

Junwei Liang

Name in Chinese: 梁俊卫

Language Technologies Institute

Carnegie Mellon University

5000 Forbes Avenue, Pittsburgh, PA 15213

Phone: (412) 251-****

Email: junwei@cs.cmu.edu

Linkedin: [linkedin.com/in/junweiliang/](https://www.linkedin.com/in/junweiliang/)

Github: github.com/JunweiLiang

Home page: www.cs.cmu.edu/~junwei

Education

Ph.D. Candidate in Artificial Intelligence

2017 - present

School of Computer Science, Carnegie Mellon University

Advisor: Alexander Hauptmann

Thesis Proposal: Joint Analysis and Prediction of Human Actions and Paths in Video

M.S. in Artificial Intelligence

2015 - 2017

School of Computer Science, Carnegie Mellon University

Advisor: Alexander Hauptmann

B.S. in Computer Science

2011 - 2015

School of Information, Renmin University of China

Advisor: Qin Jin

Awards

Rising Star (云帆奖-明日之星) , World Artificial Intelligence Conference	2020
Baidu Scholarship (10 recipients globally)	2019
Best Demo Award at CBMI2019	2019
Yahoo! Fellowship	2016-2018
Winner, Automated Streams Analysis for Public Safety Challenge (\$30k prize)	2020
Winner, TRECVID ActEV Challenge	2019
Winner, TRECVID Ad-hoc Video Search Challenge, no annotation track	2016
CMU LTI Student Research Symposium Best Paper Honorable Mentions	2018
Google Cloud COVID-19 Research Grant (\$6200)	2020
Google Cloud University Research Grant (\$22000)	2020
Carnegie Mellon Graduate Research Fellowship	2015-2021
Nominee, Carnegie Mellon University's, Google Ph.D. Fellowship	2020
Nominee, Carnegie Mellon University's, Microsoft Ph.D. Fellowship	2020
IJCAI, NIST TRECVID, NIST PSCR, CVPR, ICMR student travel grants	2015-2019
Best Undergraduate Thesis (Top 5%)	2015
Second Prize, the National Undergraduates Computer Design Competition of China	2014
National Prize (Top 10%), National Undergraduates Innovation Project	2013

Selected Media

- **Carnegie Mellon University News.** Amateur Drone Videos Could Aid in Natural Disaster Damage Assessment, August 28, 2020. [\[link\]](#)

- **AZO Robotics.** *New AI System Helps Detect Damage Caused to Buildings by Hurricanes*, August 31, 2020.
- **Washington Post.** *Lewd cheerleader videos, sexist rules: Ex-employees decry Washington's NFL team workplace* (Featured in the video analytics), August 26, 2020. [\[link\]](#)
- **CBS.** *Researchers At Carnegie Mellon University Develop Video System To Locate Mass Shooters Using Smartphones*, November 20, 2019.
- **post-gazette.** *CMU develops video system that can locate mass shooter*, November 20, 2019.
- **GIZMODO.** *Smartphone Videos Can Now Be Analyzed and Used to Pinpoint the Location of a Shooter*, November 21, 2019.
- **DailyMail.** *Active shooters can be located within minutes by new software that analyzes smartphone video from the scene and can even identify the type of gun*, November 20, 2019.
- **Techspot.** *Researchers develop system that can pinpoint a shooter's location using smartphone videos*, November 21, 2019.
- **New York Times.** *Who Killed the Kiev Protesters? A 3-D Model Holds the Clues* (Featured in the video analytics), May 30, 2018.
- **读芯术.** 卡内基梅隆大学梁俊卫：视频中行人的多种未来轨迹预测, August, 2020.
- **Baidu.** 乘风破浪的AI技术青年——首届WAIC云帆奖名单公布, July 11, 2020.
- **China.com.cn.** 人大高瓴人工智能学院“高屋建瓴-青年说”首期开讲, Jan 6, 2020.
- **Baidu.** AI界的中国力量！百度奖学金助力中国AI人才绽放光芒！, Jan 5, 2020.
- **量子位.** 李飞飞团队造出“窥视未来”新AI:去哪干啥一起猜, 准确率压倒老前辈, received 30k+ views in a week, Feb 13, 2019.
- **机器之心.** 遇见未来！李飞飞等提出端到端系统Next预测未来路径与活动, Feb 14, 2019.

