Lu Jiang (Ph.D)

Last updated August 2025

Contact Information

Mountain View,
California 94043, USA

roadjiang@yahoo.com
http://www.lujiang.info

Research Interests

Generative-AI scientist with 14+ years of R&D (9 yrs industry). Principal Research Scientist @ Apple; former Research Director @ ByteDance/TikTok, Research Manager @ Google, and Adjunct Faculty @ CMU. Expertise in video generation and multimodal foundation models, with 15k+ citations (h-index 53).

Education

Carnegie Mellon University — 2011 - 2017
Ph.D. in Artificial Intelligence. (GPA: 4.12/4.33)

Free University of Brussels — 2010 - 2011
M.Sc. in Computer Science. (Erasmus Mundus Exchange Program)

Xi'an Jiaotong University
M.Sc. in Computer Science. 2010. (GPA: 3.62/4.00, Rank: 1/142)
B.Eng. in Software Engineering, 2008. (Major GPA: 3.88/4.00)

Honors & Awards

Best Paper at ICML — 2024
Best Paper at IJCA-JAIR — 2024
Best paper finalist at CVPR – 0.4% of all submitted papers (33/8161). — 2022
Best paper nomination at ACL – 1% of all submitted long papers (18/1737). — 2019
Yahoo! Fellowship. — 2016
Best performer on multimedia event detection at NIST TRECVID — 2013, 2014
Various scholarships: Fuji Xerox, Samsung, IBM, Erasmus Mundus Tandem — 2007 - 2010

Research Experience

Principal Research Scientist at *Apple* — July 2025 - Present
- Work on visual generation essential to Apple's products.

Research Director at *ByteDance/TikTok USA* — Feb 2024 - June 2025
- Lead for the R&D of the video generation foundation model: Seaweed, co-lead of SeeDance.
- Product impact: Ranked top-tier in video generation globally (July 2025). Deployed across XX products worldwide, driving YY% growth in revenue and engagement. Powers GenAI features on Douyin, CapCut, and TikTok with record-breaking popularity, and supports research such as Omnihuman.

Staff Research Scientist and Manager at *Google Research* — May 2017 - Jan 2024
- Staff Research Scientist and Manager (2021/11–2024/01)
  - Led image and video generation effort GVT. Launched Imagen/MUSE (Google's first T2I on Cloud, YouTube, Pixel), VideoPoet, and WALT; enabled automated video ads (+$XXM revenue).
  - Research highlight: ICML Best Paper VideoPoet and MAGVIT.
- Senior Research Scientist (2019/05–2021/10)
  - Led research on robust deep learning. Improved robustness across YouTube, Search Ads, Waymo, and Google Translate (ACL best paper candidate).
  - Research highlight: IJCA-JAIR Best Paper on Confident Learning (spawned Cleanlab, valued at $XXM); CVPR Best Paper Finalist Pyramid Adversarial Training.
- Research Scientist, Google Cloud AI (2017/05–2019/04)
  - First research hire under Dr. Fei-Fei Li and Dr. Jia Li. Launched in the AutoML models.
  - Research: MentorNet, TIRG (precursor research to Multisearch).

Adjunct Professor at *Carnegie Mellon University* — Oct 2021 - Sep 2024
- Adjunct Faculty Member (part-time), Carnegie Mellon University, School of Computer Science. Supervised graduate research, taught lectures, and served on Ph.D. thesis committees.

Research Assistant at *Carnegie Mellon University* — Sep. 2011 - Apr. 2017
- Key contributor to a 5-year IARPA project on zero-shot event detection in internet videos. Developed a system that ranked top in TRECVID 2013–2015 and scaled to 100M videos.

Intern Scientist at *Yahoo Research* — May 2016 - August 2016

- Worked on large-scale personal photo and video search on Flickr.

Intern at *Google Research* Feb. 2016 - May 2016

- Built the first TensorFlow training pipeline and contributed to the launch of YouTube-8M.

Research Intern at *Microsoft Research Asia* May 2010 - August 2010

Research Assistant at *Xi'an Jiaotong University* Sep. 2008 - July 2010

SELECTED PUBLICATIONS (H-INDEX=**53**) GOOGLE SCHOLAR

Generative AI: Video and Image

[1] Ceyuan Yang, Zhijie Lin, Yang Zhao, Shanchuan Lin, et al., Lu Jiang. Seaweed-7B: Cost-Effective Training of Video Generation Foundation Model. arXiv preprint 2504.08685 — Tech report 25.

