

Lu Jiang (Ph.D)

Last updated January 2023

CONTACT INFORMATION

Google Research
1600 Amphitheatre Pkwy
Mountain View, CA, 94043

+1 (412) 897-5924
roadjiang@yahoo.com
<http://www.lujiang.info>

RESEARCH INTERESTS

My research goal is to solve real problems on big multimodal data. My research area is in the interdisciplinary field of Multimedia, Machine Learning, and Computer Vision, specifically including robust deep learning, content creation, and video understanding.

EDUCATION

Carnegie Mellon University 2011 - 2017
Ph.D. in Artificial Intelligence. (GPA: 4.12/4.33)
Advisors: Prof. Alexander Hauptmann and Prof. Teruko Mitamura
Thesis: Web-scale Multimedia Search for Internet Video Content

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), Advisor: Prof. Jun Liu
B.Eng. in Software Engineering, 2008. (Major GPA: 3.88/4.00)

HONORS & AWARDS

Best paper finalist at CVPR 2022
– 0.4% of all submitted papers (33/8161).

AI 2000 Most Influential Scholar 2021

Best paper nomination at ACL 2019
– 1% of all submitted long papers (18/1737).

Yahoo! Fellowship. 2016

Best Poster at IEEE SLT 2014

Best Paper nomination at ICMR 2014, 2015
– 2% of all submitted papers. Received two years in a row.

Best performer on multimedia event detection at NIST TRECVID 2013, 2014
–Key contributor to our winning system out of 30 companies and research institutes.

Best performer on surveillance event detection at NIST TRECVID 2011

Erasmus Mundus Tandem scholarship 2010
–Fellowship from the EU for studying in Europe. Forty-one master students in China were awarded.

Fuji Xerox Fellowship 2010

Samsung Fellowship 2009

IBM excellent student in China 2007

RESEARCH EXPERIENCE

Staff Research Scientist and Manager at *Google Research* May 2017 - Now.

- Staff Research Scientist and Manager @ Research 2021/11 - now. Research: content creation and multimodal foundation model. Product engagement and launches: Google Ads.
- Senior Research Scientist @ Research 2019/05 - 2021/10. Research: robust deep learning and multimodal learning. Product engagement and launches: YouTube, Search, Google Translation, Waymo self-driving car.
- Research Scientist @ Google Cloud AI 2017/05 - 2019/04. Research: robust learning and video understanding. Product engagement and launches: [AutoML](#) and vision models for Enterprise AI.

Adjunct Professor at *Carnegie Mellon University* Sep 2021 - Now

- [Adjunct faculty member](#) at Carnegie Mellon University, School of Computer Science, Language Technologies Institute, where Lu supervises graduate students' research, teaches lectures and serves on their Ph.D. thesis committee.

Research Assistant at *Carnegie Mellon University* Sep. 2011 - April 2017

- The key contributor to a five-year [IARPA project](#). The project is to approach an automatic method to detect the event in Internet videos without any user-generated metadata. Proposed the first-of-its-kind zero-shot system, which not only achieves the top performance in [NIST TRECVID evaluation 2013-2015](#), three years in a row, but also scales the search up to 100 million 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

- created the first training pipeline on [YouTube-8M](#) in TensorFlow, proposed novel a webly-labeled learning method that improves the state-of-the-art accuracy by Y%. In addition, slashed training time from a few days to a few hours.

Research Intern at *Microsoft Research Asia*

May 2010 - August 2010

- Designed and implemented a novel pattern/logic engine that was used in the launched product.

Research Assistant at *Xi'an Jiaotong University*

Sep. 2008 - July 2010

- Worked on a National High-Tech R&D Program aiming at discovering and managing educational resources on the Internet.

SELECTED
PUBLICATIONS
(H-INDEX=40)
[GOOGLE SCHOLAR](#)

