

JOONSEOK LEE, PH. D.

1600 Amphitheatre Parkway, Mountain View, CA 94043, USA

joonseok2010@gmail.com — <http://www.joonseok.net/>

RESEARCH INTEREST

Video understanding, Recommendation Systems, Collaborative Filtering, Climate Applications

EDUCATION

Georgia Institute of Technology, Atlanta, Georgia, USA *Aug 2015*

Ph.D. in Computer Science

Statistical Machine Learning and Visualization Lab. (Advisor: Guy Lebanon, Hongyuan Zha)

Thesis: Local Approaches for Collaborative Filtering

GPA: 4.00 / 4.00

Seoul National University, Seoul, Korea *Aug 2009*

B. S. in Computer Science & Engineering

Summa Cum Laude (Department Rank: 1/63)

GPA: 4.20 / 4.30 (Overall), 4.22 / 4.30 (Major)

Yale University, New Haven, Connecticut, USA *Jul 2009 – Aug 2009*

Exchange Program

Uppsala University, Uppsala, Sweden *Aug 2005 – Jan 2006*

Exchange Program

EMPLOYMENT

Google Research, Mountain View, California, USA *Jun 2015 – Present*

Research Engineer, Foresight Group, Research & Machine Intelligence

Select projects: Video representation learning, Multimodal deep domain understanding

Amazon, Seattle, Washington, USA *May 2014 – Aug 2014*

Machine Learning Scientist Intern (Supervisor: Guy Lebanon, Ted Sandler)

Select projects: Cold-start recommendation using latent space mapping for Amazon Instant Video

Microsoft Research, Mountain View, California, USA *Jan 2014 – May 2014*

Research Scientist Intern (Supervisor: Ariel Fuxman, Bo Zhao, Yuanhua Lv)

Select projects: Leveraging Wikipedia semantics for contextual insight search

Google Research, Mountain View, California, USA *May 2013 – Aug 2013*

Software Engineering Intern (Supervisor: Yoram Singer, Samy Bengio)

Select projects: Local low-rank matrix approximation, Local collaborative ranking

NHN (Naver) Corp., Seongnam, Korea *Apr 2007 – Jun 2010*

Software Engineer, Assistant Manager

Select projects: Game server log analysis and management system, Abuser monitoring system, Online broadcasting system

Technovision Inc., Seoul, Korea *Mar 2006 – Apr 2007*

Software Engineer

Select projects: Real-time motion tracking with PTZ camera, Unified sewer pipe monitoring system, Video editor for network camera

PUBLICATIONS

Book Chapter

1. **Joonseok Lee**. Recommendation Systems: An Industrial Application of Network Big Data for Computational Intelligence, *Big Data and Computational Intelligence in Networking*, Yulei Wu, Fei Hu, Geyong Min, Albert Y. Zomaya (ed.), CRC Press, ISBN: 9781498784863, Chapter 20, 2017.

Referred Journals

1. Sangho Suh, Sungbok Shin, **Joonseok Lee**, Chandan K. Reddy, Jaegul Choo. Localized User-Driven Topic Discovery via Boosted Ensemble of Nonnegative Matrix Factorization, *Knowledge and Information Systems (KAIS)*, 2018.
2. **Joonseok Lee**, Seungyeon Kim, Guy Lebanon, Yoram Singer, Samy Bengio. LLORMA: Local Low-Rank Matrix Approximation, *Journal of Machine Learning Research (JMLR)*, 2016.
3. **Joonseok Lee**, Hanggjun Cho, Robert Ian (Bob) McKay. A Rapid Screening and Testing Protocol for Keyboard Layout Speed Comparison, *IEEE Transactions on Human-Machine Systems*, 2014.
4. **Joonseok Lee**, Mingxuan Sun, Guy Lebanon. PREA: Personalized Recommendation Algorithms Toolkit, *Journal of Machine Learning Research (JMLR)* 13:2699-2703, 2012.

