Richard Zhang

Email (rizhang@adobe.com) • Homepage • GitHub • Scholar
Last updated [Jul 2025]

SUMMARY

Researching controllability, inference speed/compression, and data transparency for generative modeling (GenAI). Broader interests in computer vision, machine learning, and graphics, with contributions in self-supervised learning, perceptual metrics, and image forensics. Recognized as an Innovator Under 35 by MIT Tech Review in 2023.

EMPLOYMENT HISTORY

Adobe Research

Senior Research Scientist II, San Francisco, CA — Jan 2024 – Present
Senior Research Scientist — Jan 2022 – Dec 2023
Research Scientist II — Jan 2020 – Dec 2021
Research Scientist — May 2018 – Dec 2019
Research Intern, Seattle, WA — May – Aug 2017

Johns Hopkins University Applied Physics Laboratory (JHU/APL), Laurel, MD — Jul 2010 – Jul 2012
- Missile Defense Radar Engineering Group, Air & Missile Defense Dept (AMDD), *Staff Engineer*
- Electro-Optical & Infrared Systems and Technologies Group, AMDD

EDUCATION

University of California, Berkeley, Berkeley, CA
- Ph.D. in Electrical Engineering and Computer Sciences (EECS) — Aug 2012 – May 2018
  - Thesis: Image Synthesis for Self-Supervised Visual Representation Learning
  - Advisor: Prof. Alexei A. Efros

Cornell University, Ithaca, NY
- M.Eng. in Electrical & Computer Engineering (ECE) — Aug 2009 – May 2010
  - Cumulative GPA: 4.13 / 4.30
- B.S. in Electrical & Computer Engineering (ECE) — Aug 2006 – Dec 2009
  - Cumulative GPA: 4.02 / 4.30, Summa Cum Laude, Dean's List all semesters

PUBLICATIONS

CONFERENCE

[57] R. Gandikota, Z. Wu, R. Zhang, D. Bau, E. Shechtman, N. Kolkin. ***SliderSpace: Decomposing the Visual Capabilities of Diffusion Models.*** In *ICCV*, 2025.

[56] R. Po, Y. Nitzan, R. Zhang, B. Chen, T. Dao, E. Shechtman, G. Wetzstein, X. Huang. ***Long-Context State-Space Video World Models.*** In *ICCV*, 2025.

[55] T. Yin, Q. Zhang, R. Zhang, W.T. Freeman, F. Durand, E. Shechtman, X. Huang. ***From slow bidirectional to fast causal video generators.*** In *CVPR*, 2025.

[54] Y. Xu, T. Park, R. Zhang, Y. Zhou, E. Shechtman, F. Liu, J.B. Huang, D. Liu. In CVPR, 2025. ***VideoGigaGAN: Towards Detail-rich Video Super-Resolution.*** In *CVPR*, 2025.

[53] S.Y. Wang, A. Hertzmann, A.A. Efros, J.Y. Zhu, R. Zhang. ***Data Attribution for Text-to-Image Models by Unlearning Synthesized Images.*** In *NeurIPS*, 2024.

[52] T. Yin, M. Gharbi, T. Park, R. Zhang, E. Shechtman, F. Durand, W. T. Freeman. ***Improved Distribution Matching Distillation for Fast Image Synthesis.*** In *NeuIPS* (oral), 2024.

[51] J. Materzynska, J. Sivic, E. Shechtman, A. Torralba, R. Zhang, B. Russell. ***Customizing Motion in Text-to-Video Diffusion Models.*** In *ACCV*, 2024.

[50] N. Kumari, G. Su, R. Zhang, T. Park, J.Y. Zhu. ***Customizing Text-to-Image Diffusion with Camera Viewpoint Control.*** In *SIGGRAPH Asia*, 2024.

[49] Z. Wu, N. Kolkin, J. Brandt, R. Zhang, E. Shechtman. ***TurboEdit: Real-time text-based disentangled real image editing.*** In *ECCV*, 2024.

[48] M. Kang, R. Zhang, C. Barnes, S. Paris, S. Kwak, J. Park, E. Shechtman, J.Y. Zhu, T. Park. ***Diffusion2GAN: Distilling Diffusion Models into Conditional GANs.*** In *ECCV*, 2024.