- [1] Lijun Yu, Yong Cheng, Kihyuk Sohn, Jos \ddot{A} Lezama, Han Zhang, Huiwen Chang, Alexander G. Hauptmann, Ming-Hsuan Yang, Yuan Hao, Irfan Essa, [Lu Jiang](#). MAGVIT: Masked Generative Video Transformer. arXiv:2212.05199. Preprint 23
- [2] Kihyuk Sohn, Yuan Hao, Jos \ddot{A} Lezama, Luisa Polania, Huiwen Chang, Han Zhang, Irfan Essa, [Lu Jiang](#). Visual Prompt Tuning for Generative Transfer Learning. arXiv:2210.00990. Preprint 23
- [3] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, [Lu Jiang](#), Ming-Hsuan Yang, Kevin Murphy, William T. Freeman, Michael Rubinstein, Yuanzhen Li, Dilip Krishnan. Muse: Text-To-Image Generation via Masked Generative Transformers arXiv:2301.00704. Preprint 23
- [4] Jose Lezama, Tim Salimans, [Lu Jiang](#), Huiwen Chang, Jonathan Ho, Irfan Essa. Discrete Predictor-Corrector Diffusion Models for Image Synthesis. *In International Conference on Learning Representations (ICLR)*, 2023 ICLR 23
- [5] Xiang Kong, [Lu Jiang](#), Huiwen Chang, Han Zhang, Yuan Hao, Haifeng Gong, Irfan Essa. BLT: Bidirectional Layout Transformer for Controllable Layout Generation. *In European Conference on Computer Vision (ECCV)*, 2022 ECCV 22
- [6] Jose Lezama, Huiwen Chang, [Lu Jiang](#), Irfan Essa. Improved Masked Image Generation with Token-Critic. *In European Conference on Computer Vision (ECCV)*, 2022 ECCV 22
- [7] Charles Herrmann*, Kyle Sargent*, [Lu Jiang](#), Ramin Zabih, Huiwen Chang, Ce Liu, Dilip Krishnan, Deqing Sun. Pyramid Adversarial Training Improves ViT Performance. *In Computer Vision and Pattern Recognition (CVPR)*, 2022 CVPR 22 (**Best paper finalist**)
- [8] Huiwen Chang, Han Zhang, [Lu Jiang](#), Ce Liu, William T. Freeman. MaskGIT: Masked Generative Image Transformer. *In Computer Vision and Pattern Recognition (CVPR)*, 2022 CVPR 22
- [9] Chengzhi Mao, [Lu Jiang](#), Mostafa Dehghani, Carl Vondrick, Rahul Sukthankar, Irfan Essa. Discrete Representations Strengthen Vision Transformer Robustness. *In International Conference on Learning Representations (ICLR)*, 2022 ICLR 22
- [10] Kwonjoon Lee, Huiwen Chang, [Lu Jiang](#), Han Zhang, Zhuowen Tu, Ce Liu. ViT-GAN: Training gans with vision transformers. *In International Conference on Learning Representations (ICLR)*, 2022 ICLR 22 (**spotlight**)
- [11] Tianhao Zhang, Hung-Yu Tseng, [Lu Jiang](#), Weilong Yang, Honglak Lee, Irfan Essa. Text as Neural Operator: Image Manipulation by Text Instruction. *In ACM Multimedia (MM)*, 2021 MM 21
- [12] Yong Cheng, Wei Wang, [Lu Jiang](#), Wolfgang Macherey. Self-supervised and Supervised Joint Training for Resource-rich Machine Translation. *In International Conference on Machine Learning (ICML)*, 2021 ICML 21
- [13] Hung-Yu Tseng, [Lu Jiang](#), Ce Liu, Ming-Hsuan Yang, Weilong Yang. Regularizing Generative Adversarial Networks under Limited Data. *In Computer Vision and Pattern Recognition (CVPR)*, 2021 CVPR 21

