

BOLEI ZHOU

Assistant Professor
Department of Information Engineering,
Department of Computer Science and Engineering (by courtesy)
The Chinese University of Hong Kong (CUHK)
Email: bzhou@ie.cuhk.edu.hk | [Homepage](#) | [Google Scholar](#)

Research Interests

Machine perception and decision making, with a focus on enabling interpretable human-AI interaction for machines that act in the real world

Areas: Machine Learning, Computer Vision, Robotics, Visualization

Academic Position

The Chinese University of Hong Kong Sept.2018 - now
Department of Information Engineering
– Assistant Professor

Education

Massachusetts Institute of Technology Aug 2013 – May 2018
Ph.D. in Computer Science
– Thesis: Interpretable Representation Learning for Visual Intelligence
– Advisor: Antonio Torralba
– Thesis Committee: 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: Xiaou 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

MIT Technology Review's Innovators Under 35 in Asia-Pacific Region	2020
The World Artificial Intelligence Conference (WAIC) Yunfun Star Award	2020
Facebook PyRobot: Democratizing Robotics Research Award	2019
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
Postgraduate Studentship in the Chinese University of Hong Kong	2010-2012
Outstanding Undergraduate Thesis of Shanghai Jiao Tong University	2010

Press Coverage

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. Yujun Shen, Ceyuan Yang, Xiaoou Tang, **Bolei Zhou**. "InterFaceGAN: Interpreting the Disentangled Face Representation Learned by GANs" Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), Oct 2020. [[webpage](#)]
2. Ceyuan Yang, Yujun Shen, **Bolei Zhou**. "Semantic Hierarchy Emerges in Deep Generative Representations for Scene Synthesis" under second-round revision of International Journal of Computer Vision (**IJCV**). [[webpage](#)]
3. David Bau, Jun-Yan Zhu, Henrick Strobel, 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.
4. 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](#)]
5. D. Bau, H. Strobel, 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
6. **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.
7. **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](#)]
8. M.Monfort, A. Andonian, **B. Zhou**, S. Bargal, T. Yan, K. Ramakrishnan, L. Brown, Q. Fan, D. Gutfrund, 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](#)]
9. F. Zhang, **B. Zhou**, C. Ratti, Y. Liu. "Discovering place-informative scenes and objects using social media photos." Royal Society Open Science, 2018
10. 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
11. **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](#)]
12. **B. Zhou**, X. Tang, H. Zhang and X. Wang. "Measuring Crowd Collectiveness." IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2014. [[webpage](#)]

13. **B. Zhou**, X. Tang and X. Wang. "Learning Collective Crowd Behaviors with Dynamic Pedestrian-Agents." *International Journal of Computer Vision (IJCV)*, 2014. [\[webpage\]](#)
14. L.Liu, **B. Zhou**, J. Zhao, B.D.Ryan. "C-IMAGE: City Cognitive Mapping through Geo-tagged Photos." *GeoJournal*, Springer, 2016.

Book Chapters

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

17. 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. [\[PDF\]](#)
18. 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.
19. Jiapeng Zhu, Yujun Shen, Deli Zhao, **Bolei Zhou**. "In-Domain GAN Inversion for Real Image Editing ." *European Conference on Computer Vision (ECCV)*, 2020. [\[webpage\]](#)
20. 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.
21. Yujun Shen, Jinjin Gu, Xiaoou Tang, **Bolei Zhou**. "Interpreting Latent Space of GANs for Semantic Face Editing." *Computer Vision and Pattern Recognition (CVPR)*, 2020. [\[webpage\]](#)
22. Jinjin Gu, Yujun Shen, **Bolei Zhou**. "Image Processing Using Multi-Code GAN Prior." *Computer Vision and Pattern Recognition (CVPR)*, 2020. [\[webpage\]](#)
23. Ceyuan Yang*, Yinghao Xu*, Bo Dai, **Bolei Zhou**. "Temporal Pyramid Network for Action Recognition." *Computer Vision and Pattern Recognition (CVPR)*, 2020. [\[webpage\]](#)
24. Liangji Fang*, Qinrong Jiang*, Jianping Shi, **Bolei Zhou**. "Trajectory Proposal Network for Multi-modal Motion Prediction." *Computer Vision and Pattern Recognition (CVPR)*, 2020. [\[webpage\]](#)
25. 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. [\[webpage\]](#)
26. 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.
27. 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.
28. 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).
29. 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

