BOLEI ZHOU

Associate Professor
Computer Science Department
University of California, Los Angeles (UCLA)
Email: bolei@cs.ucla.edu | Homepage | Lab Page | Google Scholar

Academic Positions

University of California, Los Angeles (UCLA) Jan 2022 - now
Computer Science Department
Computational Medicine Department (Affiliated)
- Associate Professor (July 2025 - now)
- Assistant Professor (Jan 2022 - June 2025)

The Chinese University of Hong Kong Oct.2018 - Dec.2021
Department of Information Engineering
Department of Computer Science and Engineering (by courtesy)
- Assistant Professor

Education

Massachusetts Institute of Technology Aug 2013 – May 2018
Ph.D. in Computer Science
- Thesis: Interpretable Representation Learning for Visual Intelligence
- Thesis Committee: Antonio Torralba (Advisor), Aude Oliva, William T. Freeman

The Chinese University of Hong Kong Aug 2010 – July 2012
M.Phil. in Information Engineering
- Thesis: Modeling Collective Crowd Behaviors in Videos
- Advisor: Xiaoou Tang

Shanghai Jiao Tong University, China Aug 2006 – July 2010
B.Eng. in Biomedical Engineering
- Thesis: Motion Perception Model and its Application

Awards and Honors

Office of Naval Research Young Investigator Award 2025
NSF CAREER Award 2024
Intel Rising Star Faculty Award 2023
MIT Technology Review's Innovators Under 35 in Asia-Pacific Region 2020
The World Artificial Intelligence Conference (WAIC) Yunfun Star Award 2020
Facebook Ph.D. Fellowship in Computer Vision 2016-2018
BRC Fellowship Award 2017
MIT Ho-Ching and Han-Ching Fund Award 2013
MIT Greater China Computer Science Fellowship 2013
Faculty's Outstanding Thesis Award in CUHK 2012
Microsoft Research Asia Fellowship 2011

Press Coverage

UCLA Faculty Bulletin Board: Computer science professor receives NSF award for AI that interacts with real world. 2024
UCLA News: UCLA Computer Scientist Leads NSF-Funded Effort to Build Open-Source Autonomous Driving Simulation. 2023
MIT Technology Review: Innovators Under 35: Making AI models more understandable. 2020

CUHK News: CUHK Faculty Listed as One of the 20 Innovators Under 35 Asia Pacific. 2020
MIT News: Visualizing an AI model's blind spots. 2019
MIT Technology Review: A neural network can learn to organize the world it sees into concepts just like we do. 2019
MIT News: Teaching artificial intelligence to create visuals with more common sense. 2019
Neurohive: Dissecting GANs for Better Understanding and Visualization. 2018
VentureBeat: Designing AI that can track objects over time. 2018
MIT News: Helping computers fill in the gaps between video frames 2018
Quartz: Track AI's decisions back to single neurons 2017
MIT News: Peering into neural networks 2017
TechCrunch: A fully automated way to peer inside neural nets 2017
MIT CSAIL News: Scene parsing and scene classification challenges 2016
TechCrunch: AI Project designed to recognize scenes identifies objects too 2015
MIT News: Object recognition for free 2015

Journal Publications

1. Oren Avram, Berkin Durmus, Nadav Rakocz, Giulia Corradetti, Ulzee An, Muneeswar G Nittala, Prerit Terway, Akos Rudas, Zeyuan Johnson Chen, Yu Wakatsuki, and others. Accurate prediction of disease-risk factors from volumetric medical scans by a deep vision model pre-trained with 2D scans. Nature Biomedical Engineering (Nature BME), 2024

2. Jiankai Sun, Linjiang Huang, Hongsong Wang, Chuanyang Zheng, Jianing Qiu, Md Tauhidul Islam, Enze Xie, Bolei Zhou, Lei Xing, Arjun Chandrasekaran, and others. Localization and recognition of human action in 3D using transformers. Communications Engineering, 2024

3. Jianyuan Wang, Lalit Bhagat, Ceyuan Yang, Yinghao Xu, Yujun Shen, Hongdong Li, and Bolei Zhou. Spatial Steerability of GANs via Self-Supervision from Discriminator. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2024

4. Shuzhen Luo, Menghan Jiang, Sainan Zhang, Junxi Zhu, Shuangyue Yu, Israel Dominguez Silva, Tian Wang, Elliott Rouse, **Bolei Zhou**, Hyunwoo Yuk, Xianlian Zhou, Hao Su. "Experiment-free exoskeleton optimization via learning in simulation." Nature, 2024

5. Alexander Swerdlow, Runsheng Xu, and **Bolei Zhou**. "Street-View Image Generation from a Bird's-Eye View Layout." IEEE Robotics and Automation Letters (RAL), 2024

6. Shengchao Liu, Chengpeng Wang, Jiarui Lu, Weili Nie, Hanchen Wang, Zhuoxinran Li, **Bolei Zhou**, and Jian Tang. "Unsupervised Discovery of Steerable Factors When Graph Deep Generative Models Are Entangled." Transactions on Machine Learning Research (TMLR), 2024

7. Jiapeng Zhu, Yujun Shen, Yinghao Xu, Deli Zhao, Qifeng Chen, **Bolei Zhou**. "In-Domain GAN Inversion for Faithful Reconstruction and Editability." IEEE Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), 2024

8. Yuanqi Du, Xian Liu, Nilay Shah, Shengchao Liu, Jieyu Zhang, and **Bolei Zhou**. "ChemSpacE: Interpretable and Interactive Chemical Space Exploration." Transactions on Machine Learning Research (TMLR), 2023

9. Yinghao Xu, Yujun Shen, Jiapeng Zhu, Ceyuan Yang, and **Bolei Zhou**. "GH-Feat: Learning Versatile Generative Hierarchical Features from GANs." IEEE Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), 2022

10. Quanyi Li*, Zhenghao Peng*, Lan Feng, Qihang Zhang, Zhenghai Xue, and **Bolei Zhou**. "MetaDrive: Composing Diverse Driving Scenarios for Generalizable Reinforcement Learning." IEEE Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), 2022

11. Jiapeng Zhu, Deli Zhao, Bo Zhang, and **Bolei Zhou**. "Disentangled inference for gans with latently invertible autoencoder." International Journal of Computer Vision (**IJCV**), 2022

12. Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, **Bolei Zhou**, and Ming-Hsuan Yang. "GAN inversion: A survey" IEEE Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), 2022

13. Jimmy Wu, **Bolei Zhou**, Diondra Peck, Scott Hsieh, Vandana Dialani, Lester Mackey, and Genevieve Patterson. "DeepMiner: Discovering Interpretable Representations for Mammogram Classification and Explanation" Harvard Data Science Review, May 2021. [link]

14. Yujun Shen, Ceyuan Yang, Xiaoou Tang, **Bolei Zhou**. "InterFaceGAN: Interpreting the Disentangled Face Representation Learned by GANs" IEEE Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), Oct 2020. [webpage]

15. Ceyuan Yang, Yujun Shen, **Bolei Zhou**. "Semantic Hierarchy Emerges in Deep Generative Representations for Scene Synthesis" International Journal of Computer Vision (**IJCV**), Dec 2020. [webpage]

16. David Bau, Jun-Yan Zhu, Henrick Strobelt, Agata Lapedriza, **Bolei Zhou** and Antonio Torralba. "Understanding the Role of Individual Units in a Deep Neural Network." Proceedings of the National Academy of Sciences of the United States of America (**PNAS**), 2020.