[47] J. Mu, M. Gharbi, R. Zhang, E. Shechtman, N. Vasconcelos, X. Wang, T. Park. ***Editable Image Elements for Controllable Synthesis.*** In *ECCV*, 2024.

[46] Y. Nitzan, Z. Wu, R. Zhang, E. Shechtman, D. Cohen-Or, T. Park, M. Gharbi. ***Lazy Diffusion Transformer for Interactive Image Editing.*** In *ECCV*, 2024.

[45] A. Ganjdanesh, Y. Kang, Y. Liu, R. Zhang, Z. Lin, H. Huang. ***Mixture of Efficient Diffusion Experts Through Automatic Interval and Sub-Network Selection.*** In *ECCV*, 2024.

[44] T. Yin, M. Gharbi, R. Zhang, E. Shechtman, F. Durand, W. Freeman, T. Park. ***One-step Diffusion with Distribution Matching Distillation.*** In *CVPR*, 2024.

[43] Y. Chen, O. Wang, R. Zhang, E. Shechtman, X. Wang, M. Gharbi. ***Image Neural Field Diffusion Models.*** In *CVPR* (highlight), 2024.

[42] C. Ham, M. Fisher, J. Hays, N. Kolkin, Y. Lu, R. Zhang, T. Hinz. ***Personalized Residuals for Concept-Driven Text-to-Image Generation.*** In *CVPR*, 2024.

[41] S. Fu*, N. Tamir*, S. Sundaram*, L. Chai, R. Zhang, T. Dekel, P. Isola. ***DreamSim: Learning New Dimensions of Human Visual Similarity using Synthetic Data.*** In *NeurIPS* (spotlight), 2023.

[40] S.Y. Wang, A.A. Efros, J.Y. Zhu, R. Zhang. ***Evaluating Data Attribution for Text-to-Image Models.*** In *ICCV*, 2023.

[39] N. Kumari, B. Zhang, S.Y. Wang, E. Shechtman, R. Zhang, J.Y. Zhu. ***Ablating concepts in text-to-image diffusion models.*** In *ICCV*, 2023.

[38] G. Parmar, K. K. Singh, R. Zhang, Y. Li, J. Lu, J.Y. Zhu. ***Zero-shot Image-to-Image Translation.*** In *SIGGRAPH*, 2023.

[37] N. Kumari, B. Zhang, R. Zhang, E. Shechtman, J.Y. Zhu. ***Multi-Concept Customization of Text-to-Image Diffusion.*** In *CVPR*, 2023.

[36] M. Kang, J.Y. Zhu, R. Zhang, J. Park, E. Shechtman, S. Paris, T. Park. ***Scaling up GANs for Text-to-Image Synthesis.*** In *CVPR* (highlight), 2023.

[35] Y. Nitzan, M. Gharbi, R. Zhang, T. Park, J.Y. Zhu, D. Cohen-Or, E. Shechtman. ***Domain Expansion of Image Generators.*** In *CVPR*, 2023.

[34] M. Huh, H. Mohabi, R. Zhang, B. Cheung, P. Agrawal, P. Isola. ***The Low-Rank Simplicity Bias in Deep Networks.*** In *TMLR*, 2023.

[33] L. Chai, M. Gharbi, E. Shechtman, P. Isola, R. Zhang. ***Any-resolution Training for High-resolution Image Synthesis.*** In *ECCV*, 2022.

[32] D. Epstein. T. Park, R. Zhang, E. Shechtman, A. A. Efros. ***BlobGAN: Spatially Compositional Scene Representations.*** In *ECCV*, 2022.

[31] Y. Liu, Z. Shu, Y. Li, Z. Lin, R. Zhang, S.Y. Kung. ***3D-FM GAN: Towards 3D-Controllable Face Manipulation.*** In *ECCV*, 2022.

[30] D. Liu, S. Shetty, T. Hinz, M. Fisher, R. Zhang, T. Park, E. Kalogerakis. ***ASSET: Autoregressive Semantic Scene Editing with Transformers at High Resolutions.*** In *SIGGRAPH*, 2022.

[29] W. Peebles, J.Y. Zhu, R. Zhang, A. A. Efros, A. Torralba, E. Shechtman. ***GAN-Supervised Dense Visual Alignment*** In *CVPR*, 2022 (oral, best paper finalist).

