

Jun-Yan Zhu

Curriculum Vitae

MIT CSAIL, 32 Vassar Street

Cambridge, MA 02139

✉ junyanz@mit.edu

Homepage: people.csail.mit.edu/junyanz/

Research Interests

Computer Vision, Computer Graphics, and Machine Learning

Employment

2018–now **Massachusetts Institute of Technology.**

Postdoctoral researcher at MIT CSAIL

Advisors: William T. Freeman, Joshua Tenenbaum, and Antonio Torralba

Education

2013–2017 **University of California, Berkeley.**

Ph. D. in Computer Science, EECS

Thesis: *Learning to Synthesize and Manipulate Natural Images*

Advisor: Alexei A. Efros

2012–2013 **Carnegie Mellon University.**

Ph. D. student, Computer Science Department

Advisor: Alexei A. Efros

2008–2012 **Tsinghua University.**

B. E. in Computer Science and Technology

Ranked 2nd out of 140, Class of 2012

Awards

2019 ACM SIGGRAPH Real-time Live Best in Show Award

2019 ACM SIGGRAPH Real-time Live Audience Choice Award

2018 ACM SIGGRAPH Outstanding Doctoral Dissertation Award

2018 UC Berkeley EECS David J. Sakrison Memorial Prize for Outstanding Doctoral Research

2018 NVIDIA Pioneer Research Award

2015 Facebook Graduate Fellowship

2012 Outstanding Undergraduate Thesis at Tsinghua University

2012 Excellent Undergraduate Student at Tsinghua University

2009, 2010 National Scholarship, by the Ministry of Education of China

2010–2012 Singapore Technologies Engineering China Scholarship

Publications

- [1] Subramanian Sundaram, Petr Kellnhofer, Yunzhu Li, **Jun-Yan Zhu**, Antonio Torralba, and Wojciech Matusik. Learning the signatures of the human grasp using a scalable tactile glove. *Nature*, 569(7758), 2019.
- [2] Taesung Park, Ting-Chun Wang, Chris Hebert, **Jun-Yan Zhu**, Gavriil Klimov, and Ming-Yu Liu. GauGAN: Semantic image synthesis with spatially adaptive normalization. In *ACM SIGGRAPH 2019 Real-Time Live*, 2019.
- [3] David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, **Jun-Yan Zhu**, and Antonio Torralba. Semantic photo manipulation with a generative image prior. *ACM Transactions on Graphics (SIGGRAPH)*, 38(4), 2019.
- [4] Yunzhu Li, **Jun-Yan Zhu**, Russ Tedrake, and Antonio Torralba. Connecting touch and vision via cross-modal prediction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

- [5] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and **Jun-Yan Zhu**. Semantic image synthesis with spatially-adaptive normalization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [6] David Bau, **Jun-Yan Zhu**, Hendrik Strobelt, Zhou Bolei, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. GAN dissection: Visualizing and understanding generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [7] Yunzhu Li, Jiajun Wu, **Jun-Yan Zhu**, Joshua B Tenenbaum, Antonio Torralba, and Russ Tedrake. Propagation networks for model-based control under partial observation. In *International Conference on Robotics and Automation (ICRA)*, 2019.
- [8] **Jun-Yan Zhu**, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Joshua B. Tenenbaum, and William T. Freeman. Visual object networks: Image generation with disentangled 3D representations. In *Neural Information Processing System (NeurIPS)*, 2018.
- [9] Shunyu Yao, Tzu Ming Hsu, **Jun-Yan Zhu**, Jiajun Wu, Antonio Torralba, William T. Freeman, and Joshua B. Tenenbaum. 3D-aware scene manipulation via inverse graphics. In *Neural Information Processing System (NeurIPS)*, 2018.
- [10] Ting-Chun Wang, Ming-Yu Liu, **Jun-Yan Zhu**, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. In *Neural Information Processing System (NeurIPS)*, 2018.
- [11] Judy Hoffman, Eric Tzeng, Taesung Park, **Jun-Yan Zhu**, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *International Conference on Machine Learning (ICML)*, 2018.
- [12] Ting-Chun Wang, Ming-Yu Liu, **Jun-Yan Zhu**, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional GANs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [13] Chaowei Xiao*, **Jun-Yan Zhu***, Bo Li, Warren He, Mingyan Liu, and Dawn Song. Spatially transformed adversarial examples. In *International Conference on Learning Representations (ICLR)*, 2018.
- [14] Chaowei Xiao, Bo Li, **Jun-Yan Zhu**, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [15] **Jun-Yan Zhu**, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Neural Information Processing System (NeurIPS)*, 2017.
- [16] **Jun-Yan Zhu***, Taesung Park*, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *International Conference on Computer Vision (ICCV)*, 2017.
- [17] Phillip Isola, **Jun-Yan Zhu**, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [18] Richard Zhang*, **Jun-Yan Zhu***, Phillip Isola, Xinyang Geng, Angela S Lin, Tianhe Yu, and Alexei A Efros. Real-time user-guided image colorization with learned deep priors. *ACM Transactions on Graphics (SIGGRAPH)*, 2017.
- [19] Ting-Chun Wang, **Jun-Yan Zhu**, Nima Khademi Kalantari, Alexei A. Efros, and Ravi Ramamoorthi. Light field video capture using a learning-based hybrid imaging system. *ACM Transactions on Graphics (SIGGRAPH)*, 2017.
- [20] **Jun-Yan Zhu**, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative visual manipulation on the natural image manifold. In *European Conference on Computer Vision (ECCV)*, 2016.