30. 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**.
31. D. Bau, H. Strobel, W. Peebles, J. Wulff, J. Zhu, **B. Zhou**, A. Torralba. "Seeing What a GAN Cannot Generate." International Conference on Computer Vision (**ICCV**), 2019, **Oral**.
32. D. Bau, J. Zhu, H. Strobel, **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.
33. R. Xu, X. Li, **B. Zhou**, C. C. Loy. "Deep Flow-Guided Video Inpainting." Computer Vision and Pattern Recognition (**CVPR**), 2019
34. 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
35. **B. Zhou**, A. Andonian, A. Oliva, A. Torralba. "Temporal Relational Reasoning in Videos." European Conference on Computer Vision (**ECCV**), 2018.
36. **B. Zhou***, Y.Sun*, D.Bau*, A. Torralba. "Interpretable Basis Decomposition for Visual Explanation." European Conference on Computer Vision (**ECCV**), 2018.
37. T. Xiao*, Y. Liu*, **B. Zhou***, Y. Jiang, J. Sun. "Unified Perceptual Parsing for Scene Understanding." European Conference on Computer Vision (**ECCV**), 2018.
38. 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.
39. 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.
40. 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**.
41. 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.
42. 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.
43. J. Wu, D. Peck, S. Hsieh, V. Dialani, C. Lehman, **B. Zhou**, V. Syrgkanis, L. Mackey, G. Patterson. "Expert Identification of Visual Primitives used by CNNs during Mammogram Classification." SPIE Medical Imaging, 2018.
44. H. Zhao, X. Puig, **B. Zhou**, S. Fidler, and A. Torralba. "Open Vocabulary Scene Parsing." International Conference on Computer Vision (**ICCV**), 2017.
45. 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.
46. **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](#)]
47. 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](#)]
48. 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.

49. 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**.
50. **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](#)]
51. Z. Wang, **B. Zhou**, S. Jegelka. "Optimization as Estimation with Gaussian Processes in Bandit Settings." Artificial Intelligence and Statistics (**AISTATS**), 2016, **Oral**.
52. **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**.
53. **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.
54. **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](#)]
55. **B. Zhou**, L. Liu, A. Oliva and A. Torralba. "Recognizing City Identity via Attribute Analysis of Geo-tagged Images." Proceedings of 13th European Conference on Computer Vision (**ECCV**), 2014.
56. **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).
57. **B. Zhou**, X. Tang and X. Wang. "Coherent Filtering: Detecting Coherent Motions from Crowd Clusters." Proceedings of 12th European Conference on Computer Vision (**ECCV**), 2012.
58. **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).
59. M. Zhu and **B. Zhou**. "Modeling Manifold Ways of Scene Perception". Proceedings of 18th International Conference On Neural Information Processing (**ICONIP**), 2012.
60. **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.
61. **B. Zhou**, X. Hou and L. Zhang. "A phase discrepancy analysis of object motion. " Proceedings of 10th Asian Conference on Computer Vision (**ACCV**), 2010.
62. **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.

Talks

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, August 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 Carnegie 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 Scene 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 –

Department of Information Engineering

Department of Computer Science and Engineering (by courtesy)

Postdoctoral Associate at CSAIL MIT

May 2018 - Sept 2018

Computer Vision Group

- developing algorithms for visual navigation in simulated environments

Research Assistant at CSAIL MIT

Oct 2013 – May 2018

Computer Vision Group

- developing tools to understand and interpret deep neural networks
- working 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

Senior Program Committee, AAAI'21
Area Chair, CVPR'21
Organizer and Lecturer, 3rd Interpretable Machine Learning for Computer Vision at CVPR'20
Publicity Chair, ICCV'19
Organizer and Lecturer, 2nd 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, 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
Conference reviewer: CVPR'18, BMVC'17, ICCV'17, CVPR'17, ECCV'16, CVPR'16, ICCV'15, CVPR'15, ECCV'14, ACCV'14, CVPR'13
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

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

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

Ceyuan Yang, Yujun Shen, Jiankai Sun, Yinghao Xu, Jinjin Gu, Anyi Rao (co-advised with D.Lin), Hao Sun (co-advised with D.Lin)

Previous Students

Tete Xiao (Undergrad, now a PhD student at Berkeley with Trevor Darrell): Dual attention model for video understanding	2018
Bowen Pan (Undergrad, now a PhD student at MIT with Aude Oliva): Active scene perception	2018
Jimmy Wu (Master, now a PhD student at Princeton with Thomas Funkhouser): Semantic segmentation and pose estimation for robotics	2017
Kayode Ezike (Master): Optic flow for activity recognition	2017
Harini Kannan (Master): Towards more accurate eye-tracking	2016
Alex Andonian (Undergrad): Temporal relation analysis in videos	2017
Jeff Hu (Undergrad): Scaling pixel-wise annotations	2017
Jing Li (Undergrad): Self-supervision from Eye-fixations	2016