[28] N. Kumari, R. Zhang, E. Shechtman, J.Y. Zhu. ***Ensembling Off-the-shelf Models for GAN Training.*** In *CVPR*, 2022 (oral).

[27] G. Parmar, R. Zhang, J.Y. Zhu. ***On Aliased Resizing Libraries and Surprising Subtleties in FID Calculation.*** In *CVPR*, 2022.

[26] G. Parmar, Y. Li, J. Lu, R. Zhang, J.Y. Zhu, K. Singh. ***Multilayer GAN Inversion and Editing.*** In *CVPR*, 2022.

[25] S. Liu, X. Zhang, Z. Zhang, R. Zhang, J.Y. Zhu, B. Russell. ***Editing Conditional Radiance Fields.*** In *ICCV*, 2021.

[24] R. Alghofaili, M. Fisher, R. Zhang, M. Lukáč, L.F. Yu. ***Exploring Sketch-based Character Design Guided by Automatic Colorization.*** In *Graphics Interfaces*, 2021.

[23] L. Chai, J.Y. Zhu, E. Shechtman, P. Isola, R. Zhang. ***Ensembling with Deep Generative Views.*** In *CVPR*, 2021.

[22] U. Ojha, Y. Li, J. Lu, A. A. Efros, Y.J. Lee, E. Shechtman, R. Zhang. ***Few-shot Image Generation via Cross-domain Correspondence.*** In *CVPR*, 2021.

[21] J. Lin, R. Zhang, F. Ganz, S. Han, J.Y. Zhu. ***Anycost GANs for Interactive Image Synthesis and Editing.*** In *CVPR*, 2021.

[20] T. R. Shaham, M. Gharbi, R. Zhang, E. Shechtman, T. Michaeli. ***Spatially-Adaptive Pixelwise Networks for Fast Image Translation.*** In *CVPR*, 2021.

[19] P. Manocha, Z. Jin, R. Zhang, A. Finkelstein. ***CDPAM: Contrastive learning for perceptual audio similarity.*** In *ICASSP*, 2021.

[18] Y. Li, R. Zhang, J. Lu, E. Shechtman. ***Few-shot Image Generation with Elastic Weight Consolidation.*** In *NeurIPS*, 2020.

[17] T. Park, J.Y. Zhu, O. Wang, J. Lu, E. Shechtman, A. A. Efros, R. Zhang. ***Swapping Autoencoder for Deep Image Manipulation.*** In *NeurIPS*, 2020.

[16] T. Park, A. A. Efros, R. Zhang, J.Y. Zhu. ***Contrastive Learning for Unsupervised Image-to-Image Translation.*** In *ECCV*, 2020.

[15] M. Huh, R. Zhang, J.Y. Zhu, S. Paris, A. Hertzmann. ***Transforming and Projecting Images into Class-conditional Generative Networks.*** In *ECCV*, 2020 (oral).

[14] P. Manocha, A. Finkelstein, R. Zhang, N. J. Bryan, G. J. Mysore, Z. Jin. ***A Differentiable Perceptual Audio Metric Learned from Just Noticeable Differences.*** In *Interspeech*, 2020.

[13] S. Wang, O. Wang, R. Zhang, A. Owens, A. A. Efros. ***CNN-generated images are surprisingly easy to spot...for now.*** In *CVPR*, 2020 (oral).

[12] D. Smirnov, M. Fisher, V. Kim, R. Zhang, J. Solomon. ***Deep Parametric Shape Predictions using Distance Fields.*** In *CVPR*, 2020.

[11] N. Fish, R. Zhang, L. Perry, D. Cohen-Or, E. Shechtman, C. Barnes. ***Image Morphing with Perceptual Constraints and STN Alignment.*** In *CGF*, 2020.

[10] S. Wang, O. Wang, A. Owens, R. Zhang, A. A. Efros. ***Detecting Photoshopped Faces by Scripting Photoshop.*** In *ICCV*, 2019.

[9] A. Ghosh, R. Zhang, P. K. Dokania, O. Wang, A. A. Efros, P. H.S. Torr, E. Shechtman. ***Interactive Sketch & Fill: Multiclass Sketch-to-Image Translation.*** In *ICCV*, 2019.