Projects

- **Thesis Proposal: Joint Analysis and Prediction of Human Actions and Paths in Video**
 - <https://www.cs.cmu.edu/~junwei/proposal/>
- **Future Person Prediction:** <https://next.cs.cmu.edu/>
- **3D Simulation:** <https://next.cs.cmu.edu/multiverse/>
- **COVID-19 Project - Social Distancing Early Forecasting:**
 - <https://github.com/JunweiLiang/social-distancing-prediction>
- **MemexQA:** <https://memexqa.cs.cmu.edu/>
- **Video Event Reconstruction and Shooter Localization:** <https://vera.cs.cmu.edu/>
- **3D Reconstruction:** https://vera.cs.cmu.edu/VERA_3D_Reconstruction/
- **Video Analytic Toolkit:** <https://aladdin1.inf.cs.cmu.edu/human-rights>
- **CMU LTI's Attendance Management System:**
 - https://github.com/JunweiLiang/Lecture_Attendance_Management

Selected Publications

[H-INDEX=9, Google Scholar: <https://scholar.google.com/citations?hl=en&user=bMedjfUAAAAJ>]

[H-INDEX=10, Semantic Scholar: <https://www.semanticscholar.org/author/Junwei-Liang/1915796>]

1. **Junwei Liang,** Lu Jiang, Kevin Murphy, Ting Yu and Alexander Hauptmann. "The Garden of Forking Paths: **Towards Multi-Future Trajectory Prediction**". In Proceedings of the IEEE

- Conference on Computer Vision and Pattern Recognition (CVPR), 2020. [**paper mentioned in Waymo's internal document** as a proper evaluation direction and [ECCV'20 workshop keynotes](#)]
2. **Junwei Liang**, Lu Jiang and Alexander Hauptmann. "**SimAug**: Learning Robust Representations from **3D Simulation** for Pedestrian Trajectory Prediction in Unseen Cameras". In Proceedings of the European conference on computer vision (ECCV), 2020.
 3. **Junwei Liang**, Lu Jiang, Liangliang Cao, Yannis Kalantidis, Li-Jia Li and Alexander Hauptmann. "Focal Visual-Text Attention for **Memex Question Answering**". In IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2019.
 4. **Junwei Liang**, Lu Jiang, Juan Carlos Niebles, Alexander Hauptmann and Li Fei-Fei. "Peeking into the Future: **Predicting Future Person** Activities and Locations in Videos". In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. [**200+ stars on Github, #1 Tensorflow-based code on PaperWithCode in Trajectory Prediction task, top-cited paper at CVPR 2019 on pedestrian trajectory prediction**]
 5. **Junwei Liang**, Jay Aronson and Alexander Hauptmann. "**Shooter Localization** Using Social Media Videos". In Proceedings of the 27th **ACM International Conference on Multimedia**, 2019. [**Widely reported by news outlets like CBS Pittsburgh**]
 6. **Junwei Liang**, Lu Jiang, Liangliang Cao, Li-Jia Li and Alexander Hauptmann. "Focal Visual-Text Attention for **Visual Question Answering**". In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018. [**Spotlight Paper**]
 7. **Junwei Liang**, Lu Jiang, Deyu Meng and Alexander Hauptmann. "Leveraging Multi-modal Prior Knowledge for Large-scale Concept Learning in Noisy Web Data". In proceeding of the ACM International Conference on Multimedia Retrieval (ICMR), 2017.
 8. **Junwei Liang**, Desai Fan, Han Lu, Poyao Huang, Jia Chen, Lu Jiang and Alexander Hauptmann. "An Event Reconstruction Tool for Conflict Monitoring Using Social Media." In Proceedings of the Thirty-First **AAAI** Conference on Artificial Intelligence, 2017. [Demo Paper]
 9. **Junwei Liang**, Lu Jiang and Alexander Hauptmann, "Temporal Localization of **Audio Events** for Conflict Monitoring in Social Media." In Proceedings of the 42nd IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP 2017.
 10. **Junwei Liang**, Lu Jiang and Alexander Hauptmann. "Webly-supervised learning of multimodal video detectors." In Proceedings of the Thirty-First **AAAI** Conference on Artificial Intelligence, 2017. [Demo Paper]
 11. **Junwei Liang**, Lu Jiang, Deyu Meng and Alexander Hauptmann. "Learning to **detect concepts** from webly-labeled video data." In **IJCAI**, 2016. [**Oral Paper**]