17. Jiankai Sun, Lantao Yu, Pinqian Dong, Bo Lu, **Bolei Zhou**. "Adversarial Inverse Reinforcement Learning with Self-attention Dynamics Model." IEEE Robotics and Automation Letters (**RA-L**) and **ICRA**, 2021.

18. Bowen Pan, Jiankai Sun, Ho Yin Tiga Leung, Alex Andonian, **Bolei Zhou**. "Cross-view Semantic Segmentation for Sensing Surroundings." IEEE Robotics and Automation Letters (**RA-L**) and **IROS**, 2020. [webpage]

19. D. Bau, H. Strobelt, W. Peebles, J. Wulff, **B. Zhou**, J. Zhu, A. Torralba. "Semantic Photo Manipulation with a Generative Image Prior." ACM Transactions on Graphics (**TOG**) and **SIGGRAPH**, 2019

20. **B. Zhou**, H. Zhao, X. Puig, T. Xiao, S. Fidler, A. Barriuso and A. Torralba. "Semantic Understanding of Scenes through ADE20K Dataset." International Journal of Computer Vision (**IJCV**), 2018.

21. **B. Zhou***, D. Bau*, A. Oliva, A. Torralba. "Interpreting Deep Visual Representations via Network Dissection." IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2018
[webpage]

22. M.Monfort, A. Andonian, **B. Zhou**, S. Bargal, T. Yan, K. Ramakrishnan, L. Brown, Q. Fan, D. Gutfruend, C. Vondrick, A. Oliva. "Moments in Time Dataset: One Million Videos for Event Understanding." IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2018
[webpage]

23. F. Zhang, **B. Zhou**, C. Ratti, Y. Liu. "Discovering place-informative scenes and objects using social media photos." Royal Society Open Science, 2018

24. F. Zhang, **B. Zhou**, L.Liu, Y. Liu, HH. Fung, H. Lin, C. Ratti. "Measuring human perceptions of a large-scale urban region using machine learning." Landscape and Urban Planning 180, 148-160

25. **B. Zhou**, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba. "Places: A 10 million Image Database for Scene Recognition." IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2017.
[Places dataset] [PlacesCNN] [demo]

26. **B. Zhou**, X. Tang, H. Zhang and X. Wang. "Measuring Crowd Collectiveness." IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2014.
[webpage]

27. **B. Zhou**, X. Tang and X. Wang. "Learning Collective Crowd Behaviors with Dynamic Pedestrian-Agents." International Journal of Computer Vision (**IJCV**), 2014.
[webpage]

28. L.Liu, **B. Zhou**, J. Zhao, B.D.Ryan. "C-IMAGE: City Cognitive Mapping through Geo-tagged Photos." GeoJournal, Springer, 2016.

Book Chapters

29. **Bolei Zhou**. "Interpreting Generative Adversarial Networks for Interactive Image Generation." xxAI - Beyond explainable Artificial Intelligence, Springer, 2022

30. **Bolei Zhou**. "Scene Classification." Computer Vision: A Reference Guide, Editor-in-chief: Ikeuchi, Katsushi, Springer, 2020.

31. **Bolei Zhou**, D. Bau, A. Oliva, A. Torralba. "Comparing the Interpretability of Deep Networks via Network Dissection." Explainable AI: Interpreting, Explaining and Visualizing Deep Learning, Springer, 2019.

Conference Publications

32. Honglin He, Yukai Ma, Brad Squicciarini, Wayne Wu, and Bolei Zhou. Learning Sidewalk Autopilot from Multi-Scale Imitation with Corrective Behavior Expansion. In IEEE International Conference on Robotics and Automation (ICRA), 2026

33. Yukai Ma, Honglin He, Selina Song, Wayne Wu, and Bolei Zhou. AURA: Multi-modal Shared Autonomy for Urban Navigation. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2026

34. Sicheng Mo, Thao Nguyen, Richard Zhang, Nicholas Kolkin, Siddharth Srinivasan Iyer, Eli Shechtman, Krishna Kumar Singh, Yong Jae Lee, Bolei Zhou, and Yuheng Li. Group Diffusion: Enhancing Image Generation by Unlocking Cross-Sample Collaboration. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2026

35. Tianyi Xie, Yunuo Chen, Yaowei Guo, Yin Yang, Bolei Zhou, Demetri Terzopoulos, Ying Jiang, and Chenfanfu Jiang. AnimaMimic: Imitating 3D Animation from Video Priors. In IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2026

36. Andy Wang, Xu Yan, Brandon McMahan, Michael Zhou, Yuyang Yuan, Johannes Y. Lee, Ali Shreif, Matthew Li, Zhenghao Peng, Bolei Zhou, Yuchen Cui, and Jonathan C. Kao. DiSCo: Diffusion Sequence Copilots for Shared Autonomy. In Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (ICLR), 2026

37. Mingxuan Liu, Honglin He, Elisa Ricci, Wayne Wu, and Bolei Zhou. UrbanVerse: Scaling Urban Simulation by Watching City-Tour Videos. In International Conference on Learning Representation (ICLR), 2026

38. Honglin He, Yukai Ma, Brad Squicciarini, Wayne Wu, and Bolei Zhou. From Seeing to Experiencing: Scaling Navigation Foundation Models with Reinforcement Learning. In International Conference on Learning Representation (ICLR), 2026

39. Zhizheng Liu, Joe Lin, Wayne Wu, and Bolei Zhou. Joint Optimization for 4D Human-Scene Reconstruction in the Wild. In International Conference on Learning Representation (ICLR), 2026

40. Zhenghao Peng, Yuxin Liu, and Bolei Zhou. SceneStreamer: Continuous Scenario Generation as Next Token Group Prediction. In International Conference on Learning Representation (ICLR), 2026

41. Haoyuan Cai, Zhenghao Peng, and Bolei Zhou. Predictive Preference Learning from Human Interventions. In Advances in Neural Information Processing Systems (NeurIPS), 2025, Spotlight

42. Yuxin Liu, Zhenghao Peng, Xuanhao Cui, and Bolei Zhou. Adv-BMT: Bidirectional Motion Transformer for Safety-Critical Traffic Scenario Generation. In Advances in Neural Information Processing Systems (NeurIPS), 2025

43. Zewei Zhou, Tianhui Cai, Seth Z Zhao, Yun Zhang, Zhiyu Huang, Bolei Zhou, and Jiaqi Ma. AutoVLA: A Vision-Language-Action Model for End-to-End Autonomous Driving with Adaptive Reasoning and Reinforcement Fine-Tuning. In Advances in Neural Information Processing Systems (NeurIPS), 2025

44. Sicheng Mo, Thao Nguyen, Xun Huang, Siddharth Srinivasan Iyer, Yijun Li, Yuchen Liu, Abhishek Tandon, Eli Shechtman, Krishna Kumar Singh, Yong Jae Lee, Bolei Zhou, and Yuheng Li. X-Fusion: Introducing New Modality to Frozen Large Language Models. International Conference on Computer Vision (ICCV), 2025

45. Ziyang Leng, Jiawei Yang, Wenlong Yi, and Bolei Zhou. Occupancy Learning with Spatiotemporal Memory. International Conference on Computer Vision (ICCV), 2025

