

# Yuke Zhu

GDC 3.422, 2317 Speedway,  
Austin, TX, 78712 USA  
<https://cs.utexas.edu/~yukez>  
yukez@cs.utexas.edu

**ACADEMIC EMPLOYMENT**     **Assistant Professor**, Department of Computer Science     2020 – Present  
The University of Texas at Austin, Austin, TX, USA

**EDUCATION**     **Stanford University**, Stanford, CA, USA     2015 – 2019  
*Ph.D.* in Computer Science

**Stanford University**, Stanford, CA, USA     2013 – 2015  
*Master of Science* in Computer Science

**Simon Fraser University**, Vancouver, BC, Canada     2011 – 2013  
*Bachelor of Science* in Computer Science (First Class with Distinction)

**Zhejiang University**, Hangzhou, China     2009 – 2013  
*Bachelor of Engineering* in Computer Science and Technology (Dual Degree Program)

**HONORS & AWARDS**

- CoRL Early-Career Keynote Speaker     2023
- CoRL Best Paper Award (Finalist)     2023
- CoRL Best Systems Paper Award (Finalist)     2023
- JP Morgan Faculty Research Award     2023
- RSS Best Paper Award (Finalist)     2023
- AAAI New Faculty Highlights     2023
- NeurIPS Outstanding Paper Award     2022
- RSS Best Student Paper Award (Finalist)     2022
- IEEE ICRA Outstanding Learning Paper Award     2022
- Amazon Research Award     2022
- NSF CAREER Award     2022
- IEEE ICRA Best Multi-Robotic Systems Paper Award (Finalist)     2021
- Amazon Research Award     2021
- IROS Best Cognitive Robotics Paper Award (Finalist)     2019
- IEEE ICRA Best Conference Paper Award     2019
- IEEE ICRA Best Cognitive Robotics Paper Award (Finalist)     2019
- RSS Pioneers Selected Cohort     2018

## REFEREED PUBLICATIONS

\* Equal contribution. † Equal advising.

[90] Zhenyu Jiang, Hanwen Jiang, **Yuke Zhu**. Doduo: Dense Visual Correspondence from Unsupervised Semantic-Aware Flow. *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

- [89] Huihan Liu, Shivin Dass, Roberto Martín-Martín, **Yuke Zhu**. Model-Based Runtime Monitoring with Interactive Imitation Learning. *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [88] Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [87] Weikang Wan, Yifeng Zhu\*, Rutav Shah\*, **Yuke Zhu**. LOTUS: Continual Imitation Learning for Robot Manipulation Through Unsupervised Skill Discovery. *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [86] Jake Grigsby, Linxi Fan, **Yuke Zhu**. AMAGO: Scalable In-Context Reinforcement Learning for Adaptive Agents. *International Conference on Learning Representations (ICLR)*, 2024. **Spotlight Presentation**.
- [85] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, **Yuke Zhu**, Linxi Fan†, Anima Anandkumar†. Eureka: Human-Level Reward Design via Coding Large Language Models. *International Conference on Learning Representations (ICLR)*, 2024.
- [84] Hanwen Jiang, Zhenyu Jiang, Kristen Grauman, **Yuke Zhu**. Few-View Object Reconstruction with Unknown Categories and Camera Poses. *International Conference on 3D Vision (3DV)*, 2024. **Oral Presentation**.
- [83] Zizhao Wang\*, Caroline Wang\*, Xuesu Xiao, **Yuke Zhu**, Peter Stone. Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning. *AAAI Conference on Artificial Intelligence (AAAI)*, 2024. **Oral Presentation**.
- [82] Mingyo Seo, Steve Han, Kyutae Sim, Seung Hyeon Bang, Carlos Gonzalez, Luis Sentis, **Yuke Zhu**. Deep Imitation Learning for Humanoid Loco-manipulation through Human Teleoperation. *International Conference on Humanoid Robots (Humanoids)*, 2023. **Oral Presentation**.
- [81] Lucy Xiaoyang Shi\*, Yunfan Jiang\*, Jake Grigsby, Linxi Fan†, **Yuke Zhu**†. Cross-Episodic Curriculum for Transformer Agents. *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- [80] Bo Liu\*, Yifeng Zhu\*, Chongkai Gao\*, Yihao Feng, Qiang Liu, **Yuke Zhu**, Peter Stone. LIBERO: Benchmarking Knowledge Transfer in Lifelong Robot Learning. *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- [79] Zhuolin Yang, Wei Ping, Zihan Liu, Vijay Korthikanti, Weili Nie, De-An Huang, Linxi Fan, Zhiding Yu, Shiyi Lan, Bo Li, Ming-Yu Liu, **Yuke Zhu**, Mohammad Shoeybi, Bryan Catanzaro, Chaowei Xiao, Anima Anandkumar. Re-ViLM: Retrieval-Augmented Visual Language Model for Zero and Few-Shot Image Captioning. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- [78] Chen Wang, Linxi Fan, Jiankai Sun, Ruohan Zhang, Li Fei-Fei, Danfei Xu, **Yuke Zhu**†, Anima Anandkumar†. MimicPlay: Long-Horizon Imitation Learning by Watching Human Play. *Conference on Robot Learning (CoRL)*, 2023. **Best Paper Award Finalist. Best Systems Paper Award Finalist**.
- [77] Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Ireteyayo Akinola, Yashraj Narang, Linxi Fan, **Yuke Zhu**, Dieter Fox. MimicGen: A Data Generation System for Scalable Robot Learning using Human Demonstrations. *Conference on Robot Learning (CoRL)*, 2023.
- [76] Rutav Shah, Roberto Martín-Martín†, **Yuke Zhu**†. MUTEX: Learning Unified Policies from Multimodal Task Specifications. *Conference on Robot Learning (CoRL)*, 2023.

