

RUIQI GAO

ruiqigao.github.io
299 Fremont Street, San Francisco, CA 94105
ruiqingao@gmail.com • (310)592-2361

EDUCATION	University of California, Los Angeles Ph.D. in Statistics Advisor: Song-Chun Zhu	Sep 2016 – Jun 2021
	Peking University B.S. in School of Mathematics Major in Statistics	Sep 2012 – Jul 2016
RESEARCH INTERESTS	Generative modeling, representation learning.	
EXPERIENCE	Senior Research Scientist Google DeepMind	Jul 2023 – present
	Research Scientist Google DeepMind	Jul 2021 – Jul 2023
	Research Assistant Center for Vision, Cognition, Learning and Autonomy, UCLA Advisor: Prof. Song-Chun Zhu	Sep 2016 – Jun 2021
	Research Intern Google Research, Brain team Mentors: Diederik P. Kingma and Ben Poole	Jun 2020 – Sep 2020
	Research Intern Google Research, Brain team Mentors: Zhen Xu and Andrew M. Dai	Jun 2019 – Sep 2019
	Research Assistant Vision Research Group, Peking University Advisor: Prof. Yizhou Wang	Sep 2015 – Jun 2016
	Research Assistant Junction of Statistics and Biology, UCLA Advisor: Prof. Jingyi Jessica Li	Jun 2015 – Sep 2015
PROJECTS	Veo 3: Google DeepMind’s state-of-the-art video generation model, with audio on	2025
	Veo 2: Google DeepMind’s state-of-the-art video generation model	2024
PUBLICATIONS	* denotes equal contributions	
	[42] S. Szymanowicz, J. Y. Zhang, P. Srinivasan, R. Gao , A. Brussee, A. Holynski, R. Martin-Brualla, J. T. Barron, P. Henzler. “Bolt3d: Generating 3d scenes in seconds”. <i>International Conference on Computer Vision (ICCV)</i> , 2025.	
	[41] R. Gao , E. Hoogeboom, J. Heek, V. De Bortoli, K. P. Murphy, T. Salimans. “Diffusion models and gaussian flow matching: Two sides of the same coin”. <i>International Conference on Learning Representations (ICLR)</i> blogpost track, 2025.	
	[40] D. Xu, R. Gao , W.-H. Zhang, X.-X. Wei, Y. N. Wu. “On Conformal Isometry of Grid Cells: Learning Distance-Preserving Position Embedding”. <i>International Conference on Learning Representations (ICLR)</i> , 2025. [Oral]	

