

Jun-Yan Zhu

Curriculum Vitae

5000 Forbes Ave
Pittsburgh, PA 15213

✉ junyanz@cs.cmu.edu

Homepage: <https://www.cs.cmu.edu/~junyanz/>

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

Employment

- 2020–present **Carnegie Mellon University.**
Assistant Professor at the School of Computer Science
- 2019–2020 **Adobe Research**
Research Scientist at Creative Intelligence Lab
- 2018–2019 **MIT CSAIL**
Postdoc with William T. Freeman, Joshua Tenenbaum, and Antonio Torralba
- 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)
- 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

Awards

- 2019 Sony Faculty Research Award
- 2019 The 100 Greatest Innovations of 2019 by Popular Science
- 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

- [1] Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. Sketch your own GAN. In *International Conference on Computer Vision (ICCV)*, 2021.
- [2] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *International Conference on Computer Vision (ICCV)*, 2021.
- [3] Ji Lin, Richard Zhang, Frieder Ganz, Song Han, and Jun-Yan Zhu. Anycost gans for interactive image synthesis and editing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [4] Lucy Cai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Anycost gans for interactive image synthesis and editing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [5] David Bau, Jun-Yan Zhu, Hendrik 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 (PNAS)*, 2020.
- [6] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. In *Neural Information Processing System (NeurIPS)*, 2020.
- [7] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *Neural Information Processing System (NeurIPS)*, 2020.
- [8] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for conditional image synthesis. In *European Conference on Computer Vision (ECCV)*, 2020.
- [9] David Bau, Steven Liu, Tongzhou Wang, Jun-Yan Zhu, and Antonio Torralba. Rewriting a deep generative model. In *European Conference on Computer Vision (ECCV)*, 2020.
- [10] William Peebles, John Peebles, Jun-Yan Zhu, Alexei A. Efros, and Antonio Torralba. The hessian penalty: A weak prior for unsupervised disentanglement. In *European Conference on Computer Vision (ECCV)*, 2020.
- [11] Minyoung Huh, Richard Zhang, Jun-Yan Zhu, Sylvain Paris, and Aaron Hertzmann. Transforming and projecting images to class-conditional generative networks. In *European Conference on Computer Vision (ECCV)*, 2020.
- [12] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, K. Sunkavalli, R. Martin-Brualla, T. Simon, J. Saragih, M. Nießner, R. Pandey, S. Fanello, G. Wetzstein, **J.-Y. Zhu**, C. Theobalt, M. Agrawala, E. Shechtman, D. B Goldman, and M. Zollhöfer. State of the Art on Neural Rendering. *Computer Graphics Forum (EuroGraphics STAR)*, 2020.
- [13] Muyang Li, Ji Lin, Yaoyao Ding, Zhijian Liu, Jun-Yan Zhu, and Song Han. Gan Compression: Efficient architectures for interactive conditional gans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [14] 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.
- [15] David Bau, Jun-Yan Zhu, Jonas Wulff, William Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a gan cannot generate. In *International Conference on Computer Vision (ICCV)*, 2019.

- [16] 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. **Best in Show Award and Audience Choice Award**.
- [17] David Bau, Jun-Yan Zhu, Jonas Wulff, William Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a GAN cannot generate. In *International Conference on Computer Vision (ICCV)*, 2019.
- [18] 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.
- [19] 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.
- [20] 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. **Best Paper Finalist**.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] 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.

- [32] 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.
- [33] 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.
- [34] 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.
- [35] 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.
- [36] 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.
- [37] 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.
- [38] 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.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] 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.
- [43] 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.
- [44] Tao Chen, Jun-Yan Zhu, Ariel Shamir, and Shi-Min Hu. Motion-aware gradient domain video composition. *IEEE Transactions on Image Processing (TIP)*, 2013.
- [45] 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.
- [46] 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

- [45] Gaurav Parmar, Richard Zhang, Jun-Yan Zhu. On Buggy Resizing Libraries and Surprising Subtleties in FID Calculation. *arXiv preprint arXiv:2104.11222*, 2021.
- [46] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, Bryan Russell. Editing Conditional Radiance Fields. *arXiv preprint arXiv:2105.06466*, 2021.

