

Yuke Zhu

Gates Room 242, 353 Serra Mall
Stanford, CA, 94305 USA
web.stanford.edu/~yukez
yukez@cs.stanford.edu

OBJECTIVE My research objective is to teach robots to understand and to interact with the visual world. My research resides at the intersection of computer vision, machine learning, and robotics, with a focus on visual knowledge and deep reinforcement learning.

EDUCATION **Stanford University**, Stanford, CA, USA 2015 – Present
Ph.D. in Computer Science GPA: 4.0 / 4.0
Advisors: Fei-Fei Li, Silvio Savarese

Stanford University, Stanford, CA, USA 2013 – 2015
Master of Science in Computer Science GPA: 3.98 / 4.0
Advisors: Fei-Fei Li, Silvio Savarese

Simon Fraser University, Vancouver, BC, Canada 2011 – 2013
Bachelor of Science in Computer Science (first class with distinction) GPA: 4.27 / 4.3
Advisors: Greg Mori, Oliver Schulte

Zhejiang University, Hangzhou, China 2009 – 2013
Bachelor of Engineering in Computer Science and Technology GPA: 3.96 / 4.0
Ranked 1st out of 31 in the Dual Degree Program

- PUBLICATIONS**
- [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 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 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 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.

TECHNICAL REPORTS

- [3] 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. *arXiv:1807.03480*, 2018.
- [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 Multimodal Knowledge Base System for Answering Visual Queries. *arXiv:1507.05670*, 2015.

SERVICES

Conference Reviewing

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 (NIPS)
Conference on Robot Learning (CoRL)

Journal Reviewing

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)

Workshop Organizing

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 (NIPS), 2017

Outreach Program

Undergraduate Mentor, Stanford AI Undergraduate Mentoring Program 2018

SPEECHES

“Towards Generalizable Robot Learning with Perceptual Intelligence”

- Invited Talk, Chinese University of Hong Kong Shenzhen (CUHK-Shenzhen), Shenzhen, China Mar 2018
- Invited Talk, Workshop on Future Leaders of AI Retreat, New York University Shanghai (NYU-Shanghai), Shanghai, China Dec 2017
- Invited Talk, Shanghai Jiaotong University, Shanghai, China. Dec 2017

“Knowledge Acquisition for Visual Question Answering”

- Invited Talk, Stanford Semantics and Geometry Seminar, Stanford, CA, USA 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	<i>Teaching Assistant</i> , Stanford University CS231N: Convolutional Neural Networks for Visual Recognition	Winter 2014 – 2015
	<i>Teaching Assistant</i> , Stanford University CS 131: Computer Vision: Foundations and Applications	Fall 2014 – 2015
	<i>Teaching Assistant</i> , Stanford University CS 193C: Client-Side Internet Technologies	Summer 2013 – 2014
	<i>Teaching Assistant</i> , Stanford University CS 431: High-Level Vision - Behaviors, Neurons and Computational Models	Spring 2013 – 2014
EMPLOYMENT	<i>Research Intern</i> DeepMind Technologies Ltd., London, England, UK	Jun – Sept 2017
	<ul style="list-style-type: none"> • Deep reinforcement learning for vision-based robot manipulation 	
	<i>Research Intern</i> Allen Institute for Artificial Intelligence, Seattle, WA, USA	Jun – Sept 2016
	<ul style="list-style-type: none"> • Building photorealistic simulated 3D environments for visual AI and deep reinforcement learning models for visual navigation 	
	<i>Research Intern</i> Snap Inc., Venice, CA, USA	Jun – Sept 2016
	<ul style="list-style-type: none"> • Developing deep learning models for video understanding at scale in the Snap Research team 	
	<i>Software Engineer Intern</i> Twitter Inc., San Francisco, CA, USA	Apr – July 2013
	<ul style="list-style-type: none"> • Growth hacking in the Activation & Messaging team for Twitter user retention 	
	<i>Research Assistant</i> SFU Computational Logic Lab, Burnaby, BC, Canada	Dec 2011 – Apr 2013
<ul style="list-style-type: none"> • Research in statistical relational learning of efficient learning and inference with Bayesian Networks and Markov Logic Networks 		
<i>Research Assistant</i> SFU Vision and Media Lab, Burnaby, BC, Canada	Jan 2012 – Apr 2013	
<ul style="list-style-type: none"> • Research in action understanding of nursing home videos for Technology for Injury Prevention in Seniors (TIPS) program 		
<i>Software Engineer Intern</i> Qingdao Topcomm Communication Co., Ltd., Qingdao, China	July – Aug 2011	
<ul style="list-style-type: none"> • Developing router testing framework for electric energy data acquisition system 		
HONORS & AWARDS	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

Awards and Prizes

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

SELECTED PRESS COVERAGE

- [1] "Robot See, Robot Do: Bots Learn by Watching Human Behavior," by Noah Kravitz, *Nvidia Blog*. April 3, 2018.
- [2] "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.
- [4] "AI2-THOR Interactive Simulation Teaches AI About Real World," by Jeremy Hsu, *IEEE Spectrum*. Feb 15, 2018.
- [5] "Next Big Test for AI: Making Sense of the World," by Will Knight, *MIT Technology Review*. Jan 26, 2016.

MENTORING

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: Sr. Machine Learning Scientist at Tesla
- Jiren Zhu (Math, Stanford)
- Zihua Liu (CS, Stanford)
- Justin Rose (CS, Stanford)
- David Lin (Physics, Stanford)

Master students

- Andrey Kurenkov (CS, Stanford) Next: Ph.D. student at Stanford
- Julian Gao (CS, Stanford)
- Alex Kaiyi Fu (CS, Stanford)

Ph.D. students

- Danfei Xu (CS, Stanford)
- Kuan Fang (EE, Stanford)
- Jim Fan (CS, Stanford)
- Ajay Mandlekar (EE, Stanford)
- Michelle Lee (ME, Stanford)

Independent research

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