Academic Service

Conference Reviewer

- AAAI 2021
- WACV 2021
- ACM Multimedia 2017/2019/2020
- ACL 2020 Student Research Workshop
- CVPR 2020 AI for Content Creation Workshop

Journal Reviewer

- IEEE Transactions on Image Processing (TIP)
- Pattern Recognition
- IEEE Transactions on Intelligent Transportation Systems (ITS)
- IEEE Transactions on Multimedia
- IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)
- ACM TOMM
- Neurocomputing
- IEEE Access
- Defense Technology

Teaching Experience & Research Talks

Thesis Proposal: Joint Analysis and Prediction of Human Actions and Paths in Video [\[Link\]](#) 2020

Successfully proposed my CS Ph.D. thesis proposal.

Waymo Machine Learning Reading Group: Pedestrian Future Trajectory Prediction 2020

Presented my research work internally with 60 Waymo research scientists and engineers

CMU LTI Summer Seminar: Pedestrian Future Trajectory Prediction [\[Link\]](#) 2020

Presented my research work at CMU's summer research seminar for graduate students

Research Presentation: Multi-Future Trajectory Prediction [\[Talk\]](#) 2020

Live presentation for hundreds of Chinese college students.

Contributed Talk: Shooter Localization Using Social Media Videos [\[Talk\]](#) 2020

AI for Social Good, Carnegie Mellon University

Guest Lecture: Introduction to Machine Learning in Computer Vision [\[Slides\]](#) 2020

11-775 Large-scale Multimedia Analysis, hosted by Prof. Alex Hauptmann and Prof. Rita Singh, Carnegie Mellon University

Invited Talk: Real-time Video Analytics for Situation Awareness [\[Talk-2018\]](#) [\[Talk-2019\]](#) 2019

At the Public Safety Communications Research annual conference hosted by NIST PSCR.

Oral Talk: Focal Visual-Text Attention for Memex Question Answering 2018

At CMU LTI Student Research Symposium 2018. Received best paper honorable mentions award.

Spotlight Talk: Focal Visual-Text Attention for Visual Question Answering [\[Video\]](#) 2018

At CVPR 2018 in Salt Lake City.

Oral Talk: Learning to Detect Concepts from Webly-Labeled Video Data **2016**

At IJCAI 2016 in New York City.

Teaching Assistant **2018**

LTI Colloquium, hosted by Prof. Alex Hauptmann, Carnegie Mellon University

Teaching Assistant **2014**

Practical Course on Speech Synthesis and Building Synthetic Voices by Prof. Alan Black, Summer School of Renmin University of China

Research & Project Experience

Google Cloud AI **Summer 2020**

Research Intern

- **Viewpoint-equivariant representation learning for activity recognition**
 - Conducted research in activity recognition in videos
 - Advised by Ting Yu, Xuehan Xiong, and Prof. Liangliang Cao

Geospatial Alpha **July 2020 - Present**

Advisor

- **Provide advice on location prediction research and technologies**
- **Updates on recent breakthroughs in research fields related to location prediction or other geospatial areas**
 - <https://www.geospatialalpha.com/people>

Perception Team, Google AI **Summer 2019**

Research Intern

- **Future person activity and trajectory prediction in videos [CVPR 2020]**
 - Integrated research models to a Google Cloud product
 - Used 3D simulator (carla.org) to collect multi-modal future behavioral data
 - Advised by Lu Jiang and Prof. Kevin Murphy