- [14] Youjiang Xu, Linchao Zhu, Lu Jiang, Yi Yang. Faster Meta Update Strategy for Noise-Robust Deep Learning. In *Computer Vision and Pattern Recognition (CVPR)*, 2021 CVPR 21 (oral)
- [15] Yijun Li, Lu Jiang, Ming-Hsuan Yang. Contrastive Adaptation Network for Single- and Multi-Source Domain Adaptation. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2021 WACV 21
- [16] 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
- [17] Lu Jiang, Di Huang, Mason Liu, Weilong Yang. Beyond Synthetic Noise: Deep Learning on Controlled Noisy Labels. In *International Conference on Machine Learning (ICML)*, 2020 ICML 20
- [18] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, Alexander Hauptmann. The Garden of Forking Paths: Towards Multi-Future Trajectory Prediction. In *Computer Vision and Pattern Recognition (CVPR)*, 2020 CVPR 20
- [19] Yong Cheng, Lu Jiang, Wolfgang Macherey, Jacob Eisenstein. AdvAug: Robust Adversarial Augmentation for Neural Machine Translation. In *Association for Computational Linguistics (ACL)*, 2020 ACL 20
- [20] Hsin-Ying Lee, Lu Jiang, Irfan Essa, Phuong B Le, Haifeng Gong, Ming-Hsuan Yang, Weilong Yang. Neural Design Network: Graphic Layout Generation with Constraints. In *European Conference on Computer Vision (ECCV)*, 2020 ECCV 20 (spotlight)
- [21] Junwei Liang, Lu Jiang, Alexander Hauptmann. SimAug: Learning Robust Representations from 3D Simulation for Pedestrian Trajectory Prediction in Unseen Cameras. In *European Conference on Computer Vision (ECCV)*, 2020 ECCV 20
- [22] Hung-Yu Tseng*, Hsin-Ying Lee*, Lu Jiang, Weilong Yang, Ming-Hsuan Yang. RetrieveGAN: Image Synthesis via Differentiable Patch Retrieval. In *European Conference on Computer Vision (ECCV)*, 2020 (* equal contribution) ECCV 20
- [23] Yong Chen, Lu Jiang, Wolfgang Macherey. Robust Neural Machine Translation with Doubly Adversarial Inputs. In *Association for Computational Linguistics (ACL)*, 2019 ACL 19 (best paper nomination)
- [24] 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 *Computer Vision and Pattern Recognition (CVPR)*, 2019 CVPR 19 (oral)
- [25] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander Hauptmann, Li Fei-Fei. Peeking into the future: Predicting Future Person Activities and Locations in Videos. In *Computer Vision and Pattern Recognition (CVPR)*, 2019 CVPR 19
- [26] Guoliang Kang, Lu Jiang, Yi Yang, Alexander Hauptmann. Contrastive Adaptation Network for Unsupervised Domain Adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, 2019 CVPR 19
- [27] 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 *International Conference on Learning Representations (ICLR)*, 2019 ICLR 19
- [28] Junwei Liang, Lu Jiang, Liangliang Cao, Yannis Kalantidis, Li-Jia Li, Alexander Hauptmann. Focal Visual-Text Attention for Memex Question Answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019 TPAMI 19
- [29] 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 *International Conference on Machine Learning (ICML)*, 2018 ICML 18
- [30] Junwei Liang, Lu Jiang, Liangliang Cao, Li-Jia Li, Alexander Hauptmann. Focal Visual-Text Attention for Visual Question Answering. In *Computer Vision and Pattern Recognition (CVPR)*, 2018 CVPR 18 (spotlight)
- [31] Zelun Luo, Jun-Ting Hsieh, Lu Jiang, Juan Carlos Niebles, Li Fei-Fei. Graph Distillation for Action Detection with Privileged Information in RGB-D Videos. In *European Conference on Computer Vision (ECCV)*, 2018 ECCV 18

- [32] Lu Jiang, Yannis Kalantidis, Liangliang Cao, Sachin, Farfadi, Jiliang Tang, Alex Hauptmann. Delving Deep into Personal Photo and Video Search. *In Web Search and Data Mining (WSDM)*, 2017. WSDM 17
- [33] Junwei Liang, Lu Jiang, Deyu Meng, Alexander Hauptmann. Learning to Detect Concepts from Webly-Labeled Video Data. *In Joint Conference on Artificial Intelligence (IJCAI)*, 2016. IJCAI 16
- [34] Lu Jiang, Shoou-I Yu, Deyu Meng, Yi Yang, Teruko Mitamura, Alexander Hauptmann. Fast and Accurate Content-based Semantic Search in 100M Internet Videos. *In ACM Multimedia (MM)*, 2015. MM 15
- [35] Lu Jiang, Shoou-I Yu, Deyu Meng, Teruko Mitamura, Alexander Hauptmann. Bridging the Ultimate Semantic Gap: A Semantic Search Engine for Internet Videos. *In ACM International Conference on Multimedia Retrieval (ICMR)*, 2015. ICMR 15
(**Best paper candidate**)
- [36] Qian Zhao, Deyu Meng, Lu Jiang, Qi Xie, Zongben Xu, Alexander Hauptmann. Self-paced Learning for Matrix Factorization. *In Conference on Artificial Intelligence (AAAI)*, 2015. AAAI 15
(**oral**)
- [37] Lu Jiang, Deyu Meng, Shoou-I Yu, Zhen-Zhong Lan, Shiguang Shan, Alexander Hauptmann. Self-paced Learning with Diversity. *In Neural Information Processing Systems (NIPS)*, 2014. NIPS 14
- [38] Lu Jiang, Deyu Meng, Teruko Mitamura, Alexander Hauptmann. Easy Samples First: Self-paced Reranking for Zero-Example Multimedia Search. *In ACM Multimedia (MM)*, 2014. MM 14
- [39] 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**)
- [40] Lu Jiang, Wei Tong, Deyu Meng, Alexander Hauptmann. Towards Efficient Learning of Optimal Spatial Bag-of-Words Representations. *In ACM International Conference on Multimedia Retrieval (ICMR)*. 2014. ICMR 14
(**Best paper candidate**)
- [41] Lu Jiang, Yajie Miao, Yi Yang, Zhen-Zhong Lan, Alexander Hauptmann. Viral Video Style: A Closer Look at Viral Videos on YouTube. *In ACM International Conference on Multimedia Retrieval (ICMR)*, 2014. ICMR 14
- [42] Shoou-I Yu, Lu Jiang, et al. CMU-Informedia@TRECVID 2014. *In NIST TRECVID Video Retrieval Evaluation Workshop (TRECVID)*, 2014. TRECVID 14
(**Best system**)
- [43] Lu Jiang, Alexander Hauptmann, Guang Xiang. Leveraging High-level and Low-level Features for Multimedia Event Detection. *In ACM Multimedia (MM)*, 2012. MM 12
- [44] Jun Liu, Lu Jiang, Zhaohui Wu, Qinghua Zheng, Yanan Qian. Mining Learning-Dependency between Knowledge Units from Text. *The International Journal on Very Large Data Bases (VLDBJ)*, 20(3): 335-345, 2011. VLDBJ 11