46. Zewei Zhou, Hao Xiang, Zhaoliang Zheng, Seth Z. Zhao, Mingyue Lei, Yun Zhang, Tianhui Cai, Xinyi Liu, Johnson Liu, Maheswari Bajji, Xin Xia, Zhiyu Huang, Bolei Zhou, Jiaqi Ma. V2XPnP: Vehicle-to-Everything Spatio-Temporal Fusion for Multi-Agent Perception and Prediction. International Conference on Computer Vision (ICCV), 2025

47. Cheng-Fu Yang*, Da Yin*, Wenbo Hu, Heng Ji, Nanyun Peng, Bolei Zhou, Kai-Wei Chang. Verbalized Representation Learning for Interpretable Few-Shot Generalization. International Conference on Computer Vision (ICCV), 2025

48. Haoyuan Cai, Zhenghao Peng, and Bolei Zhou. AIM: Robot-Gated Interactive Imitation Learning with Adaptive Intervention Mechanism. International Conference on Machine Learning (ICML), 2025

49. Ziyang Xie, Zhizheng Liu, Zhenghao Peng, Wayne Wu, and Bolei Zhou. Vid2Sim: Realistic and Interactive Simulation from Video for Urban Navigation. IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR), 2025

50. Weizhen Wang, Chenda Duan, Zhenghao Peng, Yuxin Liu, and Bolei Zhou. Embodied Scene Understanding for Vision Language Models via MetaVQA. IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR), 2025

51. Wayne Wu*, Honglin He*, Chaoyuan Zhang, Jack He, Seth Z. Zhao, Ran Gong, Quanyi Li, and Bolei Zhou. Towards Autonomous Micromobility through Scalable Urban Simulation. IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR), 2025 , Highlight

52. Zhenghao Peng, Zhizheng Liu, and Bolei Zhou. Data-Efficient Learning from Human Interventions for Mobile Robots. IEEE International Conference on Robotics and Automation (ICRA), 2025

53. Wayne Wu, Honglin He, Jack He, Yiran Wang, Chenda Duan, Zhizheng Liu, Quanyi Li, and Bolei Zhou. MetaUrban: An Embodied AI Simulation Platform for Urban Micromobility. International Conference on Learning Representations (ICLR), 2025, Spotlight

54. Zhizheng Liu, Joe Lin, Wayne Wu, and Bolei Zhou. Learning to Generate Diverse Pedestrian Movements from Web Videos with Noisy Labels. International Conference on Learning Representations (ICLR), 2025

55. Qihang Zhang, Yinghao Xu, Chaoyang Wang, Hsin-Ying Lee, Gordon Wetzstein, Bolei Zhou, and Ceyuan Yang. 3DitScene: Editing Any Scene via Language-guided Disentangled Gaussian Splatting. International Conference on Learning Representations (ICLR), 2025

56. Yunsong Zhou, Michael Simon, Zhenghao Peng, Sicheng Mo, Hongzi Zhu, Minyi Guo, and Bolei Zhou. SimGen: Simulator-conditioned Driving Scene Generation. Advances in Neural Information Processing Systems (NeurIPS), 2024

57. Kuan Heng Lin, Sicheng Mo, Ben Klingher, Fangzhou Mu, and Bolei Zhou. Ctrl-X: Controlling Structure and Appearance for Text-To-Image Generation Without Guidance. Advances in Neural Information Processing Systems (NeurIPS), 2024

58. Brandon J McMahan, Zhenghao Peng, Bolei Zhou, and Jonathan C Kao. Shared Autonomy with IDA: Interventional Diffusion Assistance. Advances in Neural Information Processing Systems (NeurIPS), 2024

59. Sicheng Mo, Fangzhou Mu, Kuan Heng Lin, Yanli Liu, Bochen Guan, Yin Li, and **Bolei Zhou**. "FreeControl: Training-Free Spatial Control of Any Text-to-Image Diffusion Model with Any Condition." IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR), 2024.

60. Qihang Zhang, Chaoyang Wang, Aliaksandr Siarohin, Peiye Zhuang, Yinghao Xu, Ceyuan Yang, Dahua Lin, **Bolei Zhou**, Sergey Tulyakov, and Hsin-Ying Lee. "Scenewiz3d: Towards text-guided 3d scene composition." IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR) , 2024.

61. Qihang Zhang, Yinghao Xu, Yujun Shen, Bo Dai, **Bolei Zhou**, and Ceyuan Yang. "Berfscene: Bev-conditioned equivariant radiance fields for infinite 3d scene generation." IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR), 2024.

62. Yinghao Xu, Wang Yifan, Alexander W Bergman, Menglei Chai, **Bolei Zhou**, and Gordon Wetzstein. "Efficient 3d articulated human generation with layered surface volumes." The 11th International Conference on 3D Vision (3DV) , 2024, Spotlight

63. Zhenghao Peng, Wenjie Mo, Chenda Duan, Quanyi Li, and **Bolei Zhou**. "Learning from Active Human Involvement through Proxy Value Propagation." In Neural Information Processing Systems (**NeurIPS**), 2023, Spotlight

64. Quanyi Li, Zhenghao Peng, Lan Feng, Zhizheng Liu, Chenda Duan, Wenjie Mo, and **Bolei Zhou**. "ScenarioNet: Open-Source Platform for Large-Scale Traffic Scenario Simulation and Modeling." Neural Information Processing Systems, Datasets and Benchmarks (**NeurIPS**), 2023

65. Linrui Zhang, Zhenghao Peng, Quanyi Li, **Bolei Zhou**. "CAT: Closed-loop Adversarial Training for Safe End-to-End Driving ." 7th Annual Conference on Robot Learning (**CoRL**), 2023

66. Ceyuan Yang, Yujun Shen, Zhiyi Zhang, Yinghao Xu, Jiapeng Zhu, Zhirong Wu, and **Bolei Zhou**. "One-shot generative domain adaptation." International Conference on Computer Vision (**ICCV**), 2023

67. Yinghao Xu, Menglei Chai, Zifan Shi, Sida Peng, Ivan Skorokhodov, Aliaksandr Siarohin, Ceyuan Yang, Yujun Shen, Hsin-Ying Lee, **Bolei Zhou**, and Sergey Tulyakov. "DisCoScene: Spatially Disentangled Generative Radiance Fields for Controllable 3D-aware Scene Synthesis." IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), 2023, **Spotlight**

68. Runsheng Xu, Xin Xia, Jinlong Li, Hanzhao Li, Shuo Zhang, Zhengzhong Tu, Zonglin Meng, Hao Xiang, Xiaoyu Dong, Rui Song, Hongkai Yu, **Bolei Zhou**, and Jiaqi Ma. "V2V4Real: A large-scale real-world dataset for Vehicle-to-Vehicle Cooperative Perception." IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), 2023, **Spotlight**

69. Zhenghai Xue, Zhenghao Peng, Quanyi Li, Zhihan Liu, and **Bolei Zhou**. "Guarded Policy Optimization with Imperfect Online Demonstrations." International Conference on Learning Representations (**ICLR**), 2023, **Spotlight**

70. Qihang Zhang, Ceyuan Yang, Yujun Shen, Yinghao Xu, and **Bolei Zhou**. "Towards Smooth Video Composition." International Conference on Learning Representations (**ICLR**), 2023