- [39] E. Hoogeboom, T. Mensink, J. Heek, K. Lamerigts, **R. Gao**, T. Salimans. “Simpler Diffusion (SiD2): 1.5 FID on ImageNet512 with pixel-space diffusion”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [38] R. Wu, **R. Gao**, B. Poole, A. Trevithick, C. Zheng, J. T. Barron, A. Holynski. “CAT4D: Create Anything in 4D with Multi-View Video Diffusion Models”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. [Oral]
- [37] A. Trevithick, R. Paiss, P. Henzler, D. Verbin, R. Wu, H. Alzayer, **R. Gao**, B. Poole, J. T. Barron, A. Holynski, R. Ramamoorthi, P. P. Srinivasan. “SimVS: Simulating World Inconsistencies for Robust View Synthesis”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [36] Z. Zhu, Y. Zou, C. M. Jiang, B. Sun, V. Casser, X. Huang, J. Wang, Z. Yang, **R. Gao**, L. Guibas, M. Tan, D. Anguelov. “SceneCrafter: Controllable Multi-View Driving Scene Editing”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [35] **R. Gao***, A. Holynski*, P. Henzler, A. Brussee, R. Martin-Brualla, P. Srinivasan, J. T. Barron, B. Poole*. “CAT3D: Create Anything in 3D with Multi-View Diffusion Models”. *Conference on Neural Information Processing Systems (NeurIPS)*, 2024. [Oral]
- [34] S. Xie, Z. Xiao, D. P. Kingma, T. Hou, Y. N. Wu, K. P. Murphy, T. Salimans, B. Poole, **R. Gao**. “EM Distillation for One-step Diffusion Models”. *Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [33] A. Comas-Massagué, D. Qiu, M. Chai, M. Bühler, A. Raj, **R. Gao**, Q. Xu, M. Matthews, P. Gotardo, S. Orts-Escolano, T. Beeler. “MagicMirror: Fast and High-Quality Avatar Generation with a Constrained Search Space”. *European Conference on Computer Vision (ECCV)*, 2024.
- [32] Y. Zhu, J. Xie, Y. Wu, **R. Gao**. “Learning Energy-Based Models by Cooperative Diffusion Recovery Likelihood”. *International Conference on Learning Representations (ICLR)*, 2024.
- [31] Y. Zhu, Z. Dou, H. Zheng, Y. Zhang, Y. N. Wu, **R. Gao**. “Think Twice Before You Act: Improving Inverse Problem Solving With MCMC”. *arXiv*, 2024.
- [30] C.-Y. Cheng, **R. Gao**, F. Huang, Y. Li. “CoLay: Controllable Layout Generation through Multi-conditional Latent Diffusion”. *arXiv*, 2024.
- [29] P. Yu*, D. Zhang*, H. He*, X. Ma, R. Miao, Y. Lu, Y. Zhang, D. Kong, **R. Gao**, J. Xie, G. Cheng, Y. N. Wu. “Latent Energy-Based Odyssey: Black-Box Optimization via Expanded Exploration in the Energy-Based Latent Space”. *arXiv*, 2024.
- [28] R. Wu*, B. Mildenhall*, P. Henzler, K. Park, **R. Gao**, D. Watson, P. P. Srinivasan, D. Verbin, J. T. Barron, B. Poole, A. Holynski*. “ReconFusion: 3D Reconstruction with Diffusion Priors”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [27] D. P. Kingma, **R. Gao**. “Understanding Diffusion Objectives as the ELBO with Simple Data Augmentation”. *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. [Oral]
- [26] P. Yu, Y. Zhu, S. Xie, X. S. Ma, **R. Gao**, S.-C. Zhu, Y. N. Wu. “Learning Energy-Based Prior Model with Diffusion-Amortized MCMC”. *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- [25] C. Meng, R. Rombach, **R. Gao**, D. P. Kingma, S. Ermon, J. Ho, T. Salimans. “On Distillation of Guided Diffusion Models”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [Award candidate]
- [24] J. Ho*, W. Chan*, C. Saharia*, J. Whang*, **R. Gao**, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, T. Salimans. “Imagen Video: High Definition Video Generation with Diffusion Models”. *arXiv*, 2022.
- [23] D. Xu*, **R. Gao***, W.-H. Zhang, X.-X. Wei, Y. N. Wu. “Conformal Isometry of Lie Group Representation in Recurrent Network of Grid Cells”. *NeurIPS workshop on Symmetry and Geometry in Neural Representations*, 2022.