[2] Yuwei Guo, Ceyuan Yang, Ziyan Yang, Zhibei Ma, Zhijie Lin, Zhenheng Yang, Dahua Lin, underlineLu Jiang. Long Context Tuning for Video Generation. *In ICCV*, 2025. — ICCV 25.

[3] Shanchuan Lin, Xin Xia, Yuxi Ren, Ceyuan Yang, Xuefeng Xiao, Lu Jiang. Diffusion Adversarial Post-Training for One-Step Video Generation. *In ICML*, 2025. — ICML 25.

[4] Jianyi Wang, Zhijie Lin, Meng Wei, Yang Zhao, Ceyuan Yang, Fei Xiao, Chen Change Loy, Lu Jiang. SeedVR: Seeding Infinity in Diffusion Transformer Towards Generic Video Restoration . *In CVPR*, 2025. — CVPR 25 (**Highlight**)

[5] Dan Kondratyuk, Lijun Yu, et al., Lu Jiang. VideoPoet: A Large Language Model for Zero-Shot Video Generation. *In ICML*, 2024. — ICML 24 (**Best paper**)

[6] Lijun Yu, Jose Lezama, Nitesh B. Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G. Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A. Ross, Lu Jiang. Language Model Beats Diffusion – Tokenizer is Key to Visual Generation. *In ICLR*, 2024. — ICLR 24

[7] Lijun Yu, Yong Cheng, Zhiruo Wang, Vivek Kumar, Wolfgang Macherey, Yanping Huang, David A Ross, Irfan Essa, Yonatan Bisk, Ming-Hsuan Yang, Kevin Murphy, Alexander G Hauptmann, Lu Jiang. SPAE: Semantic Pyramid AutoEncoder for Multimodal Generation with Frozen LLMs. *In NeurIPS*, 2023. — NeurIPS 23 (**Spotlight**)

[8] Kihyuk Sohn, Nataniel Ruiz, Kimin Lee, Daniel Castro Chin, Irina Blok, Huiwen Chang, Jarred Barber, Lu Jiang, Glenn Entis, Yuanzhen Li, Yuan Hao, Irfan Essa, et al. StyleDrop: Text-to-Image Generation in Any Style. *In NeurIPS*, 2023. — NeurIPS 23

[9] Lijun Yu, Yong Cheng, Kihyuk Sohn, Jose Lezama, Han Zhang, Huiwen Chang, Alexander G. Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, Lu Jiang. MAGVIT: Masked Generative Video Transformer. *In CVPR* 2023. — CVPR 23 (**Highlight**)

[10] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, William T. Freeman. MaskGIT: Masked Generative Image Transformer. *In CVPR*, 2022. — CVPR 22

[11] Kwonjoon Lee, Huiwen Chang, Lu Jiang, Han Zhang, Zhuowen Tu, Ce Liu. ViTGAN: Training GANs with vision transformers. *In ICLR*, 2022. — ICLR 22 (**Spotlight**)

Robust and Trustworthy Deep Learning

[13] Charles Herrmann, Kyle Sargent, Lu Jiang, Ramin Zabih, Huiwen Chang, Ce Liu, Dilip Krishnan, Deqing Sun. Pyramid Adversarial Training Improves ViT Performance. *In CVPR*, 2022 — CVPR 22 (**Best paper finalist**)

[14] Yong Chen, Lu Jiang, Wolfgang Macherey. Robust Neural Machine Translation with Doubly Adversarial Inputs. *In ACL*, 2019. — ACL 19 (**Best paper nomination**)

[15] Northcutt, Curtis, Lu Jiang, Isaac Chuang. Confident learning: Estimating uncertainty in dataset labels. *Journal of Artificial Intelligence Research 70 (2021): 1373-1411.* — JAIR 21 (**Best paper**)

[16] Youjiang Xu, Linchao Zhu, Lu Jiang, Yi Yang. Faster Meta Update Strategy for Noise-Robust Deep Learning. *In CVPR*, 2021. — CVPR 21 (**Oral**)

[17] Guoliang Kang, Lu Jiang, Yunchao Wei, Yi Yang, Alexander G. Hauptmann. Controllable and Progressive Image Extrapolation. *In IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2021 — TPAMI 20