- [21] Ting-Chun Wang, **Jun-Yan Zhu**, Ebi Hiroaki, Manmohan Chandraker, Alexei A. Efros, and Ravi Ramamoorthi. A 4D light-field dataset and CNN architectures for material recognition. In *European Conference on Computer Vision (ECCV)*, 2016.
- [22] **Jun-Yan Zhu**, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Learning a discriminative model for the perception of realism in composite images. In *International Conference on Computer Vision (ICCV)*, 2015.
- [23] **Jun-Yan Zhu**, Jiajun Wu, Yan Xu, Eric Chang, and Zhuowen Tu. Unsupervised object class discovery via saliency-guided multiple class learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2015.
- [24] **Jun-Yan Zhu**, Aseem Agarwala, Alexei A Efros, Eli Shechtman, and Jue Wang. Mirror mirror: Crowdsourcing better portraits. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 2014.
- [25] **Jun-Yan Zhu**, Yong Jae Lee, and Alexei A Efros. AverageExplorer: Interactive exploration and alignment of visual data collections. *ACM Transactions on Graphics (SIGGRAPH)*, 2014.
- [26] Jiajun Wu, Yibiao Zhao, **Jun-Yan Zhu**, Siwei Luo, and Zhuowen Tu. MILCut: A sweeping line multiple instance learning paradigm for interactive image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [27] Jiajun Wu, **Jun-Yan Zhu**, and Zhuowen Tu. Reverse image segmentation: A high-level solution to a low-level task. In *British Machine Vision Conference (BMVC)*, 2014.
- [28] Yan Xu, **Jun-Yan Zhu**, Eric I. Chang, Maode Lai, and Zhuowen Tu. Weakly supervised histopathology cancer image segmentation and classification. *Medical Image Analysis*, 2014.
- [29] Tao Chen, **Jun-Yan Zhu**, Ariel Shamir, and Shi-Min Hu. Motion-aware gradient domain video composition. *IEEE Transactions on Image Processing (TIP)*, 2013.
- [30] **Jun-Yan Zhu**, Jiajun Wu, Yichen Wei, Eric Chang, and Zhuowen Tu. Unsupervised object class discovery via saliency-guided multiple class learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [31] Yan Xu*, **Jun-Yan Zhu***, Eric Chang, and Zhuowen Tu. Multiple clustered instance learning for histopathology cancer image classification, segmentation, and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

Preprints

- [32] Tongzhou Wang, **Jun-Yan Zhu**, Antonio Torralba, and Alexei A. Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.

Teaching

- 2018 **Co-instructor**, Deep Learning (800 enrolled students), Udacity.
with Sebastian Thrun, Ian Goodfellow, Andrew Trask, and Udacity Deep Learning Team.
- 2018 **Guest Lecturer**, 6.819/6.869 Advances in Computer Vision, MIT.
- 2016, 2017 **Teaching Assistant**, Image Manipulation and Computational Photography, UC Berkeley.
Taught four 90-minute lectures on point processing, image warping, and image morphing.

