

Employment

- 2026–present **Carnegie Mellon University**
Michael B. Donohue Associate Professor of Computer Science and Robotics.
Director of [Generative Intelligence Lab](#)
- 2020–2026 **Carnegie Mellon University**
Michael B. Donohue Assistant Professor of Computer Science and Robotics.
- 2019–2020 **Adobe Research**
Research Scientist at Creative Intelligence Lab
- 2018–2019 **MIT CSAIL**
Postdoc with 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
- 2008–2012 **Tsinghua University**
B. E. in Computer Science and Technology

Selected Publications

See the full list on Google Scholar [Page](#)

- [1] Gaurav Parmar, Or Patashnik, Daniil Ostashev, Kuan-Chieh Wang, Kfir Aberman, Srinivasa Narasimhan, and Jun-Yan Zhu. Scaling group inference for diverse and high-quality generation. In *International Conference on Learning Representations (ICLR)*, 2026.
- [2] Nupur Kumari, Sheng-Yu Wang, Nanxuan Zhao, Yotam Nitzan, Yuheng Li, Krishna Kumar Singh, Richard Zhang, Eli Shechtman, Jun-Yan Zhu, and Xun Huang. Learning an image editing model without image editing pairs. In *International Conference on Learning Representations (ICLR)*, 2026.
- [3] Joonghyuk Shin, Zhengqi Li, Richard Zhang, Jun-Yan Zhu, Jaesik Park, Eli Shechtman, and Xun Huang. MotionStream: Real-time video generation with interactive motion controls. In *International Conference on Learning Representations (ICLR)*, 2026.
- [4] Guying Lin, Kemeng Huang, Michael Liu, Ruihan Gao, Hanke Chen, Lyuhao Chen, Beijia Lu, Taku Komura, Yuan Liu, Jun-Yan Zhu, and Minchen Li. PAT3D: Physics-augmented text-to-3d scene generation. In *International Conference on Learning Representations (ICLR)*, 2026.
- [5] Sheng-Yu Wang, Aaron Hertzmann, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Fast data attribution for text-to-image models. In *Neural Information Processing Systems (NeurIPS)*, 2025.
- [6] Gaurav Parmar, Or Patashnik, Kuan-Chieh Wang, Daniil Ostashev, Srinivasa Narasimhan, Jun-Yan Zhu, Daniel Cohen-Or, and Kfir Aberman. Visual composer: Object-level visual prompts for compositional image generation. In *ACM Transactions on Graphics (SIGGRAPH Asia)*, 2025.

