

BOLEI ZHOU

Assistant Professor
Department of Information Engineering
The Chinese University of Hong Kong
Email: bzhou@ie.cuhk.edu.hk
[Homepage](#)
[Google Scholar](#)

Research Interests

Computer Vision, Deep Learning, Artificial Intelligence

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., Computer Science
– Advisor: Antonio Torralba
– Thesis: Interpretable Representation Learning for Visual Intelligence
– Thesis Committee: Antonio Torralba, Aude Oliva, William T. Freeman

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

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

Awards and Honors

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

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. **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.
2. **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 (arXiv:1711.05611) [[webpage](#)]
3. M.Monfort, **B. Zhou**, S. Bargal, A. Andonian, T. Yan, K. Ramakrishnan, L. Brown, Q. Fan, D. Gutfrund, C. Vondrick, A. Oliva. "Moments in Time Dataset: One Million Videos for Event Understanding." under major revision of the IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2018 (arXiv:1801.03150) [[webpage](#)]
4. **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](#)]
5. **B. Zhou**, X. Tang, H. Zhang and X. Wang. "Measuring Crowd Collectiveness." IEEE transaction on Pattern Analysis and Machine Intelligence (**TPAMI**), 2014. [[webpage](#)]
6. **B. Zhou**, X. Tang and X. Wang. "Learning Collective Crowd Behaviors with Dynamic Pedestrian-Agents." International Journal of Computer Vision (**IJCV**), 2014. [[webpage](#)]
7. L.Liu, **B. Zhou**, J. Zhao, B.D.Ryan. "C-IMAGE: City Cognitive Mapping through Geo-tagged Photos." GeoJournal, Springer, 2016.

Conference Publications

8. **B. Zhou**, A. Andonian, A. Oliva, A. Torralba. "Temporal Relational Reasoning in Videos." European Conference on Computer Vision (**ECCV**), 2018.
9. **B. Zhou***, Y.Sun*, D.Bau*, A. Torralba. "Interpretable Basis Decomposition for Visual Explanation." European Conference on Computer Vision (**ECCV**), 2018.
10. T. Xiao*, Y. Liu*, **B. Zhou***, Y. Jiang, J. Sun. "Unified Perceptual Parsing for Scene Understanding." European Conference on Computer Vision (**ECCV**), 2018.
11. 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.
12. 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.
13. 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'18**), 2018, **Oral**.
14. 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'18**), 2018.
15. 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'18**), 2018.
16. 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.

17. H. Zhao, X. Puig, **B. Zhou**, S. Fidler, and A. Torralba. "Open Vocabulary Scene Parsing." International Conference on Computer Vision (ICCV'17), 2017.
18. 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'17), 2017.
19. **B. Zhou**, H. Zhao, X. Puig, S. Fidler, A. Barriuso and A. Torralba. "Scene Parsing through ADE20K Dataset." Computer Vision and Pattern Recognition (CVPR'17), 2017. [[webpage](#)][[ADE dataset](#)]
20. 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'17) **Oral** (3% acceptance rate), 2017. * indicates equal contribution. [[webpage](#)]
21. 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'17), 2017.
22. J. Wong, V. Kee, T. Le, S. Wagner, G. Mariottini, A. Schneider, L. Hamilton, R. Chiaplakatty, 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'17) **Oral**.
23. **B. Zhou**, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. "Learning Deep Features for Discriminative Localization." Computer Vision and Pattern Recognition (CVPR'16), 2017. [[webpage](#)]
24. Z. Wang, **B. Zhou**, S. Jegelka. "Optimization as Estimation with Gaussian Processes in Bandit Settings." Artificial Intelligence and Statistics (AISTATS'16) **Oral** (6% acceptance rate).
25. **B. Zhou**, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. "Object Detectors Emerge in Deep Scene CNNs." International Conference on Learning Representations (ICLR'15) **Oral**.
26. **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'15).
27. **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'14), **Spotlight** (3.7% acceptance rate). [[webpage](#)] [[demo](#)]
28. **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'14).
29. **B. Zhou**, X. Tang and X. Wang. "Measuring Crowd Collectiveness." Proceedings of 26th IEEE Conference on Computer Vision and Pattern Recognition (CVPR'13), **Oral** (3% acceptance rate).
30. **B. Zhou**, X. Tang and X. Wang. "Coherent Filtering: Detecting Coherent Motions from Crowd Clusters." Proceedings of 12th European Conference on Computer Vision (ECCV'12).
31. **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'12), **Oral** (2.5% acceptance rate).
32. M. Zhu and **B. Zhou**. "Modeling Manifold Ways of Scene Perception". Proceedings of 18th International Conference On Neural Information Processing (ICONIP'12).
33. **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'11).
34. **B. Zhou**, X. Hou and L. Zhang. "A phase discrepancy analysis of object motion. " Proceedings of 10th Asian Conference on Computer Vision (ACCV'10).
35. **B. Zhou** and L. Zhang. "Scene Gist: a holistic generative model of natural image". in Proceedings of 9th Asian Conference on Computer Vision (ACCV'09).