Google Cloud AI **May 2018 - Dec 2018**

Student Researcher

- **Future person activity and trajectory prediction in videos [CVPR 2019]**
 - Paper translated and reported by multiple Chinese media (量子位 & 机器之心) with 30k+ views in a week
 - Github received 200+ stars
 - Advised by Lu Jiang and Prof. Juan Carlos Nieves

Informedia Lab, Carnegie Mellon University**2015 - present**

Research Assistant, Advisor: Alexander Hauptmann

- **Deep Intermodal Video Analytics (DIVA) by IARPA, Sep 2017 – Present [1]**
 - The DIVA project's focus is on developing systems for activity detection in extended videos. Our CMU team's (INF & MUDSML) system achieved the best performance [2] on the activity detection challenge in surveillance videos hosted by NIST & IARPA. I have developed and released the code and model for Object Detection & Tracking [3].
 - [1] <https://www.iarpa.gov/index.php/research-programs/diva>
 - [2] https://actev.nist.gov/prizechallenge#tab_leaderboard
 - [3] https://github.com/JunweiLiang/Object_Detection_Tracking
- **Public Safety Communications Research (PSCR) by NIST, Sep 2017 – Present [1]**
 - I have led our CMU team to develop crucial tools for public safety that will help first-responders, analysts to quickly understand and react to public safety events. These tools include 3D event reconstruction [2], shooter localization from social media videos and the Video Event Reconstruction and Analysis (VERA) system [3]. I have presented these tools at the annual PSCR meetings [4]. I have used VERA to process videos for human rights researchers and reporters from the **New York Times** and **Washington Post**.
 - [1] <https://www.nist.gov/ctl/pscr/real-time-video-analytics-situation-awareness>
 - [2] https://vera.cs.cmu.edu/VERA_3D_Reconstruction/
 - [3] <https://vera.cs.cmu.edu/>
 - [4] <https://www.nist.gov/publications/2018-stakeholder-meeting-presentation-day-2-room-30b>
- **Project InMind by Yahoo & CMU, 2016 – 2018 [1]**
 - The InMind Project aims to build a smart personal assistant on cellphones. I am the major contributor to the smart assistant's question-answering module. I have collected the first multi-modal question answering dataset from a real-world personal collection like Flickr [2]. Based on the dataset, I have developed a deep neural network that is state-of-the-art for real-world multi-modal media question answering. I have also built a demo app on Android phones that could answer questions about multi-media content in personal emails.
 - [1] <https://www.cmu.edu/homepage/computing/2014/winter/project-inmind.shtml>
 - [2] <https://memexqa.cs.cmu.edu/>

Multimedia Computing Lab, Renmin University of China**2013 - 2015**

Research Assistant, Advisor: Qin Jin

- Semantic Content Analysis from User Generated Videos

- Participated HUAWEI semantic concept annotation of UGC videos grand challenge 2014 and ranked 3rd in the evaluation
- Natural Language Description Generation for Images and Videos with Deep Models
 - Rank 1st in ImageCLEF 2015 “image to sentence” subtask in the evaluation
 -

Full Stack Web Developer, Renmin University of China

2012 - 2014

- Single-handedly developed the CMS website for China Undergraduate Computer Design Competition
 - Tens of thousands participates from all over the country logged in to the website and submit their work, and the website was also a platform for hundreds of reviewers to evaluate the participates’ work

Advising & Mentoring

- **Seokeon Choi**, PhD student at Korea Advanced Institute of Science and Technology (KAIST), 2020
- **Xiaoyu Zhu**, master student at CMU, 2019 - 2021
- **Diganta Misra**, undergraduate student at Kalinga Institute of Industrial Technology, 2020
- **Helen Dong**, undergraduate student at CMU, 2020
She is from the AI mentoring program at CMU.
- **Soham Ghosh**, master student at CMU, 2019
He is the lead author on a NAACL’19 paper
- **Ryota Hinami**, Ph.D. student at University of Tokyo and visiting scholar at CMU, 2018
He is the lead author of a machine learning paper co-authored with me
- **Vaibhav**, master student at CMU, 2017 - 2019
- **Xinru Yang**, master student at CMU, 2017 - 2019
She is now at Google.

Last updated: Nov. 2020