- [7] Beijia Lu, Ziyi Chen, Jing Xiao, and Jun-Yan Zhu. Input-aware sparse attention for real-time co-speech video generation. In *ACM Transactions on Graphics (SIGGRAPH Asia)*, 2025.
- [8] Ava Pun, Kangle Deng, Ruixuan Liu, Deva Ramanan, Changliu Liu, and Jun-Yan Zhu. Generating physically stable and buildable brick structures from text. In *International Conference on Computer Vision (ICCV)*, 2025. **Best Paper Award, Marr Prize**.
- [9] Nupur Kumari, Xi Yin, Jun-Yan Zhu, Ishan Misra, and Samaneh Azadi. Generating multi-image synthetic data for text-to-image customization. In *International Conference on Computer Vision (ICCV)*, 2025.
- [10] Kangle Deng, Hsueh-Ti Derek Liu, Yiheng Zhu, Xiaoxia Sun, Chong Shang, S. Kiran Bhat, Deva Ramanan, Jun-Yan Zhu, Maneesh Agrawala, and Tinghui Zhou. Efficient autoregressive shape generation via octree-based adaptive tokenization. In *International Conference on Computer Vision (ICCV)*, 2025.
- [11] Sean J Liu, Nupur Kumari, Ariel Shamir, and Jun-Yan Zhu. Generative photomontage. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [12] Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Yuwei Fang, Kwot Sin Lee, Ivan Skorokhodov, Kfir Aberman, Jun-Yan Zhu, Ming-Hsuan Yang, and Sergey Tulyakov. Multi-subject open-set personalization in video generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [13] MUYANG LI*, YUJUN LIN*, ZHEKAI ZHANG*, Tianle Cai, XIUYU LI, JUNXIAN GUO, ENZE XIE, CHENLIN MENG, JUN-YAN ZHU, and SONG HAN. Svdquant: Absorbing outliers by low-rank components for 4-bit diffusion models. In *International Conference on Learning Representations (ICLR)*, 2025.
- [14] Songwei Ge, Taesung Park, Jun-Yan Zhu, and Jia-Bin Huang. Expressive text-to-image generation with rich text. *International Journal of Computer Vision (IJCV)*, 2025.
- [15] Sheng-Yu Wang, Aaron Hertzmann, Alexei A Efros, Jun-Yan Zhu, and Richard Zhang. Data attribution for text-to-image models by unlearning synthesized images. In *Neural Information Processing Systems (NeurIPS)*, 2024.
- [16] Ruihan Gao, Kangle Den, Gengshan Yang, Wenzhen Yuan, and Jun-Yan Zhu. Tactile dreamfusion: Exploiting tactile sensing for 3d generation. In *Neural Information Processing Systems (NeurIPS)*, 2024.
- [17] Nupur Kumari, Grace Su, Richard Zhang, Taesung Park, Eli Shechtman, and Jun-Yan Zhu. Customizing text-to-image diffusion with object viewpoint control. In *SIGGRAPH Asia*, 2024.
- [18] Maxwell Jones, Sheng-Yu Wang, Nupur Kumari, David Bau, and Jun-Yan Zhu. Customizing text-to-image models with a single image pair. In *SIGGRAPH Asia*, 2024.
- [19] Or Patashnik, Rinon Gal, Daniel Cohen-Or, Jun-Yan Zhu, and Fernando De la Torre. Consolidating attention features for multi-view image editing. In *SIGGRAPH Asia*, 2024.
- [20] Minguk Kang, Richard Zhang, Connelly Barnes, Sylvain Paris, Suha Kwak, Jaesik Park, Eli Shechtman, Jun-Yan Zhu, and Taesung Park. Distilling Diffusion Models into Conditional GANs. In *European Conference on Computer Vision (ECCV)*, 2024.
- [21] Kangle Deng, Timothy Omernick, Alexander Weiss, Deva Ramanan, Jun-Yan Zhu, Tinghui Zhou, and Maneesh Agrawala. Flashtex: Fast relightable mesh texturing with lightcontrolnet. In *European Conference on Computer Vision (ECCV)*, 2024.
- [22] Aniruddha Mahapatra, Richa Mishra, Renda Li, Ziyi Chen, Boyang Ding, Shoulei Wang, Jun-Yan Zhu, Peng Chang, Mei Han, and Jing Xiao. Co-speech gesture video generation with 3d human meshes. In *European Conference on Computer Vision (ECCV)*, 2024.
- [23] Peter Schaldenbrand, Gaurav Parmar, Jun-Yan Zhu, James McCann, and Jean Oh. Cofrida: Self-supervised fine-tuning for human-robot co-painting. In *International Conference on Robotics*

and Automation (*ICRA*), 2024. **Best Paper on Human-Robot Interaction and Best Demo Award Finalist.**

- [24] Songwei Ge, Aniruddha Mahapatra, Gaurav Parmar, Jun-Yan Zhu, and Jia-Bin Huang. On the content bias in fr chet video distance. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [25] Gaurav Parmar, Taesung Park, Srinivasa Narasimhan, and Jun-Yan Zhu. One-step image translation with text-to-image models. *arXiv preprint arXiv:2403.12036*, 2024.
- [26] Sheng-Yu Wang, Alexei A Efros, Jun-Yan Zhu, and Richard Zhang. Evaluating data attribution for text-to-image models. In *International Conference on Computer Vision (ICCV)*, 2023.
- [27] Nupur Kumari, Bingliang Zhang, Sheng-Yu Wang, Eli Shechtman, Richard Zhang, and Jun-Yan Zhu. Ablating concepts in text-to-image diffusion models. In *International Conference on Computer Vision (ICCV)*, 2023.
- [28] Hyunsu Kim, Gayoung Lee, Yunjey Choi, Jin-Hwa Kim, and Jun-Yan Zhu. 3d-aware blending with generative nerfs. In *International Conference on Computer Vision (ICCV)*, 2023.
- [29] Yunji Kim, Jiyoung Lee, Jin-Hwa Kim, Jung-Woo Ha, and Jun-Yan Zhu. Dense text-to-image generation with attention modulation. In *International Conference on Computer Vision (ICCV)*, 2023.
- [30] Chonghyuk Song, Gengshan Yang, Kangle Deng, Jun-Yan Zhu, and Deva Ramanan. Total-recon: Deformable scene reconstruction for embodied view synthesis. In *International Conference on Computer Vision (ICCV)*, 2023.
- [31] Songwei Ge, Taesung Park, Jun-Yan Zhu, and Jia-Bin Huang. Expressive text-to-image generation with rich text. In *International Conference on Computer Vision (ICCV)*, 2023.
- [32] Ruihan Gao, Wenzhen Yuan, and Jun-Yan Zhu. Controllable visual-tactile synthesis. In *International Conference on Computer Vision (ICCV)*, 2023.
- [33] Tony Lee*, Michihiro Yasunaga*, Chenlin Meng*, Yifan Mai, Joon Sung Park, Yunzhi Zhang Agrim Gupta, Deepak Narayanan, Hannah Benita Teufel, Marco Bellagente, Minguk Kang, Taesung Park, Jure Leskovec, Jun-Yan Zhu, Li Fei-Fei, Jiajun Wu, Stefano Ermon, and Percy Liang. Holistic evaluation of text-to-image models. In *Neural Information Processing Systems (NeurIPS)*, 2023.
- [34] Daohan Lu*, Sheng-Yu Wang*, Nupur Kumari*, Rohan Agarwal*, Mia Tang, David Bau, and Jun-Yan Zhu. Content-based search for deep generative models. In *ACM SIGGRAPH Asia*, 2023.
- [35] Aniruddha Mahapatra, Aliaksandr Siarohin, Hsin-Ying Lee, Sergey Tulyakov, and Jun-Yan Zhu. Text-Guided Synthesis of Eulerian Cinemagraphs. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 2023.
- [36] Gaurav Parmar, Krishna Kumar Singh, Richard Zhang, Yijun Li, Jingwan Lu, and Jun-Yan Zhu. Zero-shot image-to-image translation. In *ACM SIGGRAPH*, 2023.
- [37] Muyang Li, Ji Lin, Chenlin Meng, Stefano Ermon, Song Han, and Jun-Yan Zhu. Efficient spatially sparse inference for conditional gans and diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2023.
- [38] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [39] Kangle Deng, Gengshan Yang, Deva Ramanan, and Jun-Yan Zhu. 3d-aware conditional image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

