

YIFAN JIANG

yifanjiang97@utexas.edu | Website: yifanjiang19.github.io | Google Scholar Citation: 7145

EDUCATION

The University of Texas at Austin, Austin, USA

2020 – 2024

Ph.D. in Electrical and Computer Engineering¹

- Advisor: [Zhangyang \(Atlas\) Wang](#)

Research Interests: Machine Learning, Generative Model

Huazhong University of Science and Technology, Wuhan, China

2015 – 2019

B.E. in Electronic Information Engineering

INDUSTRIAL EXPERIENCE

Meta, Menlo Park, USA

Jun. 2025 - Present

Senior Research Scientist at Meta SuperIntelligence Lab

Core contributor of the media generation model, delivered multiple key features to the [Meta AI/Vibes App](#).

- **Video Storytelling.** Advanced "Storytelling" feature in the Meta AI/Vibes from concept to the dogfooding stage.
 - Developed the training and inference infrastructure to support a **33B** text-to-videoaudio model. Implemented **FSDP / HSDP / CP** for robust training. Supported **LoRA**-based training strategy for efficiency.
 - Built data pipeline to collect and filter human-centric video data. Designed both high-level signals (VLM/LLM-based, face recognition, etc.) and low-level signals (aesthetic score, motion score, etc.).
 - Designed training recipes on both **Mid-training and SFT stages**, using human rating as signals for hyper-parameter tuning.
 - Constructed an LLM-based rewriter for prompt enhancement.
- **Person Insertion.** Led the development of the image editing feature "Add Me", successfully launching it in the Meta AI / Vibes app, beating Nano Banana Pro in both internal benchmarks and public A/B tests.
 - Architected a large-scale data filtering pipeline to support **Mid-training and SFT stages**, covering different distributions (synthetic and real).
 - Formulated training recipes for Mid-training and SFT stages, ablating different data mixture ratios. Orchestrated with a distillation team to deliver a 12-step diffusion model.
 - Constructed an evaluation pipeline utilizing human annotators.
- **Personalized Photo Generation.** Adopted reinforcement learning to improve the quality of generated contents.
 - Developed a **GRPO-based reinforcement learning** pipeline for diffusion model.
 - Collected data and trained a visual quality assessment model as the **reward model**, to evaluate content quality.

Apple, Seattle, USA

Mar. 2024 – Jun. 2025

Senior Research Scientist at AI/ML

Core contributor of the Apple 3D generative model, responsible for training and modeling.

- **3D Diffusion Model.** Architected a multi-view attention operator to improve 3D consistency. Optimized training speed utilizing PyTorch Profiler. Developed specialized noise scheduler for faster convergence.

Apple, Seattle, USA

May. 2023 – Jan. 2024

Research Intern at AI/ML Argos Team, worked with [Liangliang Cao](#)

- Designed an efficient diffusion model training algorithm for novel view synthesis applications.

Adobe, San Jose, USA

May. 2022 – Dec. 2022

Research Intern at [Marc Levoy's Team](#)

- Developed a high-quality depth estimator using multi-sensor assistance.

Google Research, Mountain View, USA

May. 2021 – May. 2022

Research Intern at [GCam](#), worked with [Tianfan Xue](#), [Bart Wronski](#), [Ben Mildenhall](#), [Peter Hedman](#), [Jon Barron](#).

- Developed a fast denoising operator and an efficient backbone.
- Designed a high-fidelity neural radiance field that improves the rendering quality.

¹Studied at Texas A&M University from Aug. 2019 to Aug. 2020; then transferred with my advisor to UT Austin

Adobe, San Jose, USA

May. 2020 – Nov. 2020

Research Intern at Applied Research Team (ART), worked with [He Zhang](#) and [Jianming Zhang](#).

- Developed a self-supervised method for image harmonization.

Bytedance AI Lab, Beijing, China

Jan. 2019 – Aug. 2019

Research Intern at Bytedance AI Lab, worked with [Xiaohui Shen](#), [Ding Liu](#), [Chen Fang](#), [Jianchao Yang](#).

- Designed a jointly image denoising and low-light enhancement algorithm.

SELECTED PUBLICATIONS

(* indicates equal contribution)

- [ICLR] **Y. Jiang**, H. Tang, J. Chang, D. Xu, L. Song, Z. Wang and L. Cao, “Efficient-3DiM: Learning a Generalizable Single-image Novel-view Synthesizer in One Day”, *International Conference on Learning Representations*, 2024
- [CVPR] D. Xu, **Y. Jiang**, P. Wang, Z. Fan, Y. Wang, and Z. Wang, “NeuralLift-360: Lifting An In-the-wild 2D Photo to A 3D Object with 360° Views”, *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. **Highlight (2.5% of 9155 submissions)**
- [CVPR] **Y. Jiang**, H. Peter, B. Mildenhall, D. Xu, J. T. Barron, and Z. Wang, “AligNeRF: High-Fidelity Neural Radiance Fields via Alignment-Aware Training”, *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023
- [ECCV] **Y. Jiang***, D. Xu*, P. Wang, Z. Fan, H. Shi, and Z. Wang. “SinNeRF: Training Neural Radiance Fields on Complex Scenes from a Single Image”, *European Conference on Computer Vision*, 2022.
- [ECCV] **Y. Jiang***, Z. Fan*, P. Wang*, X. Gong, D. Xu, and Z. Wang. “Unified Implicit Neural Stylization”, *European Conference on Computer Vision*, 2022.
- [ECCV] **Y. Jiang**, B. Wronski, B. Mildenhall, J. T. Barron, Z. Wang, and T. Xue. “Fast and High-Quality Image Denoising via Malleable Convolutions”, *European Conference on Computer Vision*, 2022.
- [NeurIPS] **Y. Jiang**, S. Chang, and Z. Wang, “TransGAN: Two Pure Transformers can Make One Strong GAN and That Can Scale Up”, *Advances in Neural Information Processing Systems*, 2021.
- [TIP] **Y. Jiang**, X. Gong, D. Liu, Y. Cheng, C. Fang, X. Shen, J. Yang, P. Zhou, and Z. Wang, “EnlightenGAN: Deep Light Enhancement without Paired Supervision”, *IEEE Transaction on Image Processing*. **IEEE SPS Young Author Best Paper Award.**

MEDIA HIGHLIGHT

- TransGAN was covered by [Quanta Magazine](#) (Mar. 2022) and was highlighted by [Top AI influencers](#) and high-profile [YouTubers](#), as well as considered as [the most influential new paper of the month \(Feb. 2021\)](#).
- AutoGAN was covered by [Synced AI Technology & Industry Review](#) (Aug. 2019), and also featured on [Towards Data Science](#) (Sep. 2019) and [Analytics Magazine](#) (Aug. 2019), etc.

COMMUNITY SERVICES

- Reviewer for: CVPR, ICCV, ECCV, ICML, NeurIPS, ICLR, Siggraph, Siggraph Aisa, IROS, WACV, IJCAI, TPAMI, TIP, IJCV, NeuroComputing, RA-L, TCI, TCSVT
- Workshop Organizer for: [ECCV RLQ-TOD Workshop 2020](#)

AWARDS

- **IEEE SPS Young Author Best Paper Award** 2024
- **University of Texas Graduate Dean’s Prestigious Fellowship** 2023
- **Apple Scholars in AI/ML Fellowship** 2023