Referred Conferences

1. Minjin Choi, Jinhong Kim, **Joonseok Lee**, Hyunjung Shim, Jongwuk Lee. Session-aware Linear Item-Item Models for Session-based Recommendation, *Proceedings of the 31st ACM Web Conference (WWW)*, 2021.
2. Minjin Choi, Yoonki Jeong, **Joonseok Lee**, Jongwuk Lee. Local Collaborative Autoencoders, *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM)*, 2021.
3. Sunghyun Park, Kangyeol Kim, Junsoo Lee, Jaegul Choo, **Joonseok Lee**, Sookyung Kim, Edward Choi. Vid-ODE: Continuous-Time Video Generation with Neural Ordinary Differential Equation, *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
4. Hyodong Lee, **Joonseok Lee**, Joe Ng, Paul Natsev. Large Scale Video Representation Learning via Relational Graph Clustering, *Proceedings of the 36th IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
5. Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, **Joonseok Lee**, Emily Denton. Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing, *Proceedings of the 3rd AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIIES)*, 2020.
6. Sookyung Kim, Sunghyun Park, Sunghyo Chung, **Joonseok Lee**, Yunsung Lee, Hyojin Kim, Mr Prabhat, Jaegul Choo. Learning to Focus and Track Extreme Climate Events, *Proceedings of the 30th British Machine Vision Conference (BMVC)*, 2019.
7. Sami Abu-El-Haija, Amol Kapoor, Bryan Perozzi, **Joonseok Lee**. N-GCN: Multi-scale Graph Convolution for Semi-supervised Node Classification, *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019.
8. Seong Jae Hwang, **Joonseok Lee**, Balakrishnan Varadarajan, Ariel Gordon, Zheng Xu, Paul Natsev. Annotating Videos at YouTube Scale, *Proceedings of the 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2019.
9. Sookyung Kim, Hyojin Kim, **Joonseok Lee**, Sangwoong Yoon, Samira Ebrahimi Kahou, Karthik Kashinath, Mr. Prabhat. Deep-Hurricane-Tracker: Tracking and Forecasting Extreme Climate Events, *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019.

10. Sookyung Kim, Sangwoong Yoon, **Joonseok Lee**, Samira Ebrahimi Kahou, Hyojin Kim, Karthik Kashinath, Mr. Prabhat. Deep-Hurricane-Tracker: Tracking Extreme Climate Events, *Climate Informatics*, 2018.
11. **Joonseok Lee**, Sami Abu-El-Haija, Balakrishnan Varadarajan, Paul Natsev. Collaborative Deep Metric Learning for Video Understanding, *Proceedings of the 24th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2018.
12. **Joonseok Lee**, Sami Abu-El-Haija. Large-Scale Content-Only Video Recommendation, *CEFRL Workshop at the International Conference on Computer Vision (ICCV)*, 2017.
13. Sangho Suh, Jaegul Choo, **Joonseok Lee**, Chandan Reddy. Local Topic Discovery via Boosted Ensemble of Nonnegative Matrix Factorization, *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI), Sister Conferences track*, 2017.
14. Sangho Suh, Jaegul Choo, **Joonseok Lee**, Chandan Reddy. Boosted L-EnsNMF: Local Topic Discovery via Ensemble of Nonnegative Matrix Factorization, *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2016. **Best Student Paper Award**
15. **Joonseok Lee**, Ariel Fuxman, Bo Zhao, Yuanhua Lv. Leveraging Knowledge Bases for Contextual Entity Exploration, *Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2015.
16. **Joonseok Lee**, Kisung Lee, Jennifer G. Kim, Sookyung Kim. Personalized Academic Research Paper Recommendation System, *Proceedings of the 6th International Workshop on Social Recommender Systems (SRS)*, 2015.
17. Seungyeon Kim, **Joonseok Lee**, Guy Lebanon, Haesun Park. Local Context Sparse Coding, *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015.
18. Seungyeon Kim, **Joonseok Lee**, Guy Lebanon, Haesun Park. Estimating Temporal Dynamics of Human Emotions, *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, 2015.
19. **Joonseok Lee**, Samy Bengio, Seungyeon Kim, Guy Lebanon, Yoram Singer. Local Collaborative Ranking, *Proceedings of the 23rd International World Wide Web Conference (WWW)*, 2014. **Best Student Paper Award**
20. **Joonseok Lee**, Seungyeon Kim, Guy Lebanon, Yoram Singer. Local Low-Rank Matrix Approximation, *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.
21. **Joonseok Lee**, Seungyeon Kim, Guy Lebanon, Yoram Singer. Matrix Approximation under Local Low-Rank Assumption, *The Learning Workshop in International Conference on Learning Representations (ICLR)*, 2013.
22. Mingxuan Sun, Fuxin Li, **Joonseok Lee**, Ke Zhou, Guy Lebanon, Hongyuan Zha. Learning Multiple-Question Decision Trees for Cold-Start Recommendation, *Proceedings of the 6th ACM International Conference on Web Search and Data Mining (WSDM)*, 2013.
23. **Joonseok Lee**, Mingxuan Sun, Seungyeon Kim, Guy Lebanon. Automatic Feature Induction for Stagewise Collaborative Filtering, *Advances in Neural Information Processing Systems (NIPS) 25*, 2012.
24. **Joonseok Lee**, Hanggjun Cho, Robert Ian (Bob) McKay. Rapid Screening of Keyboard Layouts, *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2012.
25. **Joonseok Lee**, Robert Ian (Bob) McKay. Optimizing a Personalized Cellphone Keypad, *Proceedings of the 5th International Conference on Convergence and Hybrid Information Technology (ICHIT)*, 2011.

