

Fei Xia

Room 126, Gates Building, 353 Serra Mall, Stanford, CA 94305-9515 ◊ feixia@stanford.edu ◊ fxia.me

RESEARCH STATEMENT

My research interests lie in **Computer Vision** and **Robotics**. In particular, I am interested in simulation to real world transfer and domain adaptation for vision and robotics tasks. I am also interested in 3d vision and deep learning methods for point cloud and meshes, along with its application to robotics.

EDUCATION

- Stanford University**, Stanford, CA, USA 2016.9 - 2021.9
PhD Candidate, Department of Electrical Engineering, Advisor: Silvio Savarese and Leo Guibas
GPA: 4.2/4.0, Coursework: Machine Learning, NLP with Deep Learning, Machine Learning for 3D Data, Representation Learning in Computer Vision, Deep Learning for Genomics, Probabilistic Graphical Models.
- Stanford University**, Stanford, CA, USA 2016.9 - 2019.4
MS in Department of Electrical Engineering. GPA: 4.2/4.0.
- Tsinghua University**, Beijing, China 2012.8 - 2016.7
Bachelor of Engineering, Department of Automation
- Georgia Institute of Technology**, Atlanta, GA, USA 2014.8 - 2014.12
Exchange Student in the School of Electrical and Computer Engineering

PUBLICATIONS AND MANUSCRIPTS

see [my google scholar](#) for an always up-to-date list.

- [1] **Fei Xia***, Zhiyang He*, Amir Zamir*, Sasha Sax, Jitendra Malik, Silvio Savarese. **Embodied Real-World Active Perception**. *CVPR 2018 (spotlight oral, Nvidia Pioneering Research Award)*. (*indicates equal contributions, same for the rest)
- [2] **Fei Xia**, William B. Shen, Chengshu Li, Priya Kasimbeg, Alexander Toshev, Roberto Martín-Martín, Silvio Savarese. **Interactive Gibson:A Benchmark for Interactive Navigation in Cluttered Environments**. *RA-L and ICRA20*.
- [3] William B. Shen*, **Fei Xia***, Chengshu Li*, Roberto Martín-Martín*, et al. **iGibson 1.0: a Simulation Environment for Interactive Tasks in Large Realistic Scenes**. *IROS 2021*.
- [4] **Fei Xia***, Chengshu Li*, Roberto Martín-Martín, Or Litany, Alexander Toshev, Silvio Savarese. **ReLMoGen: Leveraging Motion Generation in Reinforcement Learning for Mobile Manipulation**. *RSS workshop, ICRA 2021*.
- [5] Panos Achlioptas, Ahmed Abdelreheem, **Fei Xia**, Mohamed Elhoseiny, Leonidas Guibas. **ReferIt3D: Neural Listeners for Fine-Grained 3D Object Identification in Real-World Scenes**. *ECCV 2020, oral*
- [6] Chengshu Li, **Fei Xia**, Roberto Martín-Martín, Silvio Savarese. **HRL4IN: Hierarchical Reinforcement Learning for InteractiveNavigation with Mobile Manipulators**. *CoRL 2019*.
- [7] Noriaki Hirose, Amir Sadeghian, **Fei Xia**, Silvio Savarese. **VUNet: Traversability Estimation via Dynamic Scene View Synthesis**. *IEEE RA-L, and ICRA 2019*
- [8] Noriaki Hirose, **Fei Xia**, Roberto Martín-Martín, Amir Sadeghian, Silvio Savarese, **Deep Visual MPC-Policy Learning for Navigation**. *IEEE RA-L, and IROS 2019*
- [9] Kevin Chen, Juan Pablo de Vicente, Gabriel Sepulveda, **Fei Xia**, Alvaro Soto, Marynel Vazquez, Silvio Savarese. **A Behavioral Approach to Visual Navigation with Graph Localization Networks**. *RSS 2019*
- [10] Xinke Deng, Arsalan Mousavian, Yu Xiang, **Fei Xia**, Timothy Bretl, Dieter Fox. **PoseRBPF: A Rao-Blackwellized Particle Filter for 6D Object Pose Estimation**. *RSS 2019*
- [11] **Fei Xia**, Chengshu Li, Kevin Chen, Roberto Martin Martin, Noriaki Hirose, Amir R Zamir, Silvio Savarese. **Gibson Env V2: Embodied Simulation Environment for Interactive Navigation**. *CVPRW 2019*
- [12] Martin J. Zhang, **Fei Xia**, James Zou. **AdaFDR: a Fast, Powerful and Covariate-Adaptive Approach to Multiple Hypothesis Testing**. *Best Paper Award at RECOMB 2019*.

- [13] Martin J. Zhang, **Fei Xia**, James Zou. **Fast and covariate-adaptive method amplifies detection power in large-scale multiple hypothesis testing.** *Nature Communications*. *IF=11.9*
- [14] Soheil Feizi, Changho Suh, **Fei Xia** and David Tse. **Understanding GANs: the LQG Setting.**
- [15] **Fei Xia***, Martin Zhang*, James Zou, David Tse. **NeuralFDR: learning decision threshold from hypothesis features.** *NIPS 2017*. (*equal contributions)
- [16] Qiao Liu, **Fei Xia**, Qijin Yin, Rui Jiang. **Chromatin accessibility prediction via a hybrid deep convolutional neural network.** *Bioinformatics*, 2017. *IF=5.48*.
- [17] Govinda Kamath*, Ilan Shomorony*, **Fei Xia***, Thomas Courtade, David Tse. “**HINGE: Long-Read Assembly Achieves Optimal Repeat Resolution.**” *Genome Research* Vol 27 2017. *IF=10.1*(*equal contributions)
- [18] Ilan Shomorony, Govinda Kamath, **Fei Xia**, Thomas Courtade and David Tse, “**Partial DNA Assembly: A Rate-Distortion Perspective.**” *ISIT 2016*.
- [19] **Fei Xia**, et al. “**Human-aware mobile robot exploration and motion planner.**” *Proceeding of IEEE SoutheastCon 2015*.
- [20] Anastasia Dubrovina, **Fei Xia**, Panos Achlioptas, Mira Shalah, Leonidas Guibas. **Composite Shape Modeling via Latent Space Factorization.** *ICCV 2019*.