36. **B. Zhou** and L. Zhang. "A hierarchical model for visual perception." Proceedings of 2nd International Conference on Cognitive Neurodynamics (ICCN'09).

ArXiv Preprints

37. **B. Zhou**, A. Andonian, A. Torralba. "Revisiting the Importance of Single Units in CNNs." arXiv:1711.08496, 2017.
38. **B. Zhou**, S. Sukhbaatar, A. Szlam, R. Fergus, Yuandong Tian. "Simple Baseline for Visual Question Answering." arXiv:1512.02167, 2016.
39. D. Wei, **B. Zhou**, A. Torralba, W. Freeman. "Understanding Intra-Class Knowledge inside CNN." arXiv:1507.02379, 2015.

Workshop Papers

40. **B. Zhou***, D. Bau*, A. Oliva, and A. Torralba. "Comparing the Interpretability of the Deep Visual Representations via Network Dissection." NIPS'17 Workshop on Interpreting, Explaining, and Visualizing Deep Learning.
41. **B. Zhou***, D. Bau*, A. Oliva, and A. Torralba. "Quantifying Interpretations of Deep Visual Representations." ICML'17 Workshop on Visualization for Deep Learning.
42. D. Bau*, **B. Zhou***, A. Khosla, A. Oliva. and A. Torralba. "Network Dissection: Quantifying Interpretability of Deep Visual Representations." CVPR'17 Scene Understanding Workshop (SUNw), 2017.
43. **B. Zhou**, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. "Object Detectors Emerge from Training CNNs for Scene Recognition." CVPR'15 Scene Understanding Workshop (SUNw), 2015.
44. **B. Zhou**, X. Tang, X. Wang. "Measurability of Crowd Collectiveness in Dynamic Scenes." CVPR'13 Scene Understanding Workshop (SUNw), 2013.

Talks

Deep Scene Understanding from Passive to Active Perceptions

Tsinghua-CUHK International Doctoral Forum 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 - present

Information Engineering Department

- focus on visual scene understanding and interpretable AI systems
- teach IERG6130 Reinforcement Learning (Spring 2019)

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

Organizer and Lecturer, 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

MIT Sloan Business School and MIT CSAIL

Sept 2017 – Dec 2017

Course Facilitator for MIT online Course: Artificial Intelligence: Implications for Business Strategy

- Developed and reviewed materials of AI technologies for business applications
- Opened online sessions with more than 700 participants
- Instructor: Thomas Malone and Daniela Rus

Massachusetts Institute of Technology

Sept 2015 – Dec 2015

Teaching Assistant for MIT 6.869 Advances in Computer Vision

- 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

Chinese University of Hong Kong

Sept 2012 – Dec 2012

Teaching Assistant for CUHK IEMS5707 Multimedia Coding and Processing

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

Students Advised

Jimmy Wu (Master): 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