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

Academic Service

Area Chair/Editor

- 2021 Area chair, NeurIPS 2021
- 2021 Area chair, CVPR 2021
- 2020 Area chair, CVPR 2020
- 2019 Technical Briefs and Posters Committee member, SIGGRAPH Asia 2019
- 2019 Guest editor, International Journal of Computer Vision (IJCV)
- 2018 Technical Papers Committee member, SIGGRAPH Asia 2018

Workshop/Tutorial/Course

- 2021 Organizer, SIGGRAPH 2021 Course on Advances in Neural Rendering
- 2021 Organizer, SIGGRAPH 2021 Workshop on Measurable Creative AI
- 2021 Organizer, CVPR 2021 Workshop on Computational Measurements of Machine Creativity
- 2020 Organizer, CVPR 2020 Tutorial on Neural Rendering
- 2020 Organizer, Eurographics 2020 STAR on Neural Rendering
- 2019 Organizer, ICCV 2019 Workshop on Image and Video Synthesis
- 2019 Organizer, CVPR 2019 Tutorial on Map Synchronization
- 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

Journal and Conference Reviewer

Science, IJCV, TPAMI, ACM Transactions on Graphics
CVPR (Outstanding Reviewer Award 2017, 2019), ICCV (Outstanding Reviewer Award 2021), ECCV, SIGGRAPH, SIGGRAPH Asia, Eurographics, ICML, NeurIPS, CHI

Invited Talks

- 2021 **AI for Content Creation**
AI4All @ Carnegie Mellon program
- 2021 **Sketch Your Own Models**
CVPR 2021 Workshop on Sketch-Oriented Deep Learning
- 2021 **GANs for Everyone**
CMU RI Seminar, Pittsburgh, PA
- 2020-2021 **Understanding and Rewriting GANs**
SIGGRAPH 2021 Workshop on Measurable Creative AI
ICML 2021 Workshop on Socially Responsible Machine Learning
Stanford CS348I “Computer Graphics in the Era of AI” (Guest Lecture), Stanford, CA
Tsinghua University “Introduction to Artificial Intelligence” (Guest Lecture), Beijing, China
ISEA Workshop on Measuring Computational Creativity, Montreal, Canada
- 2020 **Efficient GANs**
MIT 6.S192 “Deep Learning for Art, Aesthetics, and Creativity”, Cambridge, MA
- 2020 **3D-Aware Image Synthesis and Editing**
IJCAI-PRICAI 2020 3D-FUTURE Workshop, Yokohama, Japan
- 2019-2020 **Visualizing and Understanding GANs**

- CVPR 2020 Workshop on Human-centric Image/Video Synthesis
 CVPR 2020 AC Workshop, San Diego, CA
 CVPR 2019 Tutorial on Deep Learning for Content Creation, Long Beach, CA
 CVPR 2019 Workshop on New Trends in Image Restoration and Enhancement, Long Beach, CA
- 2020 Semantic Photo Synthesis**
 CVPR 2020 Tutorial on Neural Rendering
 Eurographics 2020 STAR on Neural Rendering
- 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
 Massachusetts Institute of Technology, Cambridge, 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

Teaching

- 2021 **Instructor**, 16-824: Visual Learning and Recognition (Fall 2021)
- 2021 **Instructor**, 16-726: Learning-based Image Synthesis (Spring 2021)
- 2018 **Co-instructor**, Deep Learning (800 enrolled students), Udacity.
with Sebastian Thrun, Ian Goodfellow, Andrew Trask, and Udacity Deep Learning Team.

Students

Ph.D. students

Kangle Deng (co-advised with Deva Ramanan), Sheng-Yu Wang

MS students

George Cazenavette, Nupur Kumari, Muyang Li, Gaurav Parmar

Patents

- 2020 US20200242771: Semantic Image Synthesis for Generating Substantially Photorealistic Images Using Neural Networks.
- 2016 US9317781B2: Multiple cluster instance learning for image classification.
- 2015 US9224071B2: Unsupervised object class discovery via bottom up multiple class learning.

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.