

Lu Jiang (Ph.D)

Last updated July 2018

CONTACT INFORMATION

Cloud AI, Google AI
1155 Borregas Avenue
Sunnyvale, CA, 94089

+1 (412) 897-5924
lujiang@google.com
<http://www.cs.cmu.edu/~lujiang>

RESEARCH INTERESTS

My research goal is to solve real problems on big data. My research area is in the interdisciplinary field of Multimedia, Machine Learning, Computer Vision, Information Retrieval and Big Data, which specifically, includes video understanding and search, weakly supervised learning on noisy data, vision+language, cloud, etc.

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 2008 - 2011
M.Sc. in Computer Science. (GPA: 3.62/4.00, Rank: 1/142), Advisor: Prof. Jun Liu
B.Eng. in Software Engineering, 2004-2008. (Major GPA: 3.88/4.00)

HONORS & AWARDS

Yahoo! Fellowship. 2016

NIPS, AAAI, SIGMM, SIGIR, SIGKDD, SIGWEB, NSF student travel grants.

Best Poster at IEEE Spoken Language Technology 2014

Best Paper candidate at International Conference on Multimedia Retrieval 2014, 2015

–Only 2% of all submitted papers were nominated for best paper candidates. Received it twice.

Best performer on multimedia event detection at NIST TRECVID 2013, 2014

–Key contributor to our winning system. Participants included 30 companies and research institutes.

Best performer on surveillance event detection at NIST TRECVID 2011

Erasmus Mundus Tandem scholarship 2010

–Fellowship from the Europe Union for studying in Europe. Only 41 master students in China were awarded.

Fuji Xerox Fellowship 2010

Samsung Fellowship 2009

IBM excellent student in China 2007

RESEARCH EXPERIENCE

Research Scientist at *Google AI Cloud AI* May 2017 - Now

- Democratize cloud machine learning and video intelligence.

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](#), 3 years in a row, but also scales the search up to 100 million videos. The developed techniques lead to patterns and inventions.

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

- The project is on the large-scale personal photo and video search on Flickr. Proposed deep query understanding models that boost the state-of-the-art (word2vec) search accuracy by a relative 45%. In addition, analyzed big personal media search logs on Flickr, and discovered distinguishing characteristics of this novel problem.

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

- The project is on training concept detectors on big weakly-labeled data of YouTube using Tensorflow. Proposed novel webly-labeled learning method which 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 for retrieval and data mining system. The code was used in the launched product.

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

- The project is a National High-Tech R&D Program, which aims at discovering and managing the educational resources on the Internet. Designed and implemented domain term recognition and title extraction for World, PPT, PDF, and HTML. Proposed algorithms to recognize the associations in knowledge networks; designed and implemented the knowledge element association detector.

SELECTED
PUBLICATIONS
(H-INDEX=21)
GOOGLE SCHOLAR

- [1] 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
- [2] 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)
- [3] 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
- [4] Yu Wu, Linchao Zhu, Lu Jiang, Yi Yang. Decoupled Novel Object Captioner. *In ACM Multimedia (MM)*, 2018 MM 18
- [5] 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
- [6] Junwei Liang, Lu Jiang, Deyu Meng, Alexander Hauptmann. Leveraging Multimodal Prior Knowledge for Large-scale Concept Learning in Noisy Web Data. *In ACM International Conference on Multimedia Retrieval (ICMR)*, 2017. ICMR 17 (oral)
- [7] Junwei Liang, Lu Jiang, Alexander Hauptmann. Temporal Localization of Audio Events for Conflict Monitoring in Social Media. *In IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017. ICASSP 17
- [8] Zhang, Dingwen, Junwei Han, Lu Jiang, Senmao Ye, and Xiaojun Chang. Revealing event saliency in unconstrained video collection. *IEEE Transactions on Image Processing* 26, no. 4 (2017): 1746-1758. TIP 17
- [9] Lu Jiang. Web-scale Multimedia Search for Internet Video Content. *In International Conference on World Wide Web (WWW)*, 2016. WWW 16
- [10] 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
- [11] 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
- [12] 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)
- [13] 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)
- [14] Dingwen Zhang, Deyu Meng, Li Chao, Lu Jiang, Zhao Qian, Junwei Han. A self-paced multiple-instance learning framework for co-saliency detection. *In IEEE International Conference on Computer Vision (ICCV)*, 2015. ICCV 15
- [15] 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