71. Lan Feng, Quanyi Li, Zhenghao Peng, Shuhan Tan, and **Bolei Zhou**. "TrafficGen: Learning to Generate Diverse and Realistic Traffic Scenarios." IEEE International Conference on Robotics and Automation (**ICRA**), 2023

72. Hao Xiang, Runsheng Xu, Xin Xia, Zhaoliang Zheng, **Bolei Zhou**, and Jiaqi Ma. "V2XP-ASG: Generating Adversarial Scenes for Vehicle-to-Everything Perception." IEEE International Conference on Robotics and Automation (**ICRA**), 2023

73. Quanyi Li, Zhenghao Peng, Haibin Wu, Lan Feng, and **Bolei Zhou**. "Human-AI Shared Control via Policy Dissection." Neural Information Processing Systems (**NeurIPS**), 2022

74. Hao Sun, Lei Han, Rui Yang, Xiaoteng Ma, Jian Guo, and **Bolei Zhou**. "Optimistic Curiosity Exploration and Conservative Exploitation with Linear Reward Shaping." Neural Information Processing Systems (**NeurIPS**), 2022

75. Runsheng Xu, Zhengzhong Tu, Hao Xiang, Wei Shao, **Bolei Zhou**, and Jiaqi Ma. "CoBEVT: Cooperative Bird's Eye View Semantic Segmentation with Sparse Transformers." Conference on Robot Learning (**CoRL**), 2022

76. Ceyuan Yang, Yujun Shen, Yinghao Xu, Deli Zhao, Bo Dai, and **Bolei Zhou**. "Improving GANs with A Dynamic Discriminator." Neural Information Processing Systems (**NeurIPS**), 2022

77. Qihang Zhang, Zhenghao Peng, and **Bolei Zhou**. "Learning to Drive by Watching YouTube Videos: Action-Conditioned Contrastive Policy Pretraining." European Conference on Computer Vision (**ECCV**), 2022

78. Xian Liu, Yinghao Xu, Qianyi Wu, Hang Zhou, Wayne Wu, and **Bolei Zhou**. "Semantic-Aware Implicit Neural Audio-Driven Video Portrait Generation." European Conference on Computer Vision (**ECCV**), 2022, **Oral**

79. Jianyuan Wang, Ceyuan Yang, Yinghao Xu, Yujun Shen, Hongdong Li, and **Bolei Zhou**. "Improving GAN Equilibrium by Raising Spatial Awareness." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), 2022

80. Yinghao Xu, Sida Peng, Ceyuan Yang, Yujun Shen, and **Bolei Zhou**. "3D-aware Image Synthesis via Learning Structural and Textural Representations." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), 2022

81. Xian Liu, Qianyi Wu, Hang Zhou, Yinghao Xu, Rui Qian, Xinyi Lin, Xiaowei Zhou, Wayne Wu, Bo Dai, and **Bolei Zhou**. "Learning Hierarchical Cross-Modal Association for Co-Speech Gesture Generation." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), 2022

82. Yinghao Xu, Fangyun Wei, Xiao Sun, Ceyuan Yang, Yujun Shen, Bo Dai, **Bolei Zhou**, and Stephen Lin. "Cross-Model Pseudo-Labeling for Semi-Supervised Action Recognition." In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**), 2022, **Oral**

83. Quanyi Li, Zhenghao Peng, and **Bolei Zhou**. "Efficient Learning of Safe Driving Policy via Human-AI Copilot Optimization." International Conference on Learning Representations (**ICLR**) , 2022

84. Jiankai Sun, De-An Huang, Bo Lu, Yunhui Liu, **Bolei Zhou**, and Animesh Garg. "PlaTe: Visually-Grounded Planning with Transformers in Procedural Tasks." IEEE Robotics and Automation Letters (**ICRA+RAL**), 2022

85. Xian Liu, Rui Qian, Hang Zhou, Di Hu, Weiyao Lin, Ziwei Liu, **Bolei Zhou**, and Xiaowei Zhou. "Visual Sound Localization in the Wild by Cross-Modal Interference Erasing." AAAI Conference on Artificial Intelligence (**AAAI**), 2022

86. Zhenyu Li, Zehui Chen, Ang Li, Liangji Fang, Qinhong Jiang, Xianming Liu, Junjun Jiang, **Bolei Zhou**, and Hang Zhao. "SimIPU: Simple 2D Image and 3D Point Cloud Unsupervised Pre-Training for Spatial-Aware Visual Representations." AAAI Conference on Artificial Intelligence (**AAAI**), 2022

87. Zhenghao Peng, Quanyi Li, Chunxiao Liu, and **Bolei Zhou**. "Learning to Simulate Self-Driven Particles System with Coordinated Policy Optimization." Neural Information Processing Systems (**NeurIPS**), 2021. [webpage]

88. Ceyuan Yang, Yujun Shen, Yinghao Xu, and **Bolei Zhou**. "Data-Efficient Instance Generation from Instance Discrimination." Neural Information Processing Systems (**NeurIPS**), 2021. [webpage]

89. Zhenghao Peng*, Quanyi Li*, Chunxiao Liu, and **Bolei Zhou**. "Safe Driving via Expert Guided Policy Optimization." Conference on Robot Learning (**CoRL**), 2021. [webpage]

90. Yujun Shen and **Bolei Zhou**. "Closed-Form Factorization of Latent Semantics in GANs." Computer Vision and Pattern Recognition (**CVPR**), 2021, **Oral**.

91. Yinghao Xu*, Yujun Shen*, Jiapeng Zhu, Ceyuan Yang, **Bolei Zhou**. "Generative Hierarchical Features from Synthesizing Images ." Computer Vision and Pattern Recognition (**CVPR**), 2021, **Oral**.

92. Yicheng Liu, Jinghuai Zhang, Liangji Fang, Qinhong Jiang, **Bolei Zhou**. "Multimodal Motion Prediction with Stacked Transformers." Computer Vision and Pattern Recognition (**CVPR**), 2021.

93. Rui Xu, Xintao Wang, Kai Chen, **Bolei Zhou**, Chen Change Loy. "Positional Encoding as Spatial Inductive Bias in GANs." Computer Vision and Pattern Recognition (**CVPR**), 2021.

94. Ceyuan Yang, Zhirong Wu, **Bolei Zhou**, Stephen Lin. "Instance Localization for Self-supervised Detection Pretraining." Computer Vision and Pattern Recognition (**CVPR**), 2021.

95. Jiankai Sun, Rui Liu, **Bolei Zhou**. "HiABP: Hierarchical Initialized ABP for Unsupervised Representation Learning." **AAAI**, 2021.

96. Junning Huang, Sirui Xie, Jiankai Sun, Qiurui Ma, Chunxiao Liu, Dahua Lin, **Bolei Zhou**. "Learning a Decision Module by Imitating Driver's Control Behaviors." The Conference on Robot Learning (**CoRL**), 2020. [PDF]

97. Jiankai Sun, Hao Sun, Tian Han, **Bolei Zhou**. "Neural-Symbolic Program Search: Towards Automatic Autonomous Driving System Design." The Conference on Robot Learning (**CoRL**), 2020.[PDF]

98. Jiapeng Zhu, Yujun Shen, Deli Zhao, **Bolei Zhou**. "In-Domain GAN Inversion for Real Image Editing ." European Conference on Computer Vision (**ECCV**), 2020. [webpage]

99. Anyi Rao, Jiaze Wang, Linning Xu, Xuekun Jiang, Qingqiu Huang, **Bolei Zhou**, Dahua Lin. "A Unified Framework for Shot Type Classification Based on Subject Centric Lens." European Conference on Computer Vision (**ECCV**), 2020.