[8] R. Zhang. ***Making Convolutional Networks Shift-Invariant Again.*** In *ICML*, 2019.

[7] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, O. Wang. ***The Unreasonable Effectiveness of Deep Features as a Perceptual Metric.*** In *CVPR*, 2018.

[6] J.Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, E. Shechtman. ***Toward Multimodal Image-to-Image Translation.*** In *NIPS*, 2017.

[5] R. Zhang*, J.Y. Zhu*, P. Isola, X. Geng, A. S. Lin, T. Yu, A. A. Efros. ***Real-Time User-Guided Image Colorization with Learned Deep Priors.*** In *SIGGRAPH*, 2017. (*equal contribution)

[4] R. Zhang, P. Isola, A. A. Efros. ***Split-Brain Autoencoders: Unsupervised Learning by Cross-Channel Prediction.*** In *CVPR*, 2017.

[3] R. Zhang, P. Isola, A. A. Efros. ***Colorful Image Colorization.*** In *ECCV*, 2016 (oral).

[2] R. Zhang, S. Candra, K. Vetter, A. Zakhor. ***Sensor Fusion for Semantic Segmentation for Urban Scenes.*** In *ICRA*, 2015.

[1] R. Zhang and A. Zakhor. ***Automatic Identification of Window Regions on Indoor Point Clouds Using LiDAR and Cameras.*** In *WACV*, 2014.

WORKSHOP/PREPRINTS

[v] N.Y. Tamir, S. Amir, R. Itzhaky, N. Atia, S. Sundaram, S. Fu, R. Sokolovsky, P. Isola, T. Dekel, R. Zhang, M. Farber. ***What Makes for a Good Stereoscopic Image?*** In *CVPR CV4Metaverse Workshop*, 2025.

[iv] X. Wang, T. Park, Y. Zhou, E. Shechtman, R. Zhang. ***Jump Cut Smoothing for Talking Heads.*** In *ArXiv*, 2024.

[iii] D. Epstein, I. Jain, O. Wang, R. Zhang. ***Online Detection of AI-Generated Images.*** In *ICCV DFAD Workshop*, 2023.

[ii] A. Andonian, T. Park, B. Russell, P. Isola, J.Y. Zhu, R. Zhang. ***Contrastive Feature Loss for Image Prediction.*** In *ICCV AIM Workshop*, 2021.

[i] A.X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, S. Levine. ***Stochastic Adversarial Video Prediction.*** In *ArXiv*, 2018.

TECH TRANSFERS

Adobe Firefly

- Custom models, extended from Custom Diffusion [CVPR '23] — Apr 2024
- Consulted on pruning for improved inference speed — Nov 2023
- Consulted on upsampling — Feb 2023

GenAI Detection, Adobe Stock — Jan 2023

- Based on CNN and Diffusion detection [CVPR '20; ICCV Wksp '23]

Landscape Mixer, Photoshop Neural Filters — Oct 2021

- Based on Swapping Autoencoder [NeurIPS '19] paper

Smart Portrait, Photoshop Neural Filters — Oct 2021

- Consultation; developed and ran A/B testing procedure

Enhance Super Resolution, Lightroom — Mar 2021
- Consultation on feature, received recognition as one of Best Inventions of 2021 from Time

Colorize Photo, Photoshop Neural Filters, Photoshop Elements — Oct 2020, Oct 2019
- Based on Colorization papers [ECCV '16; SIGGRAPH '17]

AWARDS

35 Innovators Under 35, MIT Technology Review — Sep 2023
Best Paper Finalist, CVPR 2022 — Jul 2022
Paper Reviewing Recognitions
- ECCV, top reviewer — Oct 2022
- NeurIPS, top 10% reviewer — Dec 2020
- ECCV, top reviewer — Oct 2020
- NeurIPS, top 50% reviewer — Dec 2019
- CVPR, outstanding reviewer — Jul 2019

Best Presentation Award, SIGGRAPH Thesis Fast Forward — Jul 2018
Adobe Research Fellowship — Jan 2017
William S. Einwechter Award, Cornell University — May 2010