- [16] 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
- [17] 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)
- [17] 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)
- [18] 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
- [19] Shou-I Yu, Lu Jiang, et al. CMU-Infomedia@TRECVID 2014. *In NIST TRECVID Video Retrieval Evaluation Workshop (TRECVID)*, 2014. TRECVID 14 (best system)
- [20] Lu Jiang, Alexander Hauptmann, Guang Xiang. Leveraging High-level and Low-level Features for Multimedia Event Detection. *In ACM Multimedia (MM)*, 2012. MM 12
- [21] 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

PATTERN AND INVENTIONS

- [1] *Event Labeling through Analytic Media Processing*, Disclosure of Invention, Key Inventor. No. 0247601-15-0048, 2016.
- [2] *Large-scale video content retrieval through text query*, Invention, Second Inventor, U.S. Serial No. 62/285,256 2015.
- [3] *A Features Dictionary Generating Method for Text Classification based on LZW Compression Algorithm*. Authorization No. ZL 2008 1 0232557.2 2008 (in Chinese).
- [4] *A Deep Web Adaptive Crawling Method based on Minimum Executable Pattern*. Application No. 200810232555.3, 2008. (in Chinese).

PROFESSIONAL SERVICE

Technical Program Committee Member:
 ACM Multimedia 2013-2018 (Author Advocator Co-chair 2017)
 CVPR 2018
 AAAI 2017-2018
 IJCAI 2017
 IA-Summit 2016
 Pacific Rim Conference on Multimedia 2014
 Co-chair of CMU LTI Student Research Symposium 2015

Journal and Book Review:

Journal of Supercomputing
 Journal on Signal Processing
 Journal on Information Fusion
 IEEE Transactions on Multimedia (TMM)
 Journal of Machine Learning Research (JMLR)
 Computer Vision and Image Understanding (CVIU)
 Springer Computing (Book Chapter review)
 Journal of Intelligent Information Systems
 Electronic Commerce Research and Applications
 International Journal of Automation and Computing
 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 OR STUDENTS

Junwei Liang (CMU Master $\xrightarrow{\text{next}}$ CMU Ph.D.)
 Zelun Luo (Stanford Master $\xrightarrow{\text{next}}$ Stanford Ph.D.)
 Nam Vo (Georgia Tech Ph.D.)

Yunbo Wang (Tsinghua University Ph.D. $\xrightarrow{\text{next}}$ MIT Visiting Student)
Alejandro Newell (University of Michigan Ph.D.)
Serena Yeung (Stanford Ph.D. $\xrightarrow{\text{next}}$ Stanford Assistant Professor)
Justin Johnson (Stanford Ph.D.)

TALKS

Web-scale Multimedia Search for Internet Video Content.
Yahoo Research, San Francisco, CA, June 2016.
University of Southern California ISI, Los Angeles, CA, Jan. 2016.
Columbia University, New York City, NY, Dec. 2015.
Vision And Learning SEminar (VALSE), Beijing, China, Feb. 2014.

Learning to Detect Concepts from Webly-Labeled Video Data.
Google Research, Mountain View, CA, March 2016.

CMU@TRECVID: Large-scale Semantic Indexing.
University of Central Florida, Orlando, FL, Nov. 2014

CMU@TRECVID: Multimedia Event Detection.
National Institute of Standards and Technology, Gaithersburg, MD, Nov. 2013.

CMU@TRECVID: Surveillance Event Detection.
National Institute of Standards and Technology, Gaithersburg, MD, Nov. 2011.

Knowledge Graph Analysis.
Free University of Brussels, Brussels, Belgium, Oct. 2010.