[18] Lu Jiang, Di Huang, Mason Liu, Weilong Yang. Beyond Synthetic Noise: Deep Learning on Controlled Noisy Labels. *In ICML*, 2020. — ICML 20

[19] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, Li Fei-Fei. MentorNet: Learning Data-Driven Curriculum for Very Deep Neural Networks on Corrupted Labels. *In ICML*, 2018. — ICML 18

Video Understanding and Multimodal Learning

[20] Yunbo Wang, Lu Jiang, Ming-Hsuan Yang, Li-Jia Li, Mingsheng Long, Li Fei-Fei. Eidetic 3D LSTM: A Model for Video Prediction and Beyond. *In ICLR*, 2019. — ICLR 19

[21] Nam Vo, Lu Jiang, Chen Sun, Kevin Murphy, Li-Jia Li, Li Fei-Fei, James Hays. Composing Text and Image for Image Retrieval - An Empirical Odyssey. *In CVPR*, 2019. — CVPR 19 (**Oral**)

[22] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, Alexander Hauptmann. The Garden of Forking Paths: Towards Multi-Future Trajectory Prediction. *In CVPR*, 2020. — CVPR 20

[23] Lu Jiang, Shoou-I Yu, Deyu Meng, Teruko Mitamura, Alexander Hauptmann. Bridging the Ultimate Semantic Gap: A Semantic Search Engine for Internet Videos. *In ICMR*, 2015. — ICMR 15 (**Best paper candidate**)

[24] Yajie Miao, Lu Jiang, Hao Zhang, Florian Metze. Improvements to Speaker Adaptive Training of Deep Neural Networks. *In IEEE Spoken Language Technology (SLT)*, 2014. — SLT 14 (**Best poster**)

Curriculum Learning

[25] Lu Jiang, Deyu Meng, Shoou-I Yu, Zhen-Zhong Lan, Shiguang Shan, Alexander Hauptmann. Self-paced Learning with Diversity. *In NeurIPS*, 2014. — NeurIPS 14

[26] Qian Zhao, Deyu Meng, Lu Jiang, Qi Xie, Zongben Xu, Alexander Hauptmann. Self-paced Learning for Matrix Factorization. *In AAAI*, 2015. — AAAI 15 (**Oral**)

Professional Service

Area Chair:
- ICML (2024-2025), NeurIPS (2023-2025), ICLR (2025)
- CVPR (2023-2025), ICCV (2023,2025)
- ACM Multimedia (2019-2023), AAAI (SPC 2022-2024), WACV (2023-2024)
- Co-organized AI for Content Creation Workshop, CVPR (2020-2025)

Associate Editor:
- IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
- Transactions on Machine Learning Research (TMLR)
- Computer Vision and Image Understanding (CVIU)

Other Committees:
- America's Seed Fund (NSF SBIR) Panelist on AI (2019-Now)
- Google Faculty Award Committee (2018-2024)

Selected Lectures/Tutorials/Workshops:
- Invited Lecture, International School on Deep Learning (DeepLearn) 2025.
- Invited talk at the NTIRE workshop and the EDGE workshop at CVPR 2025.
- Keynote talk, 2024 IEEE International Conference on Multimedia and Expo (ICME 2024).
- Invited talk at the Weakly Supervised Learning workshop in ICLR 2021.
- Guest lectures on vision + language in Carnegie Mellon University (2019, 2021)

Former Interns/Students

Fortunate to work with the following talented students or interns.

Junwei Liang (CMU Ph.D. → Assistant Professor at HKUST Guangzhou) 2016-2020
Zelun Luo (Stanford Master → Stanford Ph.D.) 2018
Yunbo Wang (Tsinghua University Ph.D. → Professor Shanghai Jiaotong University) 2018
Alejandro Newell (University of Michigan Ph.D.) 2018
Nam Vo (Georgia Tech Ph.D. → Google Research) 2018
Jim Fan (Stanford Ph.D.) 2018
Yu Wu (Google Fellowship recipient. → Full Professor Wuhan University) 2017-2019
Yijun Li (UC Merced Ph.D. → Adobe Research) 2019
Hsin-Ying Lee (UC Merced Ph.D. →Snap Research) 2019-2020
Hung-Yu Tseng (University of California Merced Ph.D. → Meta Research) 2020
Chengzhi Mao (Columbia University Ph.D. → Assistant Professor at Rutegus University) 2021