- [75] Yifeng Zhu, Zhenyu Jiang, Peter Stone, **Yuke Zhu**. Learning Generalizable Manipulation Policies with Object-Centric 3D Representations. *Conference on Robot Learning (CoRL)*, 2023.
- [74] Xiaohan Zhang, Yifeng Zhu, Yan Ding, Yuqian Jiang, **Yuke Zhu**, Peter Stone, Shiqi Zhang. Symbolic State Space Optimization for Long Horizon Mobile Manipulation Planning. *International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [73] Bokui Shen\*, Zhenyu Jiang\*, Christopher Choy, Silvio Savarese, Leonidas J. Guibas, Anima Anandkumar, **Yuke Zhu**. ACID: Action-Conditional Implicit Visual Dynamics for Deformable Object Manipulation. *International Journal of Robotics Research (IJRR)*, 2023.
- [72] Yunfan Jiang, Agrim Gupta\*, Zichen Zhang\*, Guanzhi Wang\*, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, **Yuke Zhu**<sup>†</sup>, Linxi Fan<sup>†</sup>. VIMA: General Robot Manipulation with Multimodal Prompts. *International Conference on Machine Learning (ICML)*, 2023.
- [71] Huihan Liu, Soroush Nasiriany, Lance Zhang, Zhiyao Bao, **Yuke Zhu**. Robot Learning on the Job: Human-in-the-Loop Autonomy and Learning During Deployment. *Robotics: Science and Systems (RSS)*, 2023. **Best Paper Award Finalist**.
- [70] Wei Dong, Chris Choy, Charles Loop, Or Litany, **Yuke Zhu**, Anima Anandkumar. Fast Monocular Scene Reconstruction with Global-Sparse Local-Dense Grids. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [69] Mingyo Seo, Ryan Gupta, Yifeng Zhu, Alexy Skoutnev, Luis Sentis, **Yuke Zhu**. Learning to Walk by Steering: Perceptive Quadrupedal Locomotion in Dynamic Environments. *International Conference on Robotics and Automation (ICRA)*, 2023.
- [68] Cheng-Chun Hsu, Zhenyu Jiang, **Yuke Zhu**. Ditto in the House: Building Articulated Models of Indoor Scenes through Interactive Perception. *International Conference on Robotics and Automation (ICRA)*, 2023.
- [67] Yifeng Zhu, Abhishek Joshi, Peter Stone, **Yuke Zhu**. VIOLA: Imitation Learning for Vision-Based Manipulation with Object Proposal Priors. *Conference on Robot Learning (CoRL)*, 2022.
- [66] Soroush Nasiriany, Tian Gao, Ajay Mandlekar, **Yuke Zhu**. Learning and Retrieval from Prior Data for Skill-based Imitation Learning. *Conference on Robot Learning (CoRL)*, 2022.
- [65] Shuang Li, Xavier Puig, Chris Paxton, Yilun Du, Clinton Wang, Linxi Fan, Tao Chen, De-An Huang, Ekin Akyürek, Anima Anandkumar, Jacob Andreas, Igor Mordatch, Antonio Torralba, **Yuke Zhu**. Pre-Trained Language Models for Interactive Decision-Making. *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. **Oral Presentation**.
- [64] Linxi Fan, Guanzhi Wang\*, Yunfan Jiang\*, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, **Yuke Zhu**<sup>†</sup>, Anima Anandkumar<sup>†</sup>. MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge. *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. **Outstanding Paper Award**.
- [63] Zizhao Wang, Xuesu Xiao, Zifan Xu, **Yuke Zhu**, Peter Stone. Causal Dynamics Learning for Task-Independent State Abstraction. *International Conference on Machine Learning (ICML)*, 2022. **Long Talk**.

- [62] Bokui Shen, Zhenyu Jiang, Christopher Choy, Silvio Savarese, Leonidas J. Guibas, Anima Anandkumar, **Yuke Zhu**. ACID: Action-Conditional Implicit Visual Dynamics for Deformable Object Manipulation. *Robotics: Science and Systems (RSS)*, 2022. **Best Student Paper Award Finalist**.
- [61] Zhenyu Jiang, Cheng-Chun Hsu, **Yuke Zhu**. Ditto: Building Digital Twins of Articulated Objects from Interaction. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **Oral Presentation**.
- [60] Huaizu Jiang\*, Xiaojian Ma\*, Weili Nie, Zhiding Yu, **Yuke Zhu**, Anima Anandkumar. Bongard-HOI: Benchmarking Few-Shot Visual Reasoning for Human-Object Interactions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **Oral Presentation**.
- [59] Jiaxun Cui\*, Hang Qiu\*, Dian Chen, Peter Stone, **Yuke Zhu**. COOPERNAUT: End-to-End Driving with Cooperative Perception for Networked Vehicles. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [58] Josiah Wong, Viktor Makoviychuk, Anima Anandkumar, **Yuke Zhu**. OSCAR: Data-Driven Operational Space Control for Adaptive and Robust Robot Manipulation. *International Conference on Robotics and Automation (ICRA)*, 2022.
- [57] Soroush Nasiriany, Huihan Liu, **Yuke Zhu**. Augmenting Reinforcement Learning with Behavior Primitives for Diverse Manipulation Tasks. *International Conference on Robotics and Automation (ICRA)*, 2022. **Outstanding Learning Paper Award**.
- [56] Xiaohan Zhang, Yifeng Zhu, Yan Ding, **Yuke Zhu**, Peter Stone, Shiqi Zhang . Visually Grounded Task and Motion Planning for Mobile Manipulation. *International Conference on Robotics and Automation (ICRA)*, 2022.
- [55] Xiaojian Ma, Weili Nie, Zhiding Yu, Huaizu Jiang, Chaowei Xiao, **Yuke Zhu**, Song-Chun Zhu, Anima Anandkumar. RelViT: Concept-Guided Vision Transformer for Visual Relational Reasoning. *International Conference on Learning Representations (ICLR)*, 2022.
- [54] Yifeng Zhu, Peter Stone, **Yuke Zhu**. Bottom-Up Skill Discovery from Unsegmented Demonstrations for Long-Horizon Robot Manipulation. *IEEE Robotics and Automation Letters (RA-L)*, 2022.
- [53] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, **Yuke Zhu**, Roberto Martín-Martín. What Matters in Learning from Offline Human Demonstrations for Robot Manipulation. *Conference on Robot Learning (CoRL)*, 2021. **Oral Presentation**.
- [52] Youngwoon Lee, Joseph J. Lim, Anima Anandkumar, **Yuke Zhu**. Adversarial Skill Chaining for Long-Horizon Robot Manipulation via Terminal State Regularization. *Conference on Robot Learning (CoRL)*, 2021.
- [51] Paul Pu Liang, Yiwei Lyu, Xiang Fan, Zetian Wu, Yun Cheng, Jason Wu, Leslie Chen, Peter Wu, Michelle A. Lee, **Yuke Zhu**, Ruslan Salakhutdinov, Louis-Philippe Morency. MultiBench: Multiscale Benchmarks for Multimodal Representation Learning. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [50] Shiyi Lan, Zhiding Yu, Christopher Choy, Subhashree Radhakrishnan, Guilin Liu, **Yuke Zhu**, Larry S. Davis, Anima Anandkumar. DiscoBox: Weakly Supervised Instance Segmentation and Semantic Correspondence from Box Supervision. *International Conference on Computer Vision (ICCV)*, 2021.