100. Yujun Shen, Jinjin Gu, Xiaoou Tang, **Bolei Zhou**. "Interpreting Latent Space of GANs for Semantic Face Editing." Computer Vision and Pattern Recognition (**CVPR**), 2020. [webpage]

101. Jinjin Gu, Yujun Shen, **Bolei Zhou**. "Image Processing Using Multi-Code GAN Prior." Computer Vision and Pattern Recognition (**CVPR**), 2020. [webpage]

102. Ceyuan Yang*, Yinghao Xu*, Bo Dai, **Bolei Zhou**. "Temporal Pyramid Network for Action Recognition." Computer Vision and Pattern Recognition (**CVPR**), 2020. [webpage]

103. Liangji Fang*, Qinhong Jiang*, Jianping Shi, **Bolei Zhou**. "Trajectory Proposal Network for Multi-modal Motion Prediction." Computer Vision and Pattern Recognition (**CVPR**), 2020. [webpage]

104. Zhuoqian Yang, Wentao Zhu, Wayne Wu, Chen Qian, Qiang Zhou, **Bolei Zhou**, Chen Change Loy. "Video Motion Retargeting via Invariance-Driven Unsupervised Representation Disentanglement." Computer Vision and Pattern Recognition (**CVPR**), 2020. [webpage]

105. Anyi Rao, Linning Xu, Yu Xiong, Guodong Xu, Qingqiu Huang, **Bolei Zhou**, Dahua Lin. "A Local-to-Global Approach to Multi-modal Movie Scene Segmentation." Computer Vision and Pattern Recognition (**CVPR**), 2020.

106. Mingyu Ding, Zhe Wang, **Bolei Zhou**, Jianping Shi, Zhiwu Lu, Ping Luo. "Every Frame Counts: Joint Learning of Video Segmentation and Optical Flow." **AAAI**, 2020.

107. H. Sun, Z. Li, X. Liu, D. Lin, **B. Zhou**. "Policy Continuation with Hindsight Inverse Dynamics." International Conference on Computer Vision (**NeurIPS**), 2019, **Spotlight** (2% acceptance rate).

108. T. Xiao, Q. Fan, D. Gutfreund, M. Monfort, A. Oliva, **B. Zhou**. "Reasoning about Human-Object Interactions through Dual Attention Networks." International Conference on Computer Vision (**ICCV**), 2019

109. Y. Xiong, Q. Huang, L.g Guo, H. Zhou, **B. Zhou**, D. Lin. "A Graph-based Framework to Bridge Movies and Synopses." International Conference on Computer Vision (**ICCV**), 2019, **Oral**.

110. D. Bau, H. Strobelt, W. Peebles, J. Wulff, J. Zhu, **B. Zhou**, A. Torralba. "Seeing What a GAN Cannot Generate." International Conference on Computer Vision (**ICCV**), 2019, **Oral**.

111. D. Bau, J. Zhu, H. Strobelt, **B. Zhou**, J. B. Tenenbaum, W. T. Freeman, A. Torralba. "GAN Dissection: Visualizing and Understanding Generative Adversarial Networks." International Conference on Representation Learning (**ICLR**), 2019.

112. R. Xu, X. Li, **B. Zhou**, C. C. Loy. "Deep Flow-Guided Video Inpainting." Computer Vision and Pattern Recognition (**CVPR**), 2019

113. G. Yang, X. Song, C. Huang, Z. Deng, J. Shi, **B. Zhou**."DrivingStereo: A large-scale dataset for stereo matching in autonomous driving scenarios." Computer Vision and Pattern Recognition (**CVPR**), 2019

114. **B. Zhou**, A. Andonian, A. Oliva, A. Torralba. "Temporal Relational Reasoning in Videos." European Conference on Computer Vision (**ECCV**), 2018.

115. **B. Zhou***, Y.Sun*, D.Bau*, A. Torralba. "Interpretable Basis Decomposition for Visual Explanation." European Conference on Computer Vision (**ECCV**), 2018.

116. T. Xiao*, Y. Liu*, **B. Zhou***, Y. Jiang, J. Sun. "Unified Perceptual Parsing for Scene Understanding." European Conference on Computer Vision (**ECCV**), 2018.

117. W. Ma, H. Chu, **B. Zhou**, R. Urtasun, A. Torralba. "Single Image Intrinsic Decomposition without a Single Intrinsic Image." European Conference on Computer Vision (**ECCV**), 2018.

118. Y. Li, W. Ouyang, **B. Zhou**, Y. Cui, J. Shi, X. Wang. "Factorizable Net: An Efficient Subgraph based Framework for Scene Graph Generation." European Conference on Computer Vision (**ECCV**), 2018.

119. J. Wu, **B. Zhou**, R. Russell, V. Kee, S. Wagner, M. Hebert, A. Torralba, D. Johnson. "Real-Time Object Pose Estimation with Pose Interpreter Networks." International Conference on Intelligent Robots and Systems (**IROS**), 2018, **Oral**.

120. Y. Li, N. Duan, **B. Zhou**, X. Chu, W. Ouyang, X. Wang. "Visual Question Generation as Dual Task of Visual Question Answering." Computer Vision and Pattern Recognition (**CVPR**), 2018.

121. B. Pan, W. Lin, X. Fang, C. Huang, **B. Zhou**, C. Lu. "Recurrent Residual Module for Fast Inference in Videos." Computer Vision and Pattern Recognition (**CVPR**), 2018.

122. J. Wu, D. Peck, S. Hsieh, V. Dialani, C. Lehman, **B. Zhou**, V. Syrgkanis, L. Mackey, G. Patternson. "Expert Identification of Visual Primitives used by CNNs during Mammogram Classification." SPIE Medical Imaging, 2018.

123. H. Zhao, X. Puig, **B. Zhou**, S. Fidler, and A. Torralba. "Open Vocabulary Scene Parsing." International Conference on Computer Vision (**ICCV**), 2017.

124. Y. Li, W. Ouyang, **B. Zhou**, K. Wang, and X. Wang. "Scene Graph Generation from Objects, Phrases and Region Captions." International Conference on Computer Vision (**ICCV**), 2017.

125. **B. Zhou**, H. Zhao, X. Puig, S. Fidler, A. Barriuso and A. Torralba. "Scene Parsing through ADE20K Dataset." Computer Vision and Pattern Recognition (**CVPR**), 2017. [webpage][ADE dataset]

126. D. Bau*, **B. Zhou***, A. Khosla, A. Oliva. and A. Torralba. "Network Dissection: Quantifying Interpretability of Deep Visual Representations." Computer Vision and Pattern Recognition (**CVPR**) **Oral** (3% acceptance rate), 2017. * indicates equal contribution. [webpage]

127. S. Li, T. Xiao, H. Li, **B. Zhou**, D. Yue, and X. Wang. "Person Search with Natural Language Description." Computer Vision and Pattern Recognition (**CVPR**), 2017.