Workshops, Demos, and ArXiv Preprints

1. Bowen Zhang*, Hexiang Hu*, **Joonseok Lee***, Ming Zhao, Sheide Chammas, Vihan Jain, Eugene Ie, Fei Sha. A Hierarchical Multi-Modal Encoder for Moment Localization in Video Corpus,

ArXiv:2011.09046, 2020. (*Equal contribution)

2. Sunghyun Park, Kangyeol Kim, Sookyung Kim, **Joonseok Lee**, Junsoo Le, Jiwoo Lee, Jaegul Choo. Hurricane Nowcasting with Irregular Time-step using Neural-ODE and Video Prediction, *Proceedings of the International Conference on Learning Representations (ICLR), Workshop on Tackling Climate Change with ML*, 2020.
3. Sookyung Kim, Sunghyun Park, Sunghyo Chung, Yunsung Lee, Hyojin Kim, **Joonseok Lee**, Jaegul Choo, Mr Prabhat. Focus and Track: Pixel-wise Spatio-temporal Hurricane Tracking, *Proceedings of the International Conference on Machine Learning (ICML), Workshop on Climate Change AI*, 2019.
4. **Joonseok Lee**, Paul Natsev, Walter Reade, Rahul Sukthankar, George Toderici. The 2nd YouTube-8M Large-Scale Video Understanding Challenge, *Proceedings of the 15th European Conference on Computer Vision (ECCV)*, 2018.
5. Sami Abu-El-Haija, Amol Kapoor, Bryan Perozzi, **Joonseok Lee**. N-GCN: Multi-scale Graph Convolution for Semi-supervised Node Classification, *Proceedings of the 14th International Workshop on Mining and Learning with Graphs (MLG)*, 2018.
6. **Joonseok Lee**, Nisarg Kothari, Paul Natsev. Content-based Related Video Recommendations, *Advances in Neural Information Processing Systems (NIPS)*, Demonstration Track, 2016.
7. Sami Abu-El-Haija, Nisarg Kothari, **Joonseok Lee**, Paul Natsev, George Toderici, Balakrishnan Varadarajan, Sudheendra Vijayanarasimhan. YouTube-8M: A Large-Scale Video Classification Benchmark, *ArXiv Report arXiv:1609.08675*, 2016.
8. **Joonseok Lee**, Mingxuan Sun, Guy Lebanon. A Comparative Study of Collaborative Filtering Algorithms, *ArXiv Report arXiv:1205.3193*, 2012.
9. Kiyeon Lee, **Joonseok Lee**, Hyunwoong Shin, Sunghwan Kim, Jungwoo Choi. Real-Time Motion Tracking System using Omni-Directional PTZ Camera, *Small and Medium Business Administration of South Korea*, 2006.

Theses

1. **Joonseok Lee**. Local Approaches for Collaborative Filtering, *Georgia Institute of Technology*, Ph.D. Dissertation, 2015.
2. **Joonseok Lee**. Merging Algorithm for Unified Communicator, *Seoul National University*, Bachelor's Thesis, 2009.

INVITED TALKS AND PRESENTATIONS

- **Kakao Brain**, Hierarchical Multi-Modal Encoder for Moment Localization in Corpus *Dec 2020*
- **Kakao Enterprise**, Large-Scale Video Understanding *Oct 2020*
- **University of Wisconsin, Madison**, Large-Scale Video Understanding *Oct 2020*
- **Ulsan National Institute of Science and Technology**, Two-Stage Approach for Grounding Text Queries with Moments in Large-Scale Video Corpus *Oct 2020*
- **Yonsei University**, Large-Scale Video Understanding *Apr 2020*
- **Naver**, Collaborative Deep Metric Learning for Video Understanding *Aug 2018*
- **Kakao Mobility**, Collaborative Deep Metric Learning for Video Understanding *Aug 2018*
- **Louisiana State University**, Recommendation Systems *Feb 2017*
- **Smart Cloud Show**, Research at Google: With a Case of YouTube-8M *Sep 2017*
- **Global HR Forum**, My Challenge: Settling down in Silicon Valley *Nov 2016*
- **Korea University**, Recommendation Systems *Nov 2015*