- [49] Anuj Mahajan, Mikayel Samvelyan, Lei Mao, Viktor Makoviychuk, Animesh Garg, Jean Kossaifi, Shimon Whiteson, **Yuke Zhu**, Anima Anandkumar. Tesseract: Tensorised Actors for Multi-Agent Reinforcement Learning. *International Conference on Machine Learning (ICML)*, 2021.
- [48] Linxi Fan, Guanzhi Wang, De-An Huang, Zhiding Yu, Li Fei-Fei, **Yuke Zhu**, Anima Anandkumar. SECANT: Self-Expert Cloning for Zero-Shot Generalization of Visual Policies. *International Conference on Machine Learning (ICML)*, 2021.
- [47] Bo Liu, Qiang Liu, Peter Stone, Animesh Garg, **Yuke Zhu**, Anima Anandkumar. Coach-Player Multi-Agent Reinforcement Learning for Dynamic Team Composition. *International Conference on Machine Learning (ICML)*, 2021. **Long Talk**.
- [46] Zhenyu Jiang, Yifeng Zhu, Maxwell Svetlik, Kuan Fang, **Yuke Zhu**. Synergies Between Affordance and Geometry: 6-DoF Grasp Detection via Implicit Representations. *Robotics: Science and Systems (RSS)*, 2021.
- [45] Kuan Fang, **Yuke Zhu**, Silvio Savarese, Li Fei-Fei. Learning Generalizable Skills via Automated Generation of Diverse Tasks. *Robotics: Science and Systems (RSS)*, 2021.
- [44] Michelle A. Lee, Matthew Tan, **Yuke Zhu**, Jeannette Bohg. Detect, Reject, Correct: Crossmodal Compensation of Corrupted Sensors. *International Conference on Robotics and Automation (ICRA)*, 2021.
- [43] Danfei Xu, Ajay Mandlekar, Roberto Martn-Martn, **Yuke Zhu**, Silvio Savarese, Li Fei-Fei. Deep Affordance Foresight: Planning Through What Can Be Done in the Future. *International Conference on Robotics and Automation (ICRA)*, 2021.
- [42] Guanya Shi, Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, Fabio Ramos, Anima Anandkumar, **Yuke Zhu**. Fast Uncertainty Quantification for Deep Object Pose Estimation. *International Conference on Robotics and Automation (ICRA)*, 2021.
- [41] Xinlei Pan, Animesh Garg, Anima Anandkumar, **Yuke Zhu**. Emergent Hand Morphology and Control from Optimizing Robust Grasps of Diverse Objects. *International Conference on Robotics and Automation (ICRA)*, 2021.
- [40] Yifeng Zhu, Jonathan Tremblay, Stan Birchfield, **Yuke Zhu**. Hierarchical Planning for Long-Horizon Manipulation with Geometric and Symbolic Scene Graphs. *International Conference on Robotics and Automation (ICRA)*, 2021.
- [39] Albert Tung, Josiah Wong, Ajay Mandlekar, Roberto Martn-Martn, **Yuke Zhu**, Li Fei-Fei, Silvio Savarese. Learning Multi-Arm Manipulation Through Collaborative Teleoperation. *International Conference on Robotics and Automation (ICRA)*, 2021. **Best Multi-Robotic Systems Paper Award Finalist**.
- [38] Kuan Fang, **Yuke Zhu**, Silvio Savarese, Li Fei-Fei. Adaptive Procedural Task Generation for Hard-Exploration Problems. *International Conference on Learning Representations (ICLR)*, 2021.
- [37] Weili Nie, Zhiding Yu, Lei Mao, Ankit B. Patel, **Yuke Zhu**, Anima Anandkumar. Bongard-LOGO: A New Benchmark for Human-Level Concept Learning and Reasoning. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. **Spotlight Presentation**.

- [36] Xingye Da, Zhaoming Xie, David Hoeller, Byron Boots, Anima Anandkumar, **Yuke Zhu**, Buck Babich, Animesh Garg. Learning a Contact-Adaptive Controller for Robust, Efficient Legged Locomotion. *Conference on Robot Learning (CoRL)*, 2020.
- [35] Linxi Fan\*, Shyamal Buch\*, Guanzhi Wang, Ryan Cao, **Yuke Zhu**, Juan Carlos Niebles, Li Fei-Fei. RubiksNet: Learnable 3D-Shift for Efficient Video Action Recognition. *European Conference on Computer Vision (ECCV)*, 2020.
- [34] Hongyu Ren, **Yuke Zhu**, Jure Leskovec, Anima Anandkumar, Animesh Garg. OCEAN: Online Task Inference for Compositional Tasks with Context Adaptation. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020.
- [33] Yunbo Wang\*, Bo Liu\*, Jiajun Wu, **Yuke Zhu**, Simon S. Du, Li Fei-Fei, Joshua B Tenenbaum. Dual Sequential Monte Carlo: Tunneling Filtering and Planning in Continuous POMDPs. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [32] Chen Wang, Roberto Martn-Martn, Danfei Xu, Jun Lv, Cewu Lu, Li Fei-Fei, Silvio Savarese, **Yuke Zhu**. 6-PACK: Category-Level 6D Pose Tracker with Anchor-Based Keypoints. *International Conference on Robotics and Automation (ICRA)*, 2020.
- [31] Zengyi Qin, Kuan Fang, **Yuke Zhu**, Li Fei-Fei, Silvio Savarese. KETO: Learning Keypoint Representations for Tool Manipulation. *International Conference on Robotics and Automation (ICRA)*, 2020.
- [30] Michelle Lee, **Yuke Zhu**, Peter Zachares, Matthew Tan, Krishnan Srinivasan, Silvio Savarese, Li Fei-Fei, Animesh Garg, Jeannette Bohg. Making Sense of Vision and Touch: Learning Multimodal Representations for Contact-Rich Tasks. *IEEE Transactions on Robotics (T-RO)*, 2020.
- [29] Suraj Nair, **Yuke Zhu**, Silvio Savarese, Li Fei-Fei. Causal Induction from Visual Observations for Goal Directed Tasks. *NeurIPS 2019 Workshop on Causal Machine Learning*, 2019.
- [28] Danfei Xu, Roberto Martn-Martn, De-An Huang, **Yuke Zhu**, Silvio Savarese, Li Fei-Fei. Regression Planning Networks. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [27] Kuan Fang, **Yuke Zhu**, Animesh Garg, Silvio Savarese, Li Fei-Fei. Dynamics Learning with Cascaded Variational Inference for Multi-Step Manipulation. *Conference on Robot Learning (CoRL)*, 2019. **Oral Presentation.**
- [26] Kuan Fang, **Yuke Zhu**, Animesh Garg, Andrey Kurenkov, Viraj Mehta, Li Fei-Fei, Silvio Savarese. Learning Task-Oriented Grasping for Tool Manipulation from Simulated Self-Supervision. *International Journal of Robotics Research (IJRR)*, 2019.
- [25] William B. Shen, Danfei Xu, **Yuke Zhu**, Li Fei-Fei, Leonidas Guibas, Silvio Savarese. Situational Fusion of Visual Representation for Visual Navigation. *International Conference on Computer Vision (ICCV)*, 2019.
- [24] Ajay Mandlekar, Jonathan Booher, Max Spero, Albert Tung, Anchit Gupta, **Yuke Zhu**, Animesh Garg, Silvio Savarese, Li Fei-Fei. Scaling Robot Supervision to Hundreds of Hours with RoboTurk: Robotic Manipulation Dataset through Human Reasoning and Dexterity. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019. **Best Cognitive Robotics Paper Award Finalist.**