- [40] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [41] Yotam Nitzan, Michaël Gharbi, Richard Zhang, Taesung Park, Jun-Yan Zhu, Daniel Cohen-Or, and Eli Shechtman. Domain expansion of image generators. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [42] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A. Efros, and Jun-Yan Zhu. Generalizing dataset distillation via deep generative prior. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [43] Muyang Li, Ji Lin, Chenlin Meng, Stefano Ermon, Song Han, and Jun-Yan Zhu. Efficient spatially sparse inference for conditional gans and diffusion models. In *Neural Information Processing Systems (NeurIPS)*, 2022.
- [44] Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. Rewriting geometric rules of a gan. *ACM Transactions on Graphics (SIGGRAPH)*, 41(4), 2022.
- [45] William Peebles, Jun-Yan Zhu, Richard Zhang, Antonio Torralba, Alexei A. Efros, and Eli Shechtman. Gan-supervised dense visual alignment. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. **Best Paper Finalist**.
- [46] Nupur Kumari, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Ensembling off-the-shelf models for gan training. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [47] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [48] George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A. Efros, and Jun-Yan Zhu. Dataset distillation by matching training trajectories. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [49] Gaurav Parmar, Yijun Li, Richard Zhang Jingwan Lu, Jun-Yan Zhu, and Krishna Kumar Singh. Multilayer gan inversion and editing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [50] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised NeRF: Fewer views and faster training for free. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [51] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2022.
- [52] Muyang Li, Ji Lin, Yaoyao Ding, Zhijian Liu, Jun-Yan Zhu, and Song Han. Gan Compression: Efficient architectures for interactive conditional gans. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2021.
- [53] Sheng-Yu Wang, David Bau, and Jun-Yan Zhu. Sketch your own GAN. In *International Conference on Computer Vision (ICCV)*, 2021.
- [54] 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.
- [55] 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.

- [56] Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Ensembling with deep generative views. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [57] 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.
- [58] 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 Systems (NeurIPS)*, 2020.
- [59] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- [60] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision (ECCV)*, 2020.
- [61] 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.
- [62] 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.
- [63] 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.
- [64] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, Rohit Pandey, Sean Fanello, Gordon Wetzstein, Jun-Yan Zhu, Christian Theobalt, Maneesh Agrawala, Eli Shechtman, Dan B Goldman, and Michael Zollhöfer. State of the Art on Neural Rendering. *Computer Graphics Forum (EuroGraphics STAR)*, 2020.
- [65] 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.
- [66] 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.
- [67] 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.
- [68] 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.
- [69] 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.
- [70] 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.
- [71] 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.

- [72] 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**.
- [73] 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.
- [74] 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.
- [75] 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 Systems (NeurIPS)*, 2018.
- [76] 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 Systems (NeurIPS)*, 2018.
- [77] 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 Systems (NeurIPS)*, 2018.
- [78] 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.
- [79] 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.
- [80] 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.
- [81] 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.
- [82] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv preprint arXiv:1811.10959*, 2018.
- [83] 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 Systems (NeurIPS)*, 2017.
- [84] 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.
- [85] 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.
- [86] 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)*, 9(4), 2017.
- [87] 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)*, 36(4), 2017.

