT Embodied Intelligence Seminar	2021-2023
-----------------------------------	-----------