PROFESSIONAL SERVICE

Area Chair:

- ICCV (2023) CVPR (2023), WACV (2023)
- AAAI SPC (2022, 2023)
- ACM Multimedia (2019, 2020, 2021, 2022)

Associate Editor:

- IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)

Other Committee:

- ACM Multimedia (2013-2022), CVPR (2018-2023), ICML (2020-2022), NeurIPS (2020-2021), ICCV (2019,2021), ECCV (2018,2020,2022), AAAI (2017, 2018, 2020-2023), IJCAI (2017), EMNLP (2020), MSML (2021)
- Top 10% reviewer in ICML 2021, NeurIPS 2021.
- NSF SBIR/STTR Review Panel on AI and CV (2019-Now)
- Google Faculty Award Committee, Research Award Committee
- Google Research Hiring Committee

Selected Lectures/Tutorials/Workshops:

- Co-organized AI for Content Creation Workshop, CVPR 2020-2023
- Invited talk on robust deep learning for the Weakly Supervised Learning workshop in ICLR 2021
- Tutorial on Visual Search and Question Answering, ICME 2019
- Guest lectures on vision + language in Carnegie Mellon University (LTI-11-775 2019)

Journal or Book Reviewer:

- Journal of Supercomputing, Journal on Signal Processing, Journal on Information Fusion, IEEE Transactions on Multimedia (TMM), Springer Computing (Book Chapter review), Journal of Intelligent Information Systems, Journal of Machine Learning Research (JMLR), Computer Vision and Image Understanding (CVIU), Electronic Commerce Research and Applications, International Journal of Automation and Computing, Journal of Artificial Intelligence Research (JAIR), IEEE Transactions on Neural Networks and Learning Systems, IEEE Transactions on Circuits and System for Video Technology, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), ACM Transactions on Multimedia Computing, Communications and Applications

FORMER

INTERNS/STUDENTS

Fortunate to work with the following talented the following students or interns.

Junwei Liang (CMU Ph.D. $\xrightarrow{\text{now}}$ Assistant Professor at HKUST Guangzhou) 2016-2020

Zelun Luo (Stanford Master $\xrightarrow{\text{next}}$ Stanford Ph.D.) 2018

Yunbo Wang (Tsinghua University Ph.D. $\xrightarrow{\text{next}}$ Professor Shanghai Jiaotong University) 2018

Alejandro Newell (University of Michigan Ph.D.) 2018

Nam Vo (Georgia Tech Ph.D. $\xrightarrow{\text{next}}$ Google Research) 2018

Jim Fan (Stanford Ph.D.) 2018

Mason Liu (Cornell Ph.D.) 2019

Yu Wu (Google Fellowship recipient. $\xrightarrow{\text{now}}$ Full Professor Wuhan University) 2017-2019

Yijun Li (UC Merced Ph.D. $\xrightarrow{\text{next}}$ Adobe Research) 2019

Yunseok Jang (University of Michigan Ph.D.) 2019

Hsin-Ying Lee (UC Merced Ph.D. $\xrightarrow{\text{next}}$ Snap Research) 2019-2020

Hung-Yu Tseng (University of California Merced Ph.D. $\xrightarrow{\text{next}}$ Meta Research) 2020

Chengzhi Mao (Columbia University Ph.D.) 2021

Xiang Kong (Carnegie Mellon University Ph.D. $\xrightarrow{\text{next}}$ Apple Research) 2021