128. J. Wong, V. Kee, T. Le, S.Wagner, G. Mariottini, A. Schneider, L. Hamilton, R. Chiaplkatty, M. Herbert, D. Johnson, J. Wu, **B. Zhou**, and A. Torralba. "SepICP: Integrated Deep Semantic Segmentation and Pose Estimation."IEEE International Conference on Intelligent Robots and Systems (**IROS**), 2017, **Oral**.

129. **B. Zhou**, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. "Learning Deep Features for Discriminative Localization." Computer Vision and Pattern Recognition (**CVPR**), 2016. [webpage]

130. Z. Wang, **B. Zhou**, S. Jegelka. "Optimization as Estimation with Gaussian Processes in Bandit Settings." Artificial Intelligence and Statistics (**AISTATS**), 2016, **Oral**.

131. **B. Zhou**, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. "Object Detectors Emerge in Deep Scene CNNs." International Conference on Learning Representations (**ICLR**), 2015 **Oral**.

132. **B. Zhou**, V. Jagadeesh, and R. Piramuthu. "Discovering Visual Concepts from Weakly Labeled Image Collections." Proceedings of 26th IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2015.

133. **B. Zhou**, J. Xiao, A. Lapedriza, A. Torralba, and A. Aude "Learning Deep Features for Scene Recognition using Places Database." Advances in Neural Information Processing Systems 27 (**NIPS**), 2014, **Spotlight** (3.7% acceptance rate). [webpage] [demo]

134. **B. Zhou**, L. Liu, A. Oliva and A. Torralba. "Recognizing City Identity via Attribute Analysis of Geotagged Images." Proceedings of 13th European Conference on Computer Vision (**ECCV**), 2014.

135. **B. Zhou**, X. Tang and X. Wang. "Measuring Crowd Collectiveness." Proceedings of 26th IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2013, **Oral** (3% acceptance rate).

136. **B. Zhou**, X. Tang and X. Wang. "Coherent Filtering: Detecting Coherent Motions from Crowd Clutters." Proceedings of 12th European Conference on Computer Vision (**ECCV**), 2012.

137. **B. Zhou**, X. Wang and X. Tang. "Understanding Collective Crowd Behaviors:Learning a Mixture Model of Dynamic Pedestrian-Agents." Proceedings of 25th IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2012, **Oral** (2.5% acceptance rate).

138. M. Zhu and **B. Zhou**. "Modeling Manifold Ways of Scene Perception". Proceedings of 18th International Conference On Neural Information Processing (**ICONIP**), 2012.

139. **B. Zhou**, X. Wang and X. Tang. "Random Field Topic Model for Semantic Region Analysis in Crowded Scenes from Tracklets." Proceedings of 24th IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2011.

140. **B. Zhou**, X. Hou and L. Zhang. "A phase discrepancy analysis of object motion. " Proceedings of 10th Asian Conference on Computer Vision (**ACCV**), 2010.

141. **B. Zhou** and L.Zhang. "Scene Gist: a holistic generative model of natural image". in Proceedings of 9th Asian Conference on Computer Vision (**ACCV**), 2009.

Patents

1. Rui Xu, LI Xiaoxiao, Bolei Zhou, Chen Change Loy. "Video repair method and apparatus, and storage medium." US Patent 11,544,820, 2023

2. SHEN Yujun, GU Jinjin, Bolei Zhou . "Image processing method and device." US Patent App. 17/536,756, 2022

3. Quanfu Fan, Dan Gutfreund, XIAO Tete, Bolei Zhou. "Context-aware action recognition by dual attention networks." US Patent 11,195,024, 2021

4. Yu Xiong, Qingqiu Huang, GUO Lingfeng, Hang Zhou, Bolei Zhou, Dahua Lin. "Search method and device, and storage medium." US Patent App. 17/362,803, 2021

5. Bolei Zhou, Robinson Piramuthu, Vignesh Jagadeesh. "Discovering visual concepts from weakly labeled image collections." US Patent 9,965,704, 2018

Talks

Invited talk at Synthetic Data for Computer Vision Workshop in CVPR 2025
Invited talk at AVA: Accessbility, Vision, and Autonomy Workshop in CVPR 2025
Invited talk at 4th Workshop on Uncertainty Quantification for Computer Vision in CVPR 2025

Learning to Simulate Real-World Driving Scenarios and Behaviors for AI Safety
Intel Labs seminar talk, Jan.31, 2024

Learning Generalizable and Interpretable AI Agents from Humans
Caltech Computer Vision seminar talk, Feb.21, 2024

Learning Generalizable and Interpretable Embodied AI from Humans
Berkeley DeepDrive seminar talk, Jan.26, 2024

Learning Generalizable and Interpretable Embodied AI from Humans
SCAI Distinguised Lecture Guest Speaker series at Arizona State University, Oct.19, 2023

Learning Generalizable and Interpretable Embodied AI from Humans
UT Austin's Forum for AI, Sept.8, 2023

Learning to simulate real-world traffic scenarios for AI safety
CoPerception Workshop at ICRA'23, London, June 2, 2023

Human-in-the-loop Learning for Controllable AI Agents
Invited Talk at Disney Research, Los Angeles, May 22, 2023

Recreating Diverse Traffic Scenarios for Benchmarking Embodied AI Safety in Simulation
Invited Talk at UCLA Workshop on Synthetic Data, April 14, 2023

Toward Safe and Generalizable Embodied AI
Invited talk at UCSD Pixel Cafe AI Seminar, Feb 24, 2023

Human-in-the-loop Learning for Safe and Interpretable Machine Autonomy
Invited talk at Stanford Vision and Learning seminar, Feb 6, 2023

Recreating Diverse Traffic Scenarios for Benchmarking Embodied AI Safety in Simulation
Keynote Talk at IPAM Workshop on Explainable AI for the Sciences: Towards Novel Insights, Jan 9, 2023

Toward Safe and Generalizable Embodied AI for Decision-Making
Invited talk at NeurIPSâĂŹ22 Workshop on Machine Learning for Autonomous Driving, Dec.3, 2022

Human-AI Interaction via Emergent Learning
Invited talk at SoCal NLP Symposium 2022, University of California, Santa Barbara, November 18, 2022

Human-AI Interaction via Emergent Learning
Keynote talk at AI for Good "Trustworthy AI" Workshop Series by the United Nations ITU (International Telecommunication Union), October 2, 2022

Human-AI Collaboration for Interactive Image Editing and Machine Autonomy
Invited talk at UIUC Vision Seminar (Sept 6, 2022)
Invited talk at Toyota Research Institute (Sept 7, 2022)
Invited talk at Stanford Graphics Seminar (Sept 9, 2022)
Invited talk at SRI International (August 2, 2022)
Invited talk at Horizon Robotics (August 4, 2022)
Invited talk at Nuro Inc (August 10, 2022)
OSU AI seminar talk (April 13, 2022)

The story of my rejection experiences
ICCV 2021 1st Workshop on Share Stories and Lesson Learned, Oct., 2021

Human-centric AI for Computer Vision and Machine Autonomy
ICCV 2021 Tutorial on Trustworthy and Explainable Computer Vision, Oct., 2021

Understanding the human centric side of deep neural networks
Keynote talk at AIMLAI workshop at ECML-PKDD The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases, Sept 13, 2021