Academic Service

Area Chair/Editor

- 2018 Technical Papers Committee member, SIGGRAPH Asia 2018
- 2018 Guest editor, International Journal of Computer Vision (IJCV)

Workshop/Tutorial/Course Organization

- 2018 Organizer, CVPR 2018 Tutorial on Generative Adversarial Networks
- 2018 Organizer, MIT Quest Symposium on Robust, Interpretable Deep Learning Systems

2017 Instructor, ICCV 2017 Tutorial on Generative Adversarial Networks
2017 Organizer, ICML 2017 Workshop on Visualization for Deep Learning
2014 Organizer, SIGGRAPH Asia 2014 invited Course on Data-Driven Visual Computing
Conference and Journal Reviewer
CVPR (Outstanding Reviewer Award 2017, 2019), ICCV, ECCV, IJCV, TPAMI
SIGGRAPH, SIGGRAPH Asia, Eurographics, ICML, CHI

Invited Talks

2019 **Learning to Synthesize Images**

Carnegie Mellon University, Pittsburgh, PA
Massachusetts Institute of Technology, Cambridge, MA
Stanford University, Stanford, CA
The University of Maryland, College Park, MD
The University of Texas at Austin, Austin, TX
University of California San Diego, La Jolla, CA
University of Washington, Seattle, WA

2018 **Learning to Generate Images**

SIGGRAPH Dissertation Award Talk, Vancouver, Canada
UMass Machine Learning and Friends Lunch, Amherst, MA

2017-2018 **Unpaired Image-to-Image Translation**

CVPR 2018 Tutorial on GANs, Salt Lake City, UT
ICML 2017 Workshop on Implicit Models, Sydney, Australia

2017 **Learning to Synthesize and Manipulate Natural Photos**

MIT CSAIL Vision Seminar, Cambridge, MA
HKUST CSE Departmental Seminar, Hong Kong
ICCV 2017 Tutorial on GANs, Venice, Italy
O'Reilly Artificial Intelligence Conference, New York City, NY
DEVIEW Developer Conference, Seoul, Korea
Open Data Science Conference, San Francisco, CA
Y Combinator Research Conference, San Francisco, CA

2017 **On Image-to-Image Translation**

Stanford EECS Seminar, Stanford, CA
MIT CSAIL Graphics Lunch, Cambridge, MA
Facebook Fellows Research Workshop, Menlo Park, CA
Chinese University of Hong Kong CSE Seminar, Hong Kong
Seoul National University CSE Seminar, Seoul, Korea

2017 **Interactive Deep Colorization**

SIGGRAPH 2017, Los Angeles, CA
NVIDIA Innovation Theater, Los Angeles, CA
Global AI Hackathon, Seattle, WA

2016 **Visual Manipulation and Synthesis on the Natural Image Manifold**

Facebook Fellows Research Workshop, Menlo Park, CA
UC Berkeley BAIR Seminar, Berkeley, CA
Tsinghua University, Beijing, China
Microsoft Research Asia, Beijing, China
ICML 2016 Workshop on Visualization for Deep Learning, New York City, NY

2014 **Mirror Mirror: Crowdsourcing Better Portraits**

SIGGRAPH Asia 2014, Shenzhen, China

2014 **What Makes Big Visual Data Hard?**

SIGGRAPH Asia 2014 Invited Course, Shenzhen, China

2014 **AverageExplorer: Interactive Exploration and Alignment of Visual Data Collections**

SIGGRAPH 2014, Vancouver, Canada

2014 **Discovering Objects and Harvesting Visual Concepts via Weakly Supervised Learning**

UC Berkeley Visual Computing Lab Lunch, Berkeley, CA

Selected Press

2019 CNN: MIT teaches robots to ‘feel’ objects just by looking at them

2019 The Economist: Improving robots’ grasp requires a new way to measure it in humans

2019 BBC Radio: Science unwrapped - interactive science, medicine and technology (06/02/2019)

2019 Nature News: Bridging the gap between artificial vision and touch

2017 Forbes: What’s Next for Deep Learning?

2017 Distill: Using Artificial Intelligence to Augment Human Intelligence.

2016 Quartz: This digital brush paints with the memories of 275,000 landscapes.

2014 The New Yorker: One of Many, One: The Science of Composite Photography.

Past Employment

2013–2017 **Berkeley AI Research (BAIR) Lab**

Research assistant with Alexei A. Efros

2016 **Google Research**

Intern with Ce Liu, Michael Rubinstein, and William T. Freeman

2013–2017 **Adobe Research**

Intern with Eli Shechtman (’13, ’15, ’17), Oliver Wang (’17), Aseem Agarwala and Jue Wang (’13)

2012–2013 **CMU Computer Vision Group and Graphics Lab**

Research assistant with Alexei A. Efros

2011–2012 **Microsoft Research Asia**

Intern with Zhuowen Tu and Eric Chang

2010–2012 **Graphics and Geometric Computing Group, Tsinghua University**

Research assistant with Shi-Min Hu

Patents

2016 US9317781B2: Multiple cluster instance learning for image classification.

2015 US9224071B2: Unsupervised object class discovery via bottom up multiple class learning.