- **NVidia**, Local Approaches for Collaborative Filtering *Apr 2015*
- **Carnegie Mellon University**, Local Approaches for Collaborative Filtering *Apr 2015*
- **NEC Lab. America**, Local Approaches for Collaborative Filtering *Mar 2015*
- **Microsoft Research**, Local Collaborative Ranking for Recommendation Systems *Mar 2015*
- **Yahoo! Lab**, Local Collaborative Ranking for Recommendation Systems *Feb 2015*
- **Korea Advanced Institute of Science and Technology (KAIST)**, Local Approaches for Recommendation Systems *Dec 2014*
- **Center for Signal and Information Processing (CSIP)**, Georgia Institute of Technology, Local Collaborative Ranking *Oct 2014*
- **Amazon**, Cold-start Recommendation using Latent Space Mapping *Aug 2014*
- **Microsoft Research**, Leveraging Wikipedia semantics for Contextual Exploration *May 2014*
- **Naver Labs**, Local Collaborative Ranking *Apr 2014*
- **Korea Advanced Institute of Science and Technology (KAIST)**, Collaborative Filtering Approaches for Recommendation Systems *Nov 2013*
- **Seoul National University**, Collaborative Filtering Approaches for Recommendation Systems *Nov 2013*
- **Google Research**, Local Low-Rank Matrix Factorization *Aug 2013*
- **Home Depot**, Personalized Recommender Systems: Mining User Preference *Jun 2013*
- **KOCSEA 2012**, Automatic Feature Induction for Stagewise Collaborative Filtering *Dec 2012*
Best Lightning Talk Award

HONORS AND AWARDS

- Best Student Paper Award** *Dec 2016*
IEEE International Conference on Data Mining (ICDM)
- American Delegation in Heidelberg Laureate Forum 2014** *Sep 2014*
Oak Ridge Associated Universities (ORAU)
- 2014 Faces of Inclusive Excellence**, Georgia Institute of Technology *Sep 2014*
- KSEA-KUSCO Graduate Students Scholarship** *Aug 2014*
Korean-American Scientists and Engineers Association (KSEA)
- Best Poster Presentation Award**, US-Korea Conference 2014 *Aug 2014*
- Best Student Paper Award** *Apr 2014*
The 23rd ACM International World Wide Web Conference (WWW)
- Best Lightning Talk Award**, The 13th KOCSEA Technical Symposium *Dec 2012*
- Study Abroad Scholarship**, Korea Foundation for Advanced Studies *Aug 2010*

PROFESSIONAL SERVICES

Conference/Workshop Organizer

- *Program Chair*, The 3rd YouTube-8M Large-Scale Video Understanding Challenge Workshop, International Conference on Computer Vision (ICCV) 2019
- *Local Chair*, KOCSEA Symposium 2018
- *Program Chair*, The 2nd YouTube-8M Large-Scale Video Understanding Challenge Workshop, European Conference on Computer Vision (ECCV) 2018

- *Program Chair*, YouTube-8M Large-Scale Video Understanding Challenge Workshop, Conference on Computer Vision and Pattern Recognition (CVPR) 2017
- *Publicity Chair*, International Conference on Artificial Intelligence and Statistics (AISTATS) 2015
- *Program Chair Assistant*, ACM International Conference on Information and Knowledge Management (CIKM) 2012

Conference Program Committee

- Advances in Neural Information Processing Systems (NIPS) 2019, 2018, 2015, 2014
- International Conference on Machine Learning (ICML) 2019
- International Conference on Learning Representations (ICLR) 2021, 2020, 2019
- AAAI Conference on Artificial Intelligence (AAAI) 2021, 2020, 2015, 2014
- IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR) 2021, 2020
- IEEE International Conference on Computer Vision (ICCV) 2021, 2019
- European Conference on Computer Vision (ECCV) 2020
- IEEE Winter Conference on Applications of Computer Vision (WACV) 2021
- ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2021
- ACM International Conference on Web Search and Data Mining (WSDM) 2021, 2017, 2016, 2015
- ACM International World Wide Web Conferences (WWW) 2018
- IEEE International Conference on Data Mining (ICDM) 2019