Deep Learning for Scene Understanding from Recognition to Recreation
Guest Lecture at Williams College hosted by Prof. Shaoyang Ning, May 12, 2021

Interpretable latent space and inverse problem in deep generative models
ML Seminar at The Center for Machine Learning at Georgia Tech (ML@GT), Jan., 2021

From Interpretable Image Generation to Procedural Generation of Environments
Berkeley Vision Seminar, Nov., 2020

Inverting Latent Space of GANs for Real Image Editing
MIT Vision Seminar, Nov., 2020

Interpreting and Leveraging the Latent Semantics in Deep Generative Models
Invited Talk at ICML Workshop on Extending Explainable AI Beyond Deep Models and Classifiers, Aug., 2020

Exploring and Exploiting Interpretable Semantics in GANs
Lecture at CVPR Tutorial on Interpretable Machine Learning for Computer Vision, July, 2020

Understanding Visual Scenes from Recognition to Recreation
Invited Talk at University of Seattle, Seattle, Dec., 2019

Interpreting Latent Semantics in Deep Generative Models
Invited Talk at Google Research, Seattle, Dec., 2019
Invited Talk at Tencent AI Research Lab, Seattle, Dec., 2019

Understanding Latent Semantics in GANs
Lecture Talk at Interpretable Machine Learning for Computer Vision at ICCV'19, Oct.27, 2019

Interpretable Deep Representations Emerge from Recognizing and Recreating Scenes
Tech Talk at Google AI Cloud, Auguest 14, 2019

Understanding Scenes from Recognition to Recreation
CVPR'19 tutorial on Textures, Objects, Scenes: From handcrafted features, June 17, 2019

Image Manipulation from Adversarial Samples to GANs
CVPR'19 workshop on Adversarial Machine Learning in Real-World Computer Vision Systems, June 16, 2019

Objects Disentangled from Classifying and Synthesizing Scenes
CVPR'19 workshop on Weakly Supervised Learning for Real-World Computer Vision Applications, June 16, 2019

Discovering and Manipulating Facial Attributes in High-fidelity Image Synthesis
GANocracy: Workshop on Theory, Practice and Artistry of Deep Generative Modeling
MIT Quest for Intelligence, May 31 2019

Deep Scene Understanding from Passive to Active Perceptions
Tsinghua-CUHK International Doctoral Forum 2018
Tsinghua Computer Science Seminar, Dec. 2018

Deep Scene Understanding from Passive to Active Perceptions
Tsinghua-CUHK International Doctoral Forum 2018
Tsinghua Computer Science Seminar, Dec. 2018

On the Importance of Single Units in CNNs
CVPR'18 Tutorial on Interpretable Machine Learning for Computer Vision, June 2018

Interpretable Representation Learning for Visual Intelligence
Boston Machine Learning Meetup, August 2018
Open Data Science Conference (ODSC East), Boston, May 2018.
Invited talk at CS Department, Duke University, Feb 2018
Invited talk at CS Department, Brown University, March 2018.

Interpreting Deep Visual Representations via Network Dissection
Vision Sciences Society (VSS) 2018, Florida, May 2018
VASC seminar at Carneigie Mellon University, Nov.2017
New England Computer Vision Workshop, Nov.2017

Interpreting Deep Visual Representations
ICML'17 workshop on Visualization for Deep Learning, Sydney
CVPR'17 Tutorial on Deep Learning for Objects and Scenes, Hawaii

Network Dissection: Quantifying the Interpretability of Deep Visual Representations
CVPR'17, Hawaii

Understand and Leverage the Internal Representations of CNNs
Guest Lecture of the Computer Vision course at Tufts, Mar. 2017
Cornell Tech, Sept 2016
Harvard, April 2016
Boston University, Nov. 2015

Challenges in Deep Sceen Understanding
ECCV'16 ILSVRC and COCO joint workshop, Oct. 2016, Amsterdam.

Object Detectors Emerge in Deep Scene CNNs
ICLR'15, May 2015, San Diego

Learning Deep Features for Scene Recognition
NIPS'14, Dec. 2014, Montreal

Measuring Crowd Collectiveness
CVPR'13, June 2013, Portland.

Understanding Crowd Behaviors
CVPR'12, June 2012, Rhode Island.

Work Experiences

Assistant Professor at the Chinese University of Hong Kong Sept 2018 – Dec 2021
Department of Information Engineering
Department of Computer Science and Engineering (by courtesy)
- Assistant Professor

Postdoctoral Associate at CSAIL MIT May 2018 - Sept 2018
Computer Vision Group
- developed algorithms for visual navigation in simulated environments

Research Assistant at CSAIL MIT Oct 2013 – May 2018
Computer Vision Group
- developed tools to understand and interpret deep neural networks

– worked on visual scene understanding, such as scene recognition and relation reasoning.

Research Intern at Facebook AI Research May 2016 – Sept 2016
Vision Team
– developed reinforcement learning algorithms for visual navigation inside images
– worked with Larry Zitnick and Yuandong Tian

Research Intern at Facebook AI Research June 2015 – Sept 2015
Vision Team
– developed vision and text model for visual question answering
– worked with Rob Fergus, Arthur Szlam, and Yuandong Tian

Intern Research Scientist at eBay Research Labs May 2014 – Aug 2014
Vision Group
– developed novel algorithms to discover concepts from weakly labeled image collections

Summer Intern at Microsoft Research Asia May 2013 – Aug 2013
Knowledge Mining Group
– developed a tool to analyze user migration among online communities, applied to analyze the large-scale user records from Bing Search Engine

Research Assistant at The Chinese University of Hong Kong Aug 2010 – May 2013
Multimedia Laboratory
– worked on activity recognition and crowd behavior analysis in videos

Undergraduate Research Assistant at Shanghai Jiao Tong University Jan 2009 – July 2010
MOE-Microsoft Key Laboratory of Intelligent Computing
– worked on object detection in videos

Professional Services

Area Chairs for ICML'26, ICLR'26, CVPR'25, NeurIPS'25, ECCV'24, WACV'24, CVPR'24, NeurIPS'23, WACV'23, AAAI'23, ECCV'22, CVPR'22, ICCV'21, CVPR'21, BMVC'21, WACV'21

Co-organizer, CVPR 2025 Workshop on Populating Empty Cities, 2025

Co-organizer, CVPR 2025 Workshop on Real-to-Sim: Bridging the Gap between Neural Rendering and Robot Learning, 2025

Guest Editor, special issue for IEEE Robotics & Automation Magazine (RA-M) on Embodied AI: Bridging Robotics and Artificial Intelligence Toward Real-World Applications, 2025

NSF Panelist, 2024

Associate Editor, International Journal of Computer Vision (2024 -)

Associate Editor, Pattern Recognition (2019 - 2022)

External Proposal Reviewer, Singapore Ministry of Education (MOE), 2022

Organizer and Lecturer, 1st Tutorial on Human-Centered AI for Computer Vision at CVPR'22

Senior Program Committee, AAAI'22, AAAI'21

Organizer and Lecturer, 4th Tutorial on Interpretable Machine Learning for Computer Vision at CVPR'21

Organizer and Lecturer, 3rd Tutorial on Interpretable Machine Learning for Computer Vision at CVPR'20

Publicity Chair, ICCV'19

Organizer and Lecturer, 2nd Tutorial on Interpretable Machine Learning for Computer Vision at ICCV'19