AWARDS

- 2019** Qualcomm Innovation Fellowship (\$100k)
- 2019** Best Paper Award at RECOMB 2019.
- 2018** Nvidia Pioneering Resaerch Award at CVPR 2018.
- 2016** Stanford Graduate Fellowship (Michael J. Flynn Fellow), Stanford University
- 2015** Chang Jiong Scholarship (Highest honor in Dept. of Automation, Tsinghua University, 1/560)
- 2014** Fang Chongzhi Scholarship (Highest honor in Dept. of Automation, Tsinghua University, 1/560)
- 2013** National Southwest Associated University Scholarship (1/560)

RESEARCH EXPERIENCES

Stanford University, Stanford, CA, USA *2016.12 - Present*
Stanford Vision and Learning Group, Stanford AI Lab

Research Assistant, Advisor: **Prof. Silvio Savarese**, Secondary Advisor: **Prof. Leo Guibas**

Project 1: Gibson Env: Embodied active real-world perception

- Developed Gibson Env, a robotics simulator for easy transfer to real-world. First robotics simulator that enables real-world perception. Used neural network to do real-time rendering for generating photo-realistic video stream.
- Implemented a pixel level domain adaptation mechanism to map real-world images and neural network generated images to a common space for transferring to real world.
- Gibson Env[1] was published in CVPR2018 and won Nvidia Pioneering Research Award. Within a year, more than 10 publications in the robotics community successfully trained navigation policies in Gibson Environment and 3 publications transferred the policy to real robots.

Project 2: Using Gibson for Robot Learning

- Used Gibson Environment for multiple robot learning projects, including Visual Trajectory following [7], Mobile Manipulation [3,5], Interactive Navigation [2], Topological Navigation [8] and demonstrated state of the art performance in those navigation tasks.

Project 3: View Synthesis from a Single RGB Image

- Create an end-to-end deep learning based method where the geometric constraints inherent to the problem (specifically 3D rigid body transformation) are internally enforced.
- Implemented the method and tested on indoor scene datasets and showed the application on robot navigation [6].

Google Inc. , Mountain View, CA, USA
Robotics at Google

2020.6 - 2021.1

Research Intern, Host: **Dr. Alexander Toshev**, **Dr. Brian Ichter**

Main Project: Combining Classical Motion Planning and RL for Navigation and Coarse Manipulation.

- Part of the results are published in [3].

Other Google Engagements Co-organized [CVPR workshop on Embodied AI](#) and [iGibson Challenge](#) with as a Stanford-Google collaboration project.

Nvidia Research, Seattle, WA, USA
Seattle Robotics Lab

2018.6-2018.9

Research Intern, Manager: **Prof. Dieter Fox**

Intuitive Physics Modelling for Real-World Object Interactions

- Developed methods to predict pose changes after physical interactions for real world objects.

Fast Rendering for Doing Render-and-Compare on Particles

- Developed a fast rendering engine that supports CUDA-OpenGL interoperation that enables rendering of a large set of images concurrently and doing rendering-and-compare with real images. Details in [7].

Stanford University, Stanford, CA, USA

2015.7 - 2016.12

Information Systems Laboratory, Department of Electrical Engineering

Research Assistant, Advisor: **Prof. David Tse**

Project 1: HINGE: A *de novo* Sparse String Graph Assembler for PacBio Reads

- Generated **finished** assembly at accuracy 99.9% for *E.Coli* based on sparse string graph methods, with details in publication [16-17].
- Extended NSG(Not-So-Greedy) algorithm to a regime when triple repeats are all-bridged and interleaved repeats are bridged, i.e. information-theoretic bound for perfect assembly.

Project 2: NeuralFDR: learning decision threshold from hypothesis features

- Proposed a learning based method for FDR control. Developed mirroring method for FDP prediction. NeuralFDR has provable performance in FDP control.
- Implemented the method and tested on RNASeq and GWAS datasets. Details can be found in paper [14].
- Our follow-up work of NeuralFDR, which replaced the neural network with a mixture of gaussian [11], showed better results and won Best Paper Award in RECOMB 2019.

Megvii Inc., Beijing, China

2016.3-2016.7

DTR(Detection, Tracking, Re-identification) Group

Research Intern, Mentor: **Chi Zhang, Chief Scientist**

Project 1: Pedestrian Parsing Models

- Built an deep convolutional neural network model based on Holistically-Nested Edge Detection model and adapted it for pedestrian parsing.

Project 2: Pixel Level Domain Transfer for Pedestrian Re-identification

- Built a generative adversarial network that transfer from pedestrian domain to upper-cloth domain, and used that model for pedestrian re-identification.
- Both models were incorporated into company's API for downstream applications.

TECHNICAL STRENGTHS

Deep Learning Software Stacks
Programming Languages
Additional Skills

PyTorch, Tensorflow, TF & PyTorch CUDA module development
 Proficient with C/C++, Python, MATLAB, Java
 ROS, MPI, OpenMP, CUDA

last updated: 6/22/2021