

# Yuke Zhu

Gates Dell Complex, 2317 Speedway, Stop D9500  
Austin, TX, 78712 USA  
[cs.utexas.edu/~yukez](http://cs.utexas.edu/~yukez)  
[yukez@cs.utexas.edu](mailto:yukez@cs.utexas.edu)

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

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

**Visiting Scholar**, Department of Computer Science     2019 – Present  
Stanford University, Stanford, CA, USA

**EDUCATION**     **Stanford University**, Stanford, CA, USA     2015 – 2019  
*Ph.D.* in Computer Science  
Advisors: Fei-Fei Li, Silvio Savarese

**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  
Ranked 1st out of 31 in the Dual Degree Program

- PUBLICATIONS**
- [1] 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.
  - [2] 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.
  - [3] 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.
  - [4] 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.
  - [5] 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.
  - [6] 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.

- [7] 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.
- [8] 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.
- [9] 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.**
- [10] 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.
- [11] 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.**
- [12] 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 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 Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [13] 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.
- [14] 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.
- [15] **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.
- [16] 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
- [17] 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.

- [18] 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.
- [19] 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.
- [20] **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.
- [21] 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.
- [22] **Yuke Zhu**, Joseph J. Lim, Li Fei-Fei. Knowledge Acquisition for Visual Question Answering via Iterative Querying. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [23] Danfei Xu, **Yuke Zhu**, Christopher B. Choy, Li Fei-Fei. Scene Graph Generation by Iterative Message Passing. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [24] **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.
- [25] 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.
- [26] **Yuke Zhu**, Groth Oliver, Michael Bernstein, Li Fei-Fei. Visual7W: Grounded Question Answering in Images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [27] Tian Lan\*, **Yuke Zhu\***, Amir Zamir, Silvio Savarese. Action Recognition by Hierarchical Mid-level Action Elements. *International Conference on Computer Vision (ICCV)*, 2015.
- [28] **Yuke Zhu**, Alireza Fathi, Li Fei-Fei. Reasoning About Object Affordances in a Knowledge Base Representation. *European Conference on Computer Vision (ECCV)*, 2014.
- [29] Oliver Schulte, Hassan Khosravi, Arthur Kirkpatrick, Tianxiang Gao, **Yuke Zhu**. Modelling Relational Statistics With Bayes Nets. *Machine Learning Journal 94(1):105-125*, 2014.
- [30] Alfred Zong and **Yuke Zhu**. StrokeBank: Automating Personalized Chinese Handwriting Generation. *AAAI Conference on Innovative Applications of Artificial Intelligence (IAAI)*, 2014.
- [31] Zhao Song and **Yuke Zhu**. Graphical Model-based Learning in High Dimensional Feature Spaces. *The Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI)*, 2013.

- [32] **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.

## TECHNICAL REPORTS

- [1] 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. *arXiv:1909.13003*, 2019.
- [2] Linxi Fan\*, **Yuke Zhu**\*, Jiren Zhu, Zihua Liu, Orien 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.
- [3] 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.
- [4] **Yuke Zhu**, Ce Zhang, Christopher R, Li Fei-Fei. Building a Large-scale Multimodal 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

- [1] 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.
- [3] Saran Tunyasuvunakool, **Yuke Zhu**, Joshua Merel, Janos Kramar, Ziyu Wang, Nicolas Heess. Automated image processing and content curation. *US Patent Application 16/174,112*, filed October 29, 2018.

## SERVICES

### Conference Reviewer

- IEEE 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)
- International Conference on Learning Representations (ICLR)
- ACM SIGGRAPH Conference

### Journal Reviewer

- 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)

### Grant Reviewer

- *Reviewer and Panelist*, National Science Foundation, USA, 2020

### Editorial Board

- *Associate Editor*, International Conference on Robotics and Automation, 2020

### Area Chair

- *Area Chair*, IEEE Conference on Computer Vision and Pattern Recognition, 2021

### Workshop Organizer

- *Program Committee*, Workshop on Multitask and Lifelong Reinforcement Learning, International Conference on Machine Learning (ICML), 2019
- *Co-organizer*, Workshop on Perspectives in Robot Learning: Causality and Imitation, Robotics: Science and Systems (RSS), 2018
- *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

### Outreach Program

- *Undergraduate Mentor*, Stanford AI Undergraduate Mentoring Program, 2018

## SPEECHES

### *“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
- 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

**TEACHING**

*Course Assistant*, Stanford University Winter 2014 – 2015  
CS231N: Convolutional Neural Networks for Visual Recognition

*Course Assistant*, Stanford University Fall 2014 – 2015  
CS 131: Computer Vision: Foundations and Applications

*Course Assistant*, Stanford University Summer 2013 – 2014  
CS 193C: Client-Side Internet Technologies

*Course Assistant*, Stanford University Spring 2013 – 2014  
CS 431: High-Level Vision – Behaviors, Neurons and Computational Models

**EMPLOYMENT**

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

- Tenure-track faculty in UTCS and principal investigator of the Learning, Embodied Agents, and Perception (LEAP) Lab

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

- Interdisciplinary research in machine learning, computer vision, and robotics

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

- Robotics research in the Stanford People, AI & Robots 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

## HONORS & AWARDS

### Awards and Prizes

- IROS Best Paper Award on Cognitive Robotics, Finalist 2019
- IEEE ICRA Best Conference Paper Award 2019
- IEEE ICRA Best Paper Award on Cognitive Robotics, Finalist 2019
- RSS Pioneers Workshop Travel Award 2018
- DDP Outstanding Academic Achievement Award 2014
- AAAI-14 Student Scholarship 2014
- Simon Fraser University Computing Science Graduation Award 2013
- Simon Fraser University President's Honour Roll 2012, 2013
- 1st Place in Simon Fraser University 8th Winter Programming Contest 2012
- 5th Place in the ACM Pacific Northwest Programming Contest 2012
- 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

### Scholarships and Fellowships

- Tencent AI Lab PhD Fellowship 2017 – 2018
- Simon Fraser University Open Scholarship 2012, 2013
- 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

## SELECTED PRESS COVERAGE

- [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.

## ADVISING

### Master’s Students

- Andrey Kurenkov (CS, Stanford) Next: Ph.D. student at Stanford
- Julian Gao (CS, Stanford) Next: Robotics engineer at Dexterity
- Anchit Gupta (CS, Stanford) Next: Machine learning engineer at Facebook
- Alex Kaiyi Fu (CS, Stanford) Next: Software engineer at WeRide.ai

### Undergraduate Students

- Viraj Mehta (CS, Stanford) Next: Ph.D. student at CMU
- Suraj Nair (CS, Caltech) Next: Ph.D. student at Stanford
- Russell Kaplan (CS, Stanford) Next: Senior machine learning scientist at Tesla
- Jiren Zhu (Math, Stanford) Next: Quantitative developer at D. E. Shaw Group
- Joan Creus-Costa (Physics, Stanford) Next: CS Master’s student at Stanford
- Zihua Liu (CS, Stanford) Next: Stealth startup
- Andrew Kondrich (CS, Stanford)

### Independent Research

- Yurong You (CS, SJTU) Next: Ph.D. student at Cornell
- Chen Wang (CS, SJTU)
- Zengyi Qin (EE, Tsinghua)