- [23] De-An Huang, Danfei Xu, **Yuke Zhu**, Animesh Garg, Silvio Savarese, Li Fei-Fei, Juan Carlos Niebles. Continuous Relaxation of Symbolic Planner for One-Shot Imitation Learning. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [22] Michelle A. Lee\*, **Yuke Zhu**\*, Krishnan Srinivasan, Parth Shah, Silvio Savarese, Li Fei-Fei, Animesh Garg, Jeannette Bohg. Making Sense of Vision and Touch: Self-Supervised Learning of Multimodal Representations for Contact-Rich Tasks. *International Conference on Robotics and Automation (ICRA)*, 2019. **IEEE ICRA Best Conference Paper Award**.
- [21] Chen Wang, Danfei Xu, **Yuke Zhu**, Roberto Martn-Martn, Cewu Lu, Li Fei-Fei, Silvio Savarese. DenseFusion: 6D Object Pose Estimation by Iterative Dense Fusion. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [21] De-An Huang\*, Suraj Nair\*, Danfei Xu\*, **Yuke Zhu**, Animesh Garg, Li Fei-Fei, Silvio Savarese, Juan Carlos Niebles. Neural Task Graphs: Generalizing to Unseen Tasks from a Single Video Demonstration. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. **Oral Presentation**.
- [20] Linxi Fan\*, **Yuke Zhu**\*, Jiren Zhu, Zihua Liu, Orien Zeng, Anchit Gupta, Joan Creus-Costa, Silvio Savarese, Li Fei-Fei. SURREAL: Open-Source Reinforcement Learning Framework and Robot Manipulation Benchmark. *Conferences on Robot Learning (CoRL)*, 2018.
- [19] Ajay Mandlekar, **Yuke Zhu**, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao, John Emmons, Anchit Gupta, Emre Orbay, Silvio Savarese, Li Fei-Fei. RoboTurk: A Crowdsourcing Platform for Robotic Skill Learning through Imitation. *Conferences on Robot Learning (CoRL)*, 2018.
- [18] **Yuke Zhu**, Ziyu Wang, Josh Merel, Andrei Rusu, Tom Erez, Serkan Cabi, Saran Tunyasuvunakool, Jnos Kramr, Raia Hadsell, Nando de Freitas, Nicolas Heess. Reinforcement and Imitation Learning for Diverse Visuomotor Skills. *Robotics: Science and Systems (RSS)*, 2018.
- [17] Kuan Fang, **Yuke Zhu**, Animesh Garg, Virja Mehta, Andrey Kuryenkov, Li Fei-Fei, Silvio Savarese. Learning Task-Oriented Grasping for Tool Manipulation with Simulated Self-Supervision. *Robotics: Science and Systems (RSS)*, 2018
- [16] Danfei Xu\*, Suraj Nair\*, **Yuke Zhu**, Julian Gao, Animesh Garg, Li Fei-Fei, Silvio Savarese. Neural Task Programming: Learning to Generalize Across Hierarchical Tasks. *International Conference on Robotics and Automation (ICRA)*, 2018.
- [15] Bo Wang, Lin Huang, **Yuke Zhu**, Anshul Kundaje, Serafim Batzoglou, Anna Goldenberg. Vicus: Exploiting Local Structures to Improve Network-based Analysis of Biological Data. *PLOS Computational Biology*, 2017.
- [14] James Harrison\*, Animesh Garg\*, Boris Ivanovic, **Yuke Zhu**, Silvio Savarese, Li Fei-Fei, Marco Pavone. ADAPT: Zero-Shot Adaptive Policy Transfer for Stochastic Dynamical Systems. *International Symposium on Robotics Research (ISRR)*, 2017.
- [13] **Yuke Zhu**\*, Daniel Gordon\*, Eric Kolve, Dieter Fox, Li Fei-Fei, Abhinav Gupta, Roozbeh Mottaghi, Ali Farhadi. Visual Semantic Planning using Deep Successor Representations. *International Conference on Computer Vision (ICCV)*, 2017.

- [12] Ajay Mandlekar\*, **Yuke Zhu**\*, Animesh Garg\*, Li Fei-Fei, Silvio Savarese. Adversarially Robust Policy Learning through Active Construction of Physically-Plausible Perturbations. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
  - [11] **Yuke Zhu**, Joseph J. Lim, Li Fei-Fei. Knowledge Acquisition for Visual Question Answering via Iterative Querying. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
  - [10] Danfei Xu, **Yuke Zhu**, Christopher B. Choy, Li Fei-Fei. Scene Graph Generation by Iterative Message Passing. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
  - [9] **Yuke Zhu**, Roozbeh Mottaghi, Eric Kolve, Joseph J. Lim, Abhinav Gupta, Li Fei-Fei, Ali Farhadi. Target-driven Visual Navigation in Indoor Scenes using Deep Reinforcement Learning. *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
  - [8] Ranjay Krishna, **Yuke Zhu**, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalanditis, Li-Jia Li, David A. Shamma, Michael Bernstein, Li Fei-Fei. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision (IJCV)*, 2017.
  - [7] **Yuke Zhu**, Groth Oliver, Michael Bernstein, Li Fei-Fei. Visual7W: Grounded Question Answering in Images. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
  - [6] Tian Lan\*, **Yuke Zhu**\*, Amir Zamir, Silvio Savarese. Action Recognition by Hierarchical Mid-level Action Elements. *International Conference on Computer Vision (ICCV)*, 2015.
  - [5] **Yuke Zhu**, Alireza Fathi, Li Fei-Fei. Reasoning About Object Affordances in a Knowledge Base Representation. *European Conference on Computer Vision (ECCV)*, 2014.
  - [4] Oliver Schulte, Hassan Khosravi, Arthur Kirkpatrick, Tianxiang Gao, **Yuke Zhu**. Modelling Relational Statistics With Bayes Nets. *Machine Learning Journal 94(1):105-125*, 2014.
  - [3] Alfred Zong and **Yuke Zhu**. StrokeBank: Automating Personalized Chinese Handwriting Generation. *AAAI Conference on Innovative Applications of Artificial Intelligence (IAAI)*, 2014.
  - [2] Zhao Song and **Yuke Zhu**. Graphical Model-based Learning in High Dimensional Feature Spaces. *The Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, 2013.
  - [1] **Yuke Zhu**, Tian Lan, Yijian Yang, Steven Robinovitch, Greg Mori. Latent Spatio-temporal Models for Action Localization and Recognition in Nursing Home Surveillance Video. *IAPR International Conference on Machine Vision Applications (MVA)*, 2013.
- 
- [8] Soroush Nasiriany\*, Fei Xia\*, Wenhao Yu\*, Ted Xiao\*, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, **Yuke Zhu**, Andy Zeng, Karol Hausman, Nicolas Heess, Chelsea Finn, Sergey Levine, Brian Ichter\*. PIVOT: Iterative Visual Prompting Elicits Actionable Knowledge for VLMs. *arXiv:2402.07872*, 2024.