Organizer, Tutorial on Architecture Design and Interpretability of Convolutional Neural Networks at ICIP'19

Co-Organizer, Workshop on Explainable AI at CVPR'19

Co-Organizer, Workshop on Network Interpretability for Deep Learning at AAAI'19

Organizer and Lecturer, 1st Tutorial on Interpretable Machine Learning for Computer Vision at CVPR'18
Panelist, NIPS'17 Interpretable Machine Learning Symposium
Program Committee and Speaker, ICML'17 workshop on Visualization for Deep Learning
Co-Organizer, the Joint COCO and Places Recognition Challenge Workshop at ICCV'17
Organizer, the Places Challenge 2017 at ICCV'17
Organizer and Lecturer, Tutorial on Deep Learning for Objects and Scenes at CVPR'17
Organizer, 5th Scene Understanding Workshop(SUNw) at CVPR'17
Organizer, Places365 Challenge 2016 and Scene Parsing Challenge 2016 at ECCV'16
Co-organizer, ILSVRC'16 (ImageNet) challenge workshop at ECCV'16
Organizer, Places2 Scene Recognition Challenge at ICCV'15
Chair, the MIT Vision Seminar 2014-2018
Conference reviewer: CVPR, ICCV, ECCV
Journal reviewer: International Journal on Computer Vision, The Visual Computer, Computer Vision and Image Understanding, IEEE Trans on PAMI, IEEE Trans on NNLS, IEEE Trans on Image Processing, IEEE Trans on SMC., IEEE Trans on CSVT, PLOS ONE, Signal Processing: Image Communication, Pattern Recognition

Teaching Experiences

CS260R: Reinforcement Learning 2023-2026
Lecturer for UCLA graduate course

CS163: Deep Learning for Computer Vision Jan - March, 2022-2025
Lecturer for UCLA undergraduate course

CS269: Human-centered AI for Vision and Autonomy March 2022 - June 2022
Lecturer for UCLA graduate course

Multimedia Coding and Processing Jan 2021 - June 2021
Lecturer for CUHK undergraduate-level course IERG4190 & IEMS5707
- Course page: `http://bzhou.ie.cuhk.edu.hk/teaching/ierg4190iems5707/`

Reinforcement Learning Sept 2020 - Dec 2020
Lecturer for CUHK grad-level course IERG5350
- Course page: `https://cuhkrlcourse.github.io/`

Interacting with Deep Generative Models for Content Creation Oct 28, 2020
Tutorial lecturer for Open Data Science Conference (ODSC) West
- Course page: `https://github.com/zhoubolei/introGM`

AI Education for Secondary Schools at Hong Kong Jan 2020 –
Curriculum developer for the section of computer vision
- Joint effort of CUHK Faculties of Engineering and Education, funded by HK Jockey Club
- Project page: `https://cuhkjc-aiforfuture.hk/index.php/en/home/`

Simulation and Statistical Analysis Sept 2019 - Dec 2019
Lecturer for CUHK undergrad-level course IERG3050
- Course page: `https://course.ie.cuhk.edu.hk/~ierg3050/`

Seminar on Reinforcement Learning Jan 2019 – May 2019, Jan 2020 – May 2020
Lecturer for CUHK grad-level seminar course IERG6130
- Course page: `https://cuhkrlcourse.github.io/2019spring/index.html`

Artificial Intelligence: Implications for Business Strategy Sept 2017 – Dec 2017
Course Facilitator for MIT online course at Sloan Business School and MIT CSAIL
- Developed and reviewed materials of AI technologies for business applications
- Opened online sessions with more than 700 participants
- Instructor: Thomas Malone and Daniela Rus

Advances in Computer Vision Sept 2015 – Dec 2015
Teaching Assistant for MIT 6.869

- Developed course materials and assignments for deep neural networks
- Developed the Mini-Places Classification challenge as student's final project
- Instructors: Aude Oliva and Yusuf Aytar

Multimedia Coding and Processing Sept 2012 – Dec 2012
Teaching Assistant for CUHK IEMS5707

- Developed course materials and assignments, gave tutorials on image processing.
- Instructor: Xiaoou Tang

Current Students and Postdoc Researcher

Postdoc: Wayne Wu
PhD students: Zhenghao Peng, Zhizheng Liu, Sicheng Mo, Qihang Zhang
MS students: Chenda Duan, Weizhen Wang, Michael Simon, Ben Klingher, Chae Yeon Seo, Yuxin Liu, Wenlong Yi

Graduated Students

Yinghao Xu (PhD from CUHK, advisor): 3D-aware image generation 2023
Ceyuan Yang (PhD from CUHK, advisor): Generative Models 2022
Anyi Rao (PhD from CUHK, co-advisor): Intelligent Video Analysis and Creation 2022
Yujun Shen (PhD from CUHK, advisor): Interpreting Deep Generative Models 2020

Sicheng Mo (Master of Science in CS from UCLA, advisor): Training-free Controllable Generation 2024
Ben Klingher (Master of Science in CS from UCLA, advisor): Concept Editing and Spatial Control in the Latent Space of Distilled Diffusion Models 2024
Lalit Bhagat (Master of Science in CS from UCLA, advisor): Steerable Generative Models 2023
Nilay Shah (Master of Science in CS from UCLA, advisor): Unsupervised Discovery of Disentangled Directions in Deep Generative Models 2022
Hao Sun (Mphil from CUHK, advisor): Self-supervised reinforcement learning 2021
Jiankai Sun (Mphil from CUHK, advisor): Robot Learning for Perception and Decision-Making 2021

Ph.D. and Master's Thesis Committee

UCLA Ph.D. Thesis Committee Member: Tianyi Xie (CS, advisor: Demetri Terzopoulos, 2024), Pradyumna Chari (ECE, advisor: Achuta Kadambi, 2024), Rosermary He (CS, advisor: Eleazar Eskin and Daniel Tward, 2024), Sidi Lu (CS, advisor: Nanyun Peng, 2024), Noor Nakhaei (CS, advisor: Demetri Terzopoulos, 2024), Pasha Khosravi (CS, advisor: Guy Van den Broeck, 2024), Albert Zhao (CS, advisor: Stefano Soatto, 2023), Alex Hung (CS, advisor: Kyung Sung, 2023), Noyan Evirgen (ECE, advisor: Anthony Chen, 2023), Runsheng Xu (Civil and Environmental Engineering, advisor: Jiaqi Ma, 2023), Yunhao Ba (ECE, advisor: Achua Kadambi, 2023), Wuyue Lu (CS, advisor: Demetri Terzopoulos, 2022), Xiao Zeng (CS, advisor: Demetri Terzopoulos, 2022), Kaili Wang (EE dept from KU Leuven, advisor: Tinne Tuytelaars, 2021)
UCLA Master's Thesis Committee Member: BLAKE BAGNOL GELLA (ECE, advisor: Achuta Kadambi), Raunak Sinha (CS, advisor: Nanyun Peng, 2024), Zhi Li (CS, advisor: Demetri Terzopoulos, 2024), Alexander 'Sasha' Vilesov (ECE, advisor: Achuta Kadambi, 2023), Jiawei Yang (ECE, advisor: Lin Yang, 2023), Xingyu Liu (ECE, advisor: Anthony Chen, 2022)