- [22] P. Yu, S. Xie, X. Ma, B. Jia, B. Pang, **R. Gao**, Y. Zhu, S.-C. Zhu, Y. N. Wu. “Latent Diffusion Energy-based Model for Interpretable Text Modeling”. *International Conference on Machine Learning (ICML)*, 2022.
- [21] E. Nijkamp*, **R. Gao***, P. Sountsov, S. Vasudevan, B. Pang, S.-C. Zhu, Y. N. Wu. “Learning Energy-Based Model with Flow-based Backbone by Neural Transport MCMC”. *International Conference on Learning Representations (ICLR)*, 2022.
- [20] **R. Gao**, J. Xie, S. Huang, Y. Ren, S.-C. Zhu, Y. N. Wu. “Learning Vector Representation of Local Content and Matrix Representation of Local Motion, with Implications for V1”. *AAAI Conference on Artificial Intelligence (AAAI)*, 2022. **[Oral]**
- [19] **R. Gao**, J. Xie, X.-X. Wei, S.-C. Zhu, Y. N. Wu. “On Path Integration of Grid Cells: Group Representation and Isotropic Scaling”. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [18] Y. Zhu, **R. Gao**, S. Huang, S.-C. Zhu, Y. N. Wu. “Learning Neural Representation of Camera Pose with Matrix Representation of Pose Shift via View Synthesis”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [17] **R. Gao**, Y. Song, B. Poole, Y. N. Wu, D. P. Kingma. “Learning Energy-Based Models by Diffusion Recovery Likelihood”. *International Conference on Learning Representations (ICLR)*, 2021.
- [16] **R. Gao**, E. Nijkamp, D. P. Kingma, Z. Xu, A. M. Dai, Y. N. Wu. “Flow Contrastive Estimation of Energy-Based Models”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. **[Oral]** *NeurIPS workshop on Bayesian Deep Learning*, 2019. **[Spotlight]**
- [15] J. Xie*, Z. Zheng*, **R. Gao**, W. Wang, S.-C. Zhu, Y. N. Wu. “Generative VoxelNet: Learning Energy-Based Models for 3D Shape Synthesis and Analysis”. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [14] X. Xing, **R. Gao**, T. Han, S.-C. Zhu, Y. N. Wu. “Deformable Generator Networks: Unsupervised Disentanglement of Appearance and Geometry”. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2020.
- [13] J. Xie*, **R. Gao***, Z. Zheng, S.-C. Zhu, Y. N. Wu. “Motion-Based Generator Model: Unsupervised Disentanglement of Appearance, Trackable and Intrackable Motions in Dynamic Patterns”. *AAAI Conference on Artificial Intelligence (AAAI)*, 2020. **[Oral]**
- [12] J. Xie, **R. Gao**, E. Nijkamp, S.-C. Zhu, Y. N. Wu. “Representation learning: a statistical perspective”. *Annual Review of Statistics and Its Application (ARSIA)*, 2019.
- [11] **R. Gao***, J. Xie*, S.-C. Zhu, Y. N. Wu. “Learning Grid Cells as Vector Representation of Self-Position Coupled with Matrix Representation of Self-Motion”. *International Conference on Learning Representations (ICLR)*, 2019.
- [10] J. Xie*, **R. Gao***, Z. Zheng, S.-C. Zhu, Y. N. Wu. “Learning Dynamic Generator Model by Alternating Back-Propagation Through Time”. *AAAI Conference on Artificial Intelligence (AAAI)*, 2019. **[Spotlight]**
- [9] X. Xing, T. Han, **R. Gao**, S.-C. Zhu, Y. N. Wu. “Unsupervised disentanglement of appearance and geometry by deformable generator network”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [8] **R. Gao***, Y. Lu*, J. Zhou, S.-C. Zhu, Y. N. Wu. “Learning Energy-Based Models as Generative ConvNets via Multi-grid Modeling and Sampling”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. **[Spotlight]**
- [7] Y. N. Wu, **R. Gao**, T. Han, S.-C. Zhu. “A Tale of Three Probabilistic Families: Discriminative, Descriptive and Generative Models”. *Quarterly of Applied Mathematics (QAM)*, 2019.
- [6] J. Xie, Y. Lu, **R. Gao**, S.-C. Zhu, Y. N. Wu. “Cooperative Training of Descriptor and Generator Networks”. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018.

- [5] J. Xie*, Z. Zheng*, **R. Gao**, W. Wang, S.-C. Zhu, Y. N. Wu. “Learning Descriptor Networks for 3D Shape Synthesis and Analysis”. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [Oral]
- [4] J. Xie, Y. Lu, **R. Gao**, S.-C. Zhu, Y. N. Wu. “Cooperative Learning of Energy-Based Model and Latent Variable Model via MCMC Teaching”. *AAAI Conference on Artificial Intelligence (AAAI)*, 2018. [Oral]
- [3] Y. Lu, **R. Gao**, S.-C. Zhu, Y. N. Wu. “Exploring Generative Perspective of Convolutional Neural Networks by Learning Random Field Models”. *Statistics and Its Interface*, 2018.
- [2] S. Li, X. Dou, **R. Gao**, X. Ge, M. Qian, L. Wan. “A remark on copy number variation detection methods”. *PLoS One*, 2018.
- [1] **R. Gao**, J. J. Li. “Correspondence of *D. melanogaster* and *C. elegans* developmental stages revealed by alternative splicing dynamics of conserved exon”. *BMC Genomics*, 2017.