**TECHNICAL  
REPORTS**

- [7] Roya Firoozi, Johnathan Tucker, Stephen Tian, Anirudha Majumdar, Jiankai Sun, Weiyu Liu, **Yuke Zhu**, Shuran Song, Ashish Kapoor, Karol Hausman, Brian Ichter, Danny Driess, Jiajun Wu, Cewu Lu, Mac Schwager. Foundation Models in Robotics: Applications, Challenges, and the Future. *arXiv:2312.07843*, 2023.
- [6] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, **Yuke Zhu**, Linxi Fan<sup>†</sup>, Anima Anandkumar<sup>†</sup>. Voyager: An Open-Ended Embodied Agent with Large Language Models. *arXiv:2305.16291*, 2023.
- [5] Ajay Mandlekar, Danfei Xu, Roberto Martín-Martín, **Yuke Zhu**, Li Fei-Fei, Silvio Savarese. Human-in-the-Loop Imitation Learning using Remote Teleoperation. *arXiv:2012.06733*, 2020.
- [4] **Yuke Zhu**, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Soroush Nasiriany, Yifeng Zhu. robosuite: A Modular Simulation Framework and Benchmark for Robot Learning. *arXiv:2009.12293*, 2020.
- [3] Linxi Fan\*, **Yuke Zhu**\*, Jiren Zhu, Zihua Liu, Orient Zeng, Anhit Gupta, Joan Creus-Costa, Silvio Savarese, Fei-Fei Li. SURREAL-System: Fully-Integrated Stack for Distributed Deep Reinforcement Learning. *arXiv:1909.12989*, 2019.
- [2] Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, **Yuke Zhu**, Abhinav Gupta, Ali Farhadi. AI2-THOR: An Interactive 3D Environment for Visual AI. *arXiv:1712.05474*, 2017.
- [1] **Yuke Zhu**, Ce Zhang, Christopher R, Li Fei-Fei. Building a Large-scale Multi-modal Knowledge Base System for Answering Visual Queries. *arXiv:1507.05670*, 2015.

## DISSERTATIONS

- [1] **Yuke Zhu**. Closing the Perception-Action Loop: Towards General-Purpose Robot Autonomy. Ph.D. Thesis, Stanford University, 2019.

## PATENTS

- [10] Data selection based on uncertainty quantification. Jonathan Tremblay, Fabio Tozeto Ramos, **Yuke Zhu**, Anima Anandkumar, Guanya Shi. *US Patent Application 17/331,466*, filed May 26, 2021.
- [9] Using neural networks to perform object detection, instance segmentation, and semantic correspondence from bounding box supervision. Zhiding Yu, Shiyi Lan, Chris Choy, Subhashree Radhakrishnan, Guilin Liu, **Yuke Zhu**, Anima Anandkumar. *US Patent Application 17/177,068*, filed February 16, 2021.
- [8] Ajay U. Mandlekar, **Yuke Zhu**, Animesh Garg, Silvio Savarese, Fei-Fei Li. Methods and systems to remotely operate robotic devices. *International Patent Application PCT/US2020/058542*, filed November 02, 2020.
- [7] **Yuke Zhu**, Yifeng Zhu, Stanley Thomas Birchfield, Jonathan Tremblay. Machine learning model for task and motion planning. *US Patent Application 17/082,972*, filed October 28, 2020.
- [6] Animesh Garg, Hongyu Ren, **Yuke Zhu**, Anima Anandkumar. Online task inference for compositional tasks with context adaptation. *US Patent Application 16/945,753*, filed July 31, 2020.
- [5] Jianchao Yang, **Yuke Zhu**, Ning Xu, Kevin Tang, Jia Li. Systems and methods for content navigation with automated curation. *US Patent 10,733,255*, issued August 04, 2020.

- [4] Jianchao Yang, **Yuke Zhu**, Ning Xu, Kevin Dechau Tang, Jia Li. Content navigation with automated curation. *US Patent Application 16/918,343*, filed July 1, 2020.
- [3] Jianchao Yang, **Yuke Zhu**, Ning Xu, Kevin Tang, Jia Li. Automated content curation and communication. *US Patent 10,387,514*, issued August 20, 2019.
- [2] Jianchao Yang, **Yuke Zhu**, Ning Xu, Kevin Tang, Jia Li. Automated image processing and content curation. *US Patent 10,382,373*, issued August 13, 2019.
- [1] Saran Tunyasuvunakool, **Yuke Zhu**, Joshua Merel, Janos Kramar, Ziyu Wang, Nicolas Heess. Reinforcement and imitation learning for a task. *US Patent Application 16/174,112*, filed October 29, 2018.

## TEACHING

<i>Instructor</i> , The University of Texas at Austin CS 391R: Robot Learning	Fall 2023 – 2024
<i>Instructor</i> , The University of Texas at Austin CS 343: Artificial Intelligence	Spring 2022 – 2023
<i>Instructor</i> , The University of Texas at Austin CS 343: Artificial Intelligence	Spring 2021 – 2022
<i>Instructor</i> , The University of Texas at Austin CS 391R: Robot Learning	Fall 2021 – 2022
<i>Instructor</i> , The University of Texas at Austin CS 343: Artificial Intelligence	Spring 2020 – 2021
<i>Instructor</i> , The University of Texas at Austin CS 391R: Robot Learning	Fall 2020 – 2021
<i>Course Assistant</i> , Stanford University CS231N: Convolutional Neural Networks for Visual Recognition	Winter 2014 – 2015
<i>Course Assistant</i> , Stanford University CS 131: Computer Vision: Foundations and Applications	Fall 2014 – 2015
<i>Course Assistant</i> , Stanford University CS 193C: Client-Side Internet Technologies	Summer 2013 – 2014
<i>Course Assistant</i> , Stanford University CS 431: High-Level Vision – Behaviors, Neurons and Computational Models	Spring 2013 – 2014

## SERVICES

### Conference Organizer

- *Local Chair*, International Conference on Development and Learning (ICDL), 2024
- *Publications Chair*, Conference on Robot Learning (CoRL), 2023
- *Website Chair*, Robotics: Science and Systems (RSS), 2022
- *Senior Program Committee*, Conference on Lifelong Learning Agents (CoLLAs), 2022
- *Senior Program Committee*, International Joint Conference on Artificial Intelligence (IJCAI), 2023

- *Area Chair*, Conference on Robot Learning (CoRL), 2020, 2021, 2022, 2023
- *Area Chair*, International Conference on Computer Vision (ICCV), 2021
- *Area Chair*, NeurIPS Datasets and Benchmarks Track (NeurIPS), 2021
- *Area Chair*, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2021, 2022
- *Area Chair*, Conference on Neural Information Processing Systems (NeurIPS), 2022
- *Area Chair*, European Conference on Computer Vision (ECCV), 2022
- *Area Chair*, International Conference on Machine Learning (ICML), 2023
- *Area Chair*, Robotics: Science and Systems (RSS), 2023, 2024