COMMUNITY SERVICE

AREA CHAIR

European Conference on Computer Vision (ECCV)	2024
Computer Vision and Pattern Recognition (CVPR)	2020-21, 2023-24
International Conference on Learning Representations (ICLR)	2023
British Machine Vision Conference (BMVC)	2022

PAPERS REVIEWED

Computer Vision and Pattern Recognition (CVPR)	2018-19, 2022
European Conference on Computer Vision (ECCV)	2018-22
International Conference on Computer Vision (ICCV)	2017-19, 2023
Neural Information Processing Systems (NIPS, NeurIPS)	2016-21
International Conference in Machine Learning (ICML)	2019-20
Special Interest Group in Graphics (SIGGRAPH)	2017-19, 20-23
Special Interest Group in Graphics, Asia (SIGGRAPH Asia)	2017-19, 2021
International Conference on Robotics and Automation (ICRA)	2015, 2018
International Journal of Computer Vision (IJCV)	2019, 2021
Transactions in Pattern Analysis and Machine Intelligence (TPAMI)	2018
Transactions in Image Processing (TIP)	2017-18
Technical Committee on Vision and Graphics (TCVG)	2018
Pacific Graphics	2018
Eurographics	2019

WORKSHOP ORGANIZATION COMMITTEE

Sketching for Human Expressivity (SHE), at ECCV 2022	Oct 2022
Advancements in Image Manipulation (AIM), at ICCV 2019	Nov 2019
New Trends in Image Restoration and Enhancement (NTIRE), at CVPR 2019	Jul 2019

SELECTED PUBLICITY

TWiML (This Week in ML) Podcast. *Visual Generative AI Ecosystem Challenges.*	Dec 2023
Adobe Blog. *MIT Technology Review names Adobe Research's Richard Zhang a top "Innovator Under 35" for groundbreaking generative AI and image forensics work.*	Sep 2023
Time. *The Best Inventions of 2021: Adobe Super Resolution.*	Nov 2021
Adobe Research Blog. *Advancing the Science of Image Forensics.*	Jan 2020
Adobe MAX (Sneak Peek). *Project About Face.*	Nov 2019
The Verge. *Adobe's prototype AI tool automatically spots Photoshopped faces.*	Jun 2019
The New Yorker. *In the Age of A.I., Is Seeing Still Believing?*	Nov 2018
Adobe Research Blog. *With Deep Learning, Computers See Images More Like Humans Do.*	May 2018
UK Times. *Computers give the past a blast of colour.*	Apr 2016
TechCrunch. *This neural network 'hallucinates' the right colors into black and white pictures.*	Mar 2016

INVITED PRESENTATIONS

Data Attribution for Text-to-Image Models
MIT Computer Vision Seminar — Jul 2024

Incentivizing Opt-in and Enabling Opt-out for Text-to-Image Models
CVPR GenAI Media Gen. Challenge (GAMC), Trusted Computer Vision (TCV) workshops — Jun 2024
NeurIPS ML for Creativity & Design Workshop — Dec 2023

Perception, Generation, & Forensics
Video Quality Experts Group (VQEG) Panel — Jun 2023

Colourisation and Any-resolution Generation
CeADAR Tech Talk — Jan 2023

Anycost and Any-resolution Image Synthesis
CVPR New Trends in Image Restoration (NTIRE), AI for Content Creation (AICC) workshops — Jun 2022
Netflix Seminar — Aug 2022

The Unreasonable Effectiveness of Deep Features as a Perceptual Metric
JPEG Workshop on Subjective Quality Assessment — Jun 2022

Swapping Autoencoder for Deep Image Manipulation
Rework Deep Learning Summit, Generative Models Stage — Jan 2021

Deep Learning for Computer Vision and Graphics
Illinois Mathematics and Science Academy, Intersession — Jan 2021

Detecting Generated Imagery, Deep and Shallow
Learning-Based Image Synthesis, CMU — May 2021
ECCV Sensing, Understanding and Synthesizing Workshop — Aug 2020

Style and Structure Disentanglement for Image Manipulation
ECCV Advances in Image Manipulation (AIM) Workshop — Aug 2020

Analyzing CNN Artifacts in Discriminative and Generative Models
Machine Learning @ Berkeley invited seminar talk — Sep 2020
Graphics and Mixed Environment (GAMES) Webinar — Aug 2020
CVPR Area Chair Workshop — Mar 2020