ACADEMIC HONORS & AWARDS	CVPR award candidate	2023
	CVPR outstanding reviewer	2021
	Most Promising Statistician Award, UCLA	2019
	Doctoral Student Travel Award, UCLA	2017 - 2020
	Excellent College Graduate Award, Peking University	2016
	Guanghua Scholarship, Peking University	2015
	Yizheng Alumni Scholarship, Peking University	2014
	May Fourth Scholarship, Peking University	2013
	Entrance Scholarship, Peking University	2012
INVITED TALKS	Keynote talk at International Joint Conference on Neural Networks (IJCNN)	Jul 2025
	Nvidia Research, Mountain View, California.	Aug 2024
	New York University.	Aug 2024
	Generative models and uncertainty quantification, Copenhagen.	Jul 2024
	Guest lecture at CSE6243: Advanced Topics in Machine Learning, Georgia Institute of Technology.	Oct 2023
	University of Michigan, MIDAS mini-symposium.	Sep 2023
	Guest lecture at CS231n: Deep Learning for Computer Vision, Stanford University.	May 2023
	University of Texas Southwestern, Texas.	Sep 2022
	University of California Irvine, California.	Feb 2022
	Amazon, California.	Jan 2021
	Facebook AI Research, California.	Jan 2021
	Nvidia Research, Mountain View, California.	Nov 2020
	Google Brain, Mountain View, California.	Oct 2020
	Institute for Artificial Intelligence, Peking University.	Oct 2020
	Google Brain, Mountain View, California.	Sep 2019
	Guest lecture at CS276A: Pattern Recognition and Machine Learning, UCLA.	Sep 2018
	Guest lecture at CS266A: Statistical Modeling and Learning in Vision and Cognition, UCLA.	Mar 2018
PROFESSIONAL SERVICES	Conference Organization	
	Organizer, CoRL 2025 workshop on “Robotics World Modeling”.	
	Organizer, ICML 2025 workshop on “Building Physically Plausible World Models”.	
	Organizer, CVPR 2025 workshop on “Efficient and On-Device Generation”.	
	Organizer, ICML 2024 workshop on “Structured Probabilistic Inference & Generative Modeling”.	
	Organizer, CVPR 2024 workshop on “Efficient and On-Device Generation”.	
	Organizer and Speaker, NeurIPS 2023 tutorial on “Latent Diffusion Models: Is the Generative AI Revolution Happening in Latent Space?”	
	Organizer and Speaker, CVPR 2022 tutorial on “Denoising Diffusion-based Generative Modeling: Foundations and Applications”.	

Peer-reviewed Journals and Conferences

Conference on Neural Information Processing Systems (NeurIPS), 2020-2025
International Conference on Machine Learning (ICML), 2022-2025
International Conference on Learning Representations (ICLR), 2021-2025
Computer Vision and Pattern Recognition (CVPR), 2019-2024
European Conference on Computer Vision (ECCV), 2020
International Conference on Computer Vision (ICCV), 2019, 2021, 2023, 2025
AAAI Conference on Artificial Intelligence (AAAI), 2019-2022
International Conference on Artificial Intelligence and Statistics (AISTATS), 2021

**STUDENTS
MENTORED**

- Kyle Sargent, Ph.D. in Computer Science, Stanford University, 2025 Summer
- Haian Jin, Ph.D. in Computer Science, Cornell University, 2025 Summer
- Rundi Wu, Ph.D. in Computer Science, Columbia University, 2024 Summer
- David Charatan, Ph.D. in EECS, MIT, 2024 Summer
- Stan Szymanowicz, Ph.D. in Engineering Science, University of Oxford, 2024 Summer
- Sirui Xie, Ph.D. in Statistics, UCLA, 2023 Winter - 2024 Summer
- Alex Trevithick, Ph.D. in CSE, UC San Diego, 2024 Spring - present
- Armand Comas-Massagué, Ph.D. in Computer Science, 2023 Summer - 2024 Summer
- Dehong Xu, Ph.D. in Statistics, UCLA, 2020 Fall - present
- Yaxuan Zhu, Ph.D. in Statistics, UCLA, 2018 Winter - present
- Yufan Ren, Master in Computer Science, University of Lausanne, 2019 Fall - 2020 Spring
- Xiaolin Fang, Ph.D. in Computer Science, MIT, 2018 Summer
- Junpei Zhou, Master in Language Technologies, CMU, 2017 Summer
- Wenwen Si, Master in Computer Vision, CMU, 2017 Summer
- Jiawen Wu, Master in Computer Science, USC, 2017 Summer

SKILLS

C/C++, Python, Jax/Flax, TensorFlow, PyTorch
Fluent in English and Chinese