### **Workshop Organizer**

- *Co-organizer*, Towards Reliable and Deployable Learning-Based Robotic Systems, Conference on Robot Learning (CoRL), 2023
- *Co-organizer*, Causality for Robotics: Answering the Question of Why, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2023
- *Co-organizer*, 2nd Workshop on 3D Vision and Robotics, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023
- *Co-organizer*, Texas Regional Robotics Symposium (TEROS), 2022
- *Co-organizer*, Workshop on Socially Responsible Machine Learning, International Conference on Learning Representations (ICLR), 2022
- *Co-organizer*, Workshop on Visual Learning and Reasoning for Robotics, Robotics: Science and Systems (RSS), 2021
- *Co-organizer*, Tutorial on End-to-end GPU-accelerated Learning and Control for Robotics with Isaac Gym, Robotics: Science and Systems (RSS), 2021
- *Co-organizer*, Workshop on 3D Vision and Robotics, IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2021
- *Co-organizer*, Workshop on Embodied Multimodal Learning, International Conference on Learning Representations (ICLR), 2021
- *Co-organizer*, Workshop on Visual Learning and Reasoning for Robotic Manipulation, Robotics: Science and Systems (RSS), 2020
- *Co-organizer*, Workshop on Perspectives in Robot Learning: Causality and Imitation, Robotics: Science and Systems (RSS), 2018
- *Program Committee*, A Roadmap to Never-Ending RL Workshop, International Conference on Learning Representations (ICLR), 2021
- *Program Committee*, 4th Lifelong Learning Workshop, International Conference on Machine Learning (ICML), 2020
- *Program Committee*, Workshop on Multitask and Lifelong Reinforcement Learning, International Conference on Machine Learning (ICML), 2019
- *Program Committee*, RSS Pioneers, Robotics: Science and Systems (RSS), 2019
- *Program Committee*, Workshop on Vision in Practice on Autonomous Robots (ViPAR), International Conference on Computer Vision (ICCV), 2017
- *Program Committee*, Workshop on Challenges in Robot Learning, Conference on Neural Information Processing Systems (NeurIPS), 2017
- *Scientific Advisory Board*, Workshop on BEHAVIOR: Benchmark for Everyday Household Activities in Virtual, Interactive, and Ecological Environments, International Conference on Computer Vision (ICCV), 2021

### **Grant Reviewer**

- *Reviewer*, European Research Council, Europe, 2023

- *Reviewer*, United States-Israel Binational Science Foundation, USA, 2020
- *Reviewer and Panelist*, National Science Foundation, USA, 2020 – 2022
- *Reviewer*, Army Research Office, USA, 2020

### **Editorial Board**

- *Associate Editor*, International Conference on Robotics and Automation (ICRA), 2020, 2022, 2023
- *Associate Editor*, International Conference on Intelligent Robots and Systems (IROS), 2021, 2022
- *Editorial Board*, Machine Learning, 2022-2025
- *Topic Editor*, Trustworthy Machine Learning, Frontiers, 2021-2023

### **Conference Reviewer**

- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)
- International Conference on Computer Vision (ICCV)
- European Conference on Computer Vision (ECCV)
- Asian Conference on Computer Vision (ACCV)
- International Conference on Intelligent Robots and Systems (IROS)
- International Conference on Robotics and Automation (ICRA)
- International Conference on Humanoid Robots (Humanoids)
- Conference on Neural Information Processing Systems (NeurIPS)
- International Conference on Machine Learning (ICML)
- Conference on Robot Learning (CoRL)
- Robotics: Science and Systems (RSS)
- International Conference on Learning Representations (ICLR)
- ACM SIGGRAPH Conference (SIGGRAPH)
- Learning for Dynamics & Control Conference (L4DC)

### **Journal Reviewer**

- International Journal of Robotics Research (IJRR)
- IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)
- IEEE Transactions on Multimedia (T-MM)
- IEEE Robotics and Automation Letters (RA-L)
- ACM Computing Surveys (CSUR)
- Transactions on Machine Learning Research (TMLR)
- Science Robotics

### **Departmental and University Committee**

- *Robotics Graduate Portfolio Program Committee*, UT Austin, 2021, 2022, 2023
- *Doctoral Admissions Committee*, UT Austin Computer Science, 2020, 2021
- *Faculty Recruiting Committee*, UT Austin Computer Science, 2022

### **Outreach Program**

- *Inclusion@RSS Mentor*, Robotics: Science and Systems, 2022
- *Faculty Mentor*, UT-Austin CS Directed Research Program, 2020, 2021
- *Undergraduate Mentor*, Stanford AI Undergraduate Mentoring Program, 2018

**Ph.D. Advisor**

- Yifeng Zhu, Computer Science (Fall 2019–, co-advised with Peter Stone)
- Zhenyu Jiang, Computer Science (Fall 2020–)
- Soroush Nasiriany, Computer Science (Fall 2020–)
- Mingyo Seo, Electrical and Computer Engineering (Spring 2021–, co-advised with Luis Sentis)
- Huihan Liu, Computer Science (Fall 2021–)
- Rutav Shah, Computer Science (Fall 2022–, co-advised with Roberto Martín-Martín)
- Jake Grigsby, Computer Science (Fall 2022–)

**Master's Advisor**

- Braham Snyder, Computer Science (2021–2023)  
Thesis: Towards Convergent Offline Reinforcement Learning
- Cheng-Chun Hsu, Computer Science (2021–2023) Next: Ph.D. at UT-Austin  
Thesis: Building Digital Twins of Articulated Objects and Scenes through Interactive Perception
- Steve Han, Computer Science (2022–2023) Next: Roblox  
Thesis: VR Teleoperation Interface for Learning Loco-manipulation of Humanoid Robots

**Undergraduate Advisor**

- Abhishek Joshi, Computer Science and Math (2020–)
- Zhiyao Bao, Computer Science (2020–2022) Next: Meta
- Aditya Arjun, Computer Science (2021|2022) Next: Jane Street  
Thesis: Using Counterfactual Options for Hierarchical Reinforcement Learning — Extension to Robotic Grasping
- Alexy Skoutnev, Mechanical Engineering (2020–2022) Next: Ph.D. at Vanderbilt
- Pranav Atreya, Computer Science (2022–2023) Next: Ph.D. at UC Berkeley
- Lance Zhang, Computer Science (2022–)
- Kyutae Sim, Computer Science (2022–)  
Thesis: Human-in-the-Loop Learning for Humanoid Robots
- Abhi Maddukuri, Computer Science (2023–)