Making Convolutional Networks Shift-Invariant Again
Simon Fraser University, CMPT 361 Intro to Vision, Invited Lecture — Sep 2020
Berkeley AI Research (BAIR) Seminar — Aug 2019
International Conference on Machine Learning (ICML) — Jun 2019
Google Research, Cambridge, MA — May 2019

Modeling Perceptual Similarity and Shift-Invariance in Deep Networks
NAVER Labs, Tech talk — Oct 2019
University College London, Smart Geometry Processing Group seminar — Oct 2019
Oxford University, VGG seminar — Oct 2019
Scale.AI, seminar talk — Aug 2019
Toyota Technological Institute of Chicago (TTIC), Young Researcher Talk — May 2019
Massachusetts Institute of Technology (MIT), Computer Vision Seminar — Apr 2019

Deep Learning for Content Synthesis
Association for Content Editors (ACE) Tech Day with Adobe — Sep 2019
Hollywood Professional Association (HPA) Tech Retreat — Feb 2019

Image Synthesis for Self-Supervised Visual Representation Learning
Stanford University, Graphics Group; University of Michigan, Computer Vision Group — Jan 2019
Berkeley Special Topics in Deep Learning Seminar, CS 294-131 — Nov 2018
SIGGRAPH 2018 Thesis Fast Forward (3 min) — Jul 2018
Berkeley AI Research (BAIR) Seminar, Dissertation Talk — Apr 2018
Alibaba Research; Amazon AI Deep Learning; DeepScale; Facebook AML; Fyusion; Google Research; Intel Intelligent Systems; NVIDIA Research — Mar 2018

Adobe Research; Allen Institute for AI (AI2); Amazon A9; Apple Turi; eBay Research; Snap Research; WaveOne — Feb 2018

Multimodal Image-to-Image Translation
University of Washington, Graphics and Imaging Lab (GRAIL) — Jul 2018

Real-Time User-Guided Image Colorization with Learned Deep Priors
Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH) — Aug 2017
NVIDIA SIGGRAPH Innovation Theater — Aug 2017

Cross-Channel Visual Prediction
Graphics and Mixed Environment (GAMES) Webinar — Oct 2017
Global AI Hackathon Webinar — Jun 2017
Berkeley AI Research (BAIR) Seminar — Apr 2017

Colorful Image Colorization
Berkeley AI Research (BAIR) Seminar — Sep 2017
European Conference on Computer Vision (ECCV) — Oct 2016
Oxford University; INRIA Paris; INRIA Sophia Antipolis; École des Ponts ParisTech — Jun 2016

Sensor Fusion for Semantic Segmentation for Urban Scenes
Berkeley Deep Drive (BDD) Kickoff — Mar 2016
Amazon Computer Vision PhD Symposium — Oct 2015
International Conference on Robotics and Automation (ICRA) — Mar 2015

Automatic Identification of Window Regions on Indoor Point Clouds Using LiDAR and Cameras
Winter Conference on Applications of Computer Vision (WACV) — May 2014
Microsoft Research (MSR) Computer Vision Group — Jan 2014

TEACHING EXPERIENCE

Berkeley EECS Department

- CS 188 Intro to Artificial Intelligence, *Graduate Student Instructor* — Jan – May 2017
  - Instructor: Prof. Anca Dragan
- CS 280 Computer Vision, *Graduate Student Instructor* — Jan – May 2016
  - Instructor: Prof. Alexei A. Efros

Cornell ECE Department

- ECE 2100 Intro to Circuits, *Teaching Assistant* — Jan – May 2010
  - Instructor: Prof. Alyosha Molnar
- ECE 2100 Intro to Circuits, *Course Assistant* — Aug – Dec 2008
  - Instructor: Prof. John Belina

VOLUNTEER EXPERIENCE

Illinois Math and Science Academy (IMSA), *Intersession Instructor* — Jan 2014, Jan 2021
Berkeley AI Research (BAIR) Mentorship Program, *Mentor* — Aug – Dec 2017
Clarksville Middle School, Howard County Public School System, *Volunteer* — Dec 2010 – May 2011

LANGUAGES

Chinese (Mandarin) – Conversational