- [88] 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.
- [89] 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.
- [90] 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.
- [91] 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.
- [92] Jun-Yan Zhu, Aseem Agarwala, Alexei A Efros, Eli Shechtman, and Jue Wang. Mirror mirror: Crowdsourcing better portraits. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 33(6), 2014.
- [93] Jun-Yan Zhu, Yong Jae Lee, and Alexei A Efros. AverageExplorer: Interactive exploration and alignment of visual data collections. *ACM Transactions on Graphics (SIGGRAPH)*, 33(4), 2014.
- [94] 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.
- [95] 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.
- [96] Tao Chen, Jun-Yan Zhu, Ariel. Shamir, and Shi-Min Hu. Motion-aware gradient domain video composition. *IEEE Transactions on Image Processing (TIP)*, 2013.
- [97] 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.
- [98] 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.

--- Academic Service

Area Chair/Editor

- 2026 Best Paper Award Committee, CVPR 2026
- 2026 Technical Papers Sorter, SIGGRAPH Asia 2026
- 2026 Senior Area chair, CVPR 2026
- 2025 Tutorial chair, ICCV 2025
- 2025 Area chair, CVPR 2025
- 2024 Area chair, CVPR 2024
- 2024 Technical Papers Committee member, SIGGRAPH Asia 2024
- 2023 Area chair, ICCV 2023
- 2026 Best Paper Award Committee, SIGGRAPH Asia 2023
- 2023 Technical Papers Committee member, SIGGRAPH Asia 2023
- 2022 Area chair, NeurIPS 2022
- 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
- 2018-20 Guest editor, International Journal of Computer Vision (IJCV)
- 2018 Technical Papers Committee member, SIGGRAPH Asia 2018
- 2019- Evaluation Committee member, Graphics Replicability Stamp Initiative
- Workshop/Tutorial/Course**
- 2026 Organizer, CVPR 2026 Workshop on AI for Content Creation
- 2025 Organizer, SIGGRAPH 2025 Workshop AI for Creative Visual Content Generation, Editing, and, Understanding
- 2025 Organizer, CVPR 2025 Workshop on AI for Content Creation
- 2024 Organizer, CVPR 2024 Workshop on AI for Content Creation
- 2023 Organizer, SIGGRAPH 2023 Course on Diffusion Models
- 2023 Organizer, CVPR 2023 Workshop on AI for Content Creation
- 2022 Organizer, CVPR 2022 Workshop on AI for Content Creation
- 2022 Organizer, CVPR 2022 Workshop on Sketch-Oriented Deep Learning
- 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

Nature, Science, IJCV, TPAMI, ACM Transactions on Graphics, Foundations and Trends in Computer Graphics and Vision

CVPR (Outstanding Reviewer Award 2017, 2019), ICCV (Outstanding Reviewer Award 2021), ECCV, SIGGRAPH, SIGGRAPH Asia, Eurographics, ICML, NeurIPS, CHI

Teaching

- 2021- **Instructor**, 16-824: Visual Learning and Recognition (Fall 2021-2025)
- 2021- **Instructor**, 16-726: Learning-based Image Synthesis (Spring 2021-2025)
- 2018 **Co-instructor**, Deep Learning (800 enrolled students), Udacity.

Honors and Awards

- 2026 Sloan Research Fellowships
- 2026 Frontiers of Science Award
- 2025 ICCV 2025 Best Paper Award (the Marr Prize)
- 2025 Frontiers of Science Award
- 2024 Samsung AI Researcher of the Year
- 2024 ICRA Best Paper on Human-Robot Interaction
- 2024 ICRA EXPO 2024 Best Demo Award Finalist
- 2023 Packard Fellowships for Science and Engineering

- 2023 NSF CAREER Award
- 2022 Amazon Faculty Research Award
- 2022 JPMorgan Chase Faculty Research Award
- 2022 CVPR 2022 Best Paper Finalist
- 2019-2022 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
- 2019 CVPR 2019 Best Paper Finalist
- 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

Selected Press

- 2025 CNET: You Can Now Use LegoGPT AI to Create Your Own Lego Designs.
- 2024 CMU: RI Research Brings Together Humans, Robots and Generative AI To Create Art.
- 2023 Quartz: How should creators be compensated for their work training AI models?
- 2023 Reuters: Adobe, Nvidia AI imagery systems aim to resolve copyright questions.
- 2023 CMU News: Addressing Copyright, Compensation Issues in Generative AI.
- 2023 WIRED: Where the AI Art Boom Came From—and Where It's Going?
- 2020 MIT Technology Review: A neural network can learn to organize the world it sees into concepts.
- 2019 CNN: MIT teaches robots to 'feel' objects just by looking at them.
- 2019 BBC News: Robot taught to 'feel' objects by sight.
- 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.
- 2018 The New Yorker: In the Age of A.I., Is Seeing Still Believing?
- 2018 The New York Times: How an A.I. 'Cat-and-Mouse Game' Generates Believable Fake Photos.
- 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.