**Visiting Scholar Advisor**

- Tian Gao, Tsinghua University (2021–2023) Next: Ph.D. at Stanford
- Weikang Wan, Peking University (2023)
- Quantao Yang, Örebro University (2023) Next: ABB Corporate Research

**Doctoral Committee Member**

- Tongzheng Ren, Computer Science, Supervisor: Sujay Sanghavi
- Xingyu Gong, Electrical and Computer Engineering, Supervisor: Atlas Wang
- Eddy Hudson, Computer Science, Supervisor: Peter Stone
- Yifan Jiang, Electrical and Computer Engineering, Supervisor: Atlas Wang
- Yuqian Jiang, Computer Science, Supervisor: Peter Stone
- Dian Chen, Computer Science, Supervisor: Philipp Krähenbühl
- Xingyi Zhou, Computer Science, Supervisor: Philipp Krähenbühl

- Sadegh Rabiee, Computer Science, Supervisor: Joydeep Biswas
- Junhyeok Ahn, Mechanical Engineering, Supervisor: Luis Sentis
- Yagiz Savas, Aerospace Engineering, Supervisor: Ufuk Topcu
- Garrett Bingham, Computer Science, Supervisor: Risto Miikkulainen
- Tushar Nagarajan, Computer Science, Supervisor: Kristen Grauman
- Wonjoon Goo, Computer Science, Supervisor: Scott Niekum
- Yihao Feng, Computer Science, Supervisor: Qiang Liu
- Santhosh Kumar Ramakrishnan, Computer Science, Supervisor: Kristen Grauman
- Cyrus Neary, The Oden Institute, Supervisor: Ufuk Topcu

#### **External Doctoral Committee Member**

- Sasha Salter, Balliol College, University of Oxford, Supervisor: Ingmar Posner

#### **Master's Thesis Committee Member**

- Alexander Joseph Zuzow, Computer Science, Supervisor: Atlas Wang
- Charles Nimo, Computer Science, Supervisor: Ying Ding
- Sagnik Majumder, Computer Science, Supervisor: Kristen Grauman
- Priyanka Mandikal, Computer Science, Supervisor: Kristen Grauman

#### **Undergraduate Honor Thesis Committee Member**

- Tianwei Yin, Computer Science, Supervisor: Philipp Krähenbühl
- Kevin Black, Computer Science, Supervisor: Scott Niekum

#### **Stanford University**

2013 – 2020

#### **Master's Student Mentor**

- Andrey Kurenkov, Computer Science, Next: Ph.D. at Stanford
- Julian Gao, Computer Science, Next: Robotics engineer at Dexterity
- Anchit Gupta, Computer Science, Next: Machine learning engineer at Facebook
- Alex Kaiyi Fu, Computer Science, Next: Software engineer at WeRide.ai
- Josiah Wong, Mechanical Engineering: Next: Ph.D. at Stanford
- Guanzhi Wang, Computer Science: Next: Ph.D. at Caltech

#### **Undergraduate Student Mentor**

- Viraj Mehta, Computer Science, Next: Ph.D. student at CMU
- Russell Kaplan, Computer Science, Next: Senior machine learning scientist at Tesla
- Jiren Zhu, Math, Next: Quantitative developer at D. E. Shaw Group
- Joan Creus-Costa, Physics, Next: CS Master's at Stanford
- William B. Shen, Computer Science, Next: Ph.D. at Stanford
- James (Zihua) Liu, Computer Science, Next: Co-founder of include.ai
- Orien Zeng, Computer Science, Next: Software engineer at Nuro
- Andrew Kondrich, Computer Science, Next: ML research engineer at Scale AI

#### **Visiting Scholar Mentor**

- Yurong You, SJTU, Next: Ph.D. at Cornell
- Chen Wang, SJTU, Next: Ph.D. at Stanford
- Zengyi Qin, Tsinghua, Next: Ph.D. at MIT
- Oliver Groth, TU-Dresden, Next: Ph.D. at Oxford

- Suraj Nair, Caltech, Next: Ph.D. at Stanford

## EMPLOYMENT HISTORY

*Assistant Professor*, Department of Computer Science  
The University of Texas at Austin, Austin, TX, USA Aug 2020 – Present

- Tenure-track faculty and director of the Robot Perception and Learning Lab
- Affiliations: Texas Robotics, Machine Learning Laboratory, and Center for Perceptual Systems

*Senior Research Scientist* Oct 2019 – Present  
NVIDIA Research, Santa Clara, CA, USA

- Interdisciplinary research in machine learning, computer vision, and robotics
- Co-leading the Generalist Embodied Agent Research (GEAR) team

*Visiting Scholar* Sept 2019 – Aug 2020  
Stanford University, Computer Science Department, Stanford, CA, USA

- Robotics research in the Stanford People, AI & Robots (PAIR) Group

*Research Intern* Jun – Sept 2017  
DeepMind Technologies Ltd., London, England, UK

- Deep reinforcement and imitation learning for vision-based robot manipulation

*Research Intern* Jun – Sept 2016  
Allen Institute for Artificial Intelligence, Seattle, WA, USA

- Building photorealistic simulated 3D environments for visual AI and deep reinforcement learning models for visual navigation

*Research Intern* Jun – Sept 2015  
Snap Inc., Venice, CA, USA

- Developing deep learning models for video understanding at scale in the Snap Research team

*Software Engineer Intern* Apr – July 2013  
Twitter Inc., San Francisco, CA, USA

- Growth hacking in the Activation & Messaging team for Twitter user retention

*Research Assistant* Dec 2011 – Apr 2013  
SFU Computational Logic Lab, Burnaby, BC, Canada

- Research in statistical relational learning of efficient learning and inference with Bayesian Networks and Markov Logic Networks

*Research Assistant* Jan 2012 – Apr 2013  
SFU Vision and Media Lab, Burnaby, BC, Canada

- Research in action understanding of nursing home videos for Technology for Injury Prevention in Seniors (TIPS) program

## INVITED TALKS

“Pathway to Generalist Robots: Scaling Law, Data Flywheel, and Humanlike Embodiment”

- CoRL 2023 Early Career Keynote, Atlanta, GA Nov 2023
- Babuška Forum Series, Austin, TX Feb 2024

“Building Multimodal Foundation Models for Embodied Agents”