Conference External Reviewer

- AAAI Conference on Artificial Intelligence (AAAI) 2017, 2016
- ACM International Conference on Information and Knowledge Management (CIKM) 2012

Journal Reviewer

- Journal of Machine Learning Research (JMLR)
- ACM Transactions on Intelligent Systems and Technology (TIST)
- IEEE Transactions on Knowledge and Data Engineering (TKDE)
- Information Sciences (INS)
- IEEE Intelligent Systems (IS)
- Hindawi Computational Intelligence and Neuroscience (CIN)
- Journal of Network and Computer Applications (JNCA)
- IEEE Transactions on Big Data

TEACHING EXPERIENCE

Teaching Assistant, Georgia Institute of Technology

- CS 7641 Machine Learning *Fall 2014*
- CSE/ISYE 6740 Computational Data Analysis *Fall 2014*
- CS 7641 Machine Learning *Fall 2013*
- CSE/ISYE 6740 Computational Data Analysis *Fall 2013*

Guest Lecturer, Georgia Institute of Technology

- CSE 6240 Web Search and Text Mining *Spring 2015*
- CS 7641 Machine Learning *Fall 2014*

Instructor, NHN Corp.

- Programming Psychology *Jul 2008 – Aug 2008*

SELECT PROJECTS

- **Moment Localization in Video Corpus**, *Google Research* 2020
 To localize relevant segments for a text query in untrimmed and unsegmented videos at scale, we propose Hierarchical Multi-Modal EncoderR (HAMMER) that encodes a video at both coarse-grained and finegrained level to extract information at different scales based on multiple subtasks, including video retrieval, segment temporal localization, and masked language modeling.
- **Annotation for Private Videos**, *Google Research* 2018-2019
 Working on annotating private videos on Google Photos to enable and improve keyword search and creating collages. Annotating private videos is more challenging than public ones, since we are not able to watch them, and usually no labels available. Tackling with re-weighted sampling from public videos and semi-supervised learning.
- **YouTube-8M Dataset**, *Google Research* 2016-2019
 A large-scale labeled video dataset containing 8 million YouTube video IDs and associated labels from a diverse vocabulary of 4,800 visual entities. It comes with precomputed state-of-the-art vision features. Hosting annual video classification competition on Kaggle, since 2017.
- **Content-based Video Recommendation**, *Google Research* 2015-2019
 Applying deep neural networks on learning co-watch relationship between a pair of YouTube videos to produce features optimized for recommendation purpose. Demonstration on recent videos shows surprisingly nice performance to retrieve related videos.
- **Cold-start Recommendation using Latent Space Mapping**, *Amazon* 2014
 Pure collaborative filtering recommenders cannot be used for cold-start users and items. To tackle this problem, we tried to learn a mapping from user or item feature space to the latent space discovered by matrix factorization. We observed significant improvement in Amazon Instant Video recommendation with user features such as book or DVD purchases.
- **Contextual Exploration with Wikipedia semantics**, *Microsoft Research* 2014
 Contextual exploration is an entity recommendation problem given a query and context, in order to satisfy user's information need directly within an application. We leverage semantic signals from Wikipedia link structures as well as relate to the context with several graph mining techniques. A crowd-sourced experimental study indicates that the proposed method successfully mines contextually-relevant pages.
- **Local Low-Rank Matrix Approximation, Ranking**, *Google Research* 2013
 A prevalent assumption in constructing matrix approximations was that the partially observed matrix is of low-rank. Instead, we assume that the matrix is locally of low-rank, leading to a representation of the observed matrix as a weighted sum of low-rank matrices. It is easy to parallelize, making it a viable approach for large scale real-world recommender systems.
- **Learning Multiple-Question Decision Trees for Cold-Start Recommendation** 2012
 For cold-start recommendation, it is important to rapidly prole new users and generate a good initial set of recommendations through an interview process. We proposed an algorithm learning to conduct the interview process guided by a decision tree with multiple questions at each split.
- **Automatic Feature Induction for Stagewise Collaborative Filtering** 2012
 For the task of predicting missing ratings in collaborative filtering, we observe that different models have relative advantages in different regions of the input space. This motivated our approach of using stagewise linear combinations of collaborative filtering algorithms, with non-constant combination coefficients based on kernel smoothing.

Last updated: January 15, 2021.