- ICCV 2023 PERception, Decision making and REASONing through Multimodal foundational modeling (PerDream) Workshop, Online Oct 2023
- “Building Generalist Agents with Internet-scale Knowledge”*
- Salesforce Distinguished AI Speaker Seminar, Online Oct 2023
  - UC Berkeley BAIR Research Tech Talk, Online Sept 2023
- “Is Scale All We Need for Robotics Foundation Models?”*
- KAIST AI Colloquium, Online Oct 2023
  - Cornell Robotics Seminar, Online Sept 2023
- “Data-Efficient Robot Learning with Prior Data”*
- ICRA 2023 Pretraining for Robotics Workshop, London, UK May 2023
- “The Data Pyramid for General Robot Manipulation”*
- AAAI 2023 New Faculty Highlights, Washington, DC Feb 2023
  - Oregon State University, Corvallis, OR Feb 2023
  - MIT Embodied Intelligence Seminar, Boston, MA Dec 2022
  - Texas Robotics Symposium, Austin, TX Nov 2022
  - ML for Robotics with Large Datasets Workshop, Berkeley, CA Oct 2022
  - Salesforce Research, Palo Alto, CA Oct 2022
  - Ambi Robotics, Berkeley, CA Oct 2022
- “Objects, Skills, and the Quest for Compositional Robot Autonomy”*
- Amazon Web Services, Santa Clara, CA Aug 2022
  - Nuro, Mountain View, CA Aug 2022
  - Stanford Robotics Seminar, Stanford, CA Mar 2022
- “The Systems Challenge of Building Robot Autonomy”*
- UT Austin Sys/ML Workshop, Austin, TX Mar 2022
- “Visual Affordance Learning for Robot Manipulation”*
- Peking University VCL Seminar, Online Sept 2021
  - Toyota Research Institute, Los Altos, CA Aug 2021
- “Compositional Learning for Robot Autonomy via Modularity and Abstraction”*
- VinAI Research Seminar Series, Online Apr 2021
  - UIUC Robotics Seminar, Urbana-Champaign, IL Apr 2021
  - Institute for Foundations of Machine Learning Seminar, Austin, TX Mar 2021
- “Visual Imitation Learning: Generalization, Perceptual Grounding, & Abstraction”*
- RSS 2020 Workshop on Advances and Challenges in Imitation Learning for Robotics July 2020
- “Building General-Purpose Robot Autonomy: A Progressive Roadmap”*
- CS231N Guest Lecture, Stanford University, Stanford, CA June 2021
  - Samsung Forum, Samsung Strategy and Innovation Center June 2020
- “Learning Keypoint Representations for Robot Manipulation”*

- IROS 2019 Workshop on Learning Representations for Planning and Control, Macau, China Nov 2019

*“Learning How-To Knowledge from the Web”*

- 3rd International Workshop on the Applications of Knowledge Representation and Semantic Technologies in Robotics, Macau, China Nov 2019

*“Closing the Perception-Action Loop”*

- Facebook AI Research, Menlo Park, CA June 2019
- NVIDIA Research, Santa Clara, CA June 2019
- Google Brain Robotics, Mountain View, CA May 2019
- Georgia Institute of Technology, Atlanta, GA Apr 2019
- University of Southern California, Los Angeles, CA Apr 2019
- McGill University, Montreal, QC, Canada Apr 2019
- Yale University, New Haven, CT Mar 2019
- New York University, New York, NY Mar 2019
- University of Toronto, Toronto, ON, Canada Mar 2019
- Carnegie Mellon University, Robotics Institute, Pittsburgh, PA Feb 2019
- The University of Texas at Austin, Austin, TX Feb 2019
- Princeton University, Princeton, NJ Feb 2019
- Massachusetts Institute of Technology, Boston, MA Dec 2018
- Stanford University, Stanford, CA Nov 2018

*“From Disembodied Visual Recognition to Robotic Interactive Perception”*

- RSS Pioneers Workshop, Pittsburgh, PA June 2018

*“Towards Generalizable Robot Learning with Perceptual Intelligence”*

- Chinese University of Hong Kong Shenzhen, Shenzhen, China Mar 2018
- Workshop on Future Leaders of AI Retreat, Shanghai, China Dec 2017
- Shanghai Jiaotong University, Shanghai, China Dec 2017

*“Knowledge Acquisition for Visual Question Answering”*

- Stanford Semantics and Geometry Seminar, Stanford, CA Mar 2016

*“Computer Vision Algorithms for Fall Detection”*

- Technology for Injury Prevention in Seniors (TIPS) 3rd Annual Research Symposium, Vancouver, BC, Canada Nov 2012

**EARLIER  
HONORS &  
AWARDS**

- Tencent AI Lab PhD Fellowship 2017 – 2018
- AAAI-14 Student Scholarship 2014
- DDP Outstanding Academic Achievement Award 2014
- Simon Fraser University Computing Science Graduation Award 2013
- Simon Fraser University President’s Honour Roll 2012, 2013
- Simon Fraser University Open Scholarship 2012, 2013
- 1st Place in Simon Fraser University 8th Winter Programming Contest 2012
- 5th Place in the ACM Pacific Northwest Programming Contest 2012

- Simon Fraser University Entrance Scholarship 2011
- National Scholarship of China (Top 2% in Zhejiang University) 2010, 2011
- Zhejiang University First-class Academic Excellence Scholarship 2010, 2011
- Zhejiang University Research and Innovation Scholarship 2010
- Gold Medal in the 8th ACM Programming Contest of Zhejiang Province 2010
- First Prize in the 10th ACM Programming Contest of Zhejiang University 2010
- First Prize in National Olympiad of Informatics in Shandong Province 2008

**SELECTED  
PRESS  
COVERAGE**

- [12] “Eureka! NVIDIA Research Breakthrough Puts New Spin on Robot Learning,” by Angie Lee, *NVIDIA Blog*. October 20, 2023.
- [11] “Robot from the University of Texas at Austin Makes Instant Ramen,” by Bryan Ke, *Yahoo! News*. June 15, 2023.
- [10] “This AI Used GPT-4 to Become an Expert Minecraft Player,” by Devin Coldewey, *TechCrunch*. June 2, 2023.
- [9] “To Really Judge an AI’s Smarts, Give it One of These IQ Tests,” by Matthew Hutson, *IEEE Spectrum*. Feb 2, 2021.
- [8] “AI Researchers Challenge a Robot to Ride a Skateboard in Simulation,” by Khari Johnson, *VentureBeat*. Oct 6, 2020.
- [7] “RoboTurk: A Crowdsourcing Platform for Imitation Learning in Robotics,” by Ingrid Fadelli, *Tech Xplore*. Nov 21, 2018.
- [6] “Robots Learn Tasks from People with Framework Developed by Stanford Researchers,” by Sofie Bates, *Stanford News*. Oct 26, 2018.
- [5] “Robot See, Robot Do: Bots Learn by Watching Human Behavior,” by Noah Kravitz, *Nvidia Blog*. Apr 3, 2018.
- [4] “Virtual Reality Training Ground Helps Robots Prepare for the Real World,” by Luke Dormehl, *Digital Trends*. Feb 19, 2018.
- [3] “A Detailed Virtual House Will Help Robots Train to Become Your Butler,” by Jackie Snow, *MIT Technology Review*. Feb 16, 2018.
- [2] “AI2-THOR Interactive Simulation Teaches AI About Real World,” by Jeremy Hsu, *IEEE Spectrum*. Feb 15, 2018.
- [1] “Next Big Test for AI: Making Sense of the World,” by Will Knight, *MIT Technology Review*. Jan 26, 2016.