

# YaGuang Li

---

## CONTACT INFORMATION

Google DeepMind  
1600 Amphitheatre Parkway  
Mountain View, CA

*E-mail:* [yaguang@google.com](mailto:yaguang@google.com)  
*Homepage:* <http://yaguang.li>

## RESEARCH INTERESTS

Machine learning; Instruction Tuning for Large Models; Deep learning on graphs; Spatiotemporal Forecasting;

## EDUCATION

**University of Southern California**, Los Angeles, California, USA  
*Department of Computer Science*

Ph.D. in Computer Science, Aug. 2014 - May 2019, GPA: 3.95/4.0

- Advisors: Cyrus Shahabi, Yan Liu

**University of Chinese Academy of Sciences**, Beijing, China  
*Institute of Software*

M.S in Computer Science, Sep. 2011 - Jul. 2014, Ranked Top 1%

**University of Science and Technology Beijing**, China  
*Department of Computer Science*

B.Eng in Computer Science, Sep. 2007 - Jul. 2011, GPA: 3.87/4.0 (Major), Ranked Top 1%

## PREPRINTS

Gemini Team, Google

Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2312.11805*, 2023

Gemini Team, Google

Gemini: A Family of Highly Capable Multimodal Models. *arXiv preprint arXiv:2312.11805*, 2023

Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A Choquette-Choo, Aakanksha Chowdhery, Clment Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Daz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, **YaGuang Li**, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Aurko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, Yonghui Wu.

PaLM 2 Technical Report. *arXiv preprint arXiv:2305.10403*, 2023

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, **YaGuang Li**, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, Quoc Le.  
LaMDA: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022

REFEREED  
CONFERENCE  
PUBLICATIONS

Sidharth Mudgal\*, Jong Lee\*, Harish Ganapathy, **YaGuang Li**, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, Jilin Chen, Alex Beutel, Ahmad Beirami. Controlled decoding from language models. *International Conference on Machine Learning*, (ICML), 2024

Yun He, Steven Zheng, Yi Tay, Jai Gupta, Yu Du, Vamsi Aribandi, Zhe Zhao, **YaGuang Li**, Zhao Chen, Donald Metzler, Heng-Tze Cheng, Ed H Chi. Hyperprompt: Prompt-based task-conditioning of transformers. *International Conference on Machine Learning*, (ICML), 2022

Hongzhi Shi, Quanming Yao, Qi Guo, **YaGuang Li**, Lingyu Zhang, Jieping Ye, Yong Li, Yan Liu. Predicting Origin-Destination Flow via Multi-Perspective Graph Convolutional Network. *IEEE International Conference on Data Engineering*, (ICDE), 2020 (Research Track, Short Paper)

**YaGuang Li\***, Xu Geng\*, Leye Wang, Lingyu Zhang, Qiang Yang, Yan Liu, Jieping Ye (\*Equal Contribution). Spatiotemporal Multi-graph Convolution Network for Ride-hailing Demand Forecasting. *The Thirty-Third AAAI Conference on Artificial Intelligence* (AAAI), 2019 (**Oral presentation**)

Hanpeng Liu, **YaGuang Li**, Michael Tsang, Yan Liu. CoSTCo: A Neural Tensor Completion Model for Sparse Tensors, *ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (KDD), 2019 (Research Track, Oral Presentation)

Zahaib Akhtar, **YaGuang Li**, Ramesh Govindan, Emir Halepovic, Shuai Hao, Yan Liu, Subhabrata Sen. AViC: A Cache for Adaptive Bitrate Video. *The 15th International Conference on emerging Networking EXperiments and Technologies* (CoNEXT), 2019

Mingxuan Yue, **YaGuang Li**, Haoze Yang, Ritesh Ahuja, Yao-yi Chiang, and Cyrus Shahabi. DETECT: Deep Trajectory Clustering for Mobility-Behavior Analysis. *IEEE International Conference on Big Data* (IEEE Bigdata), 2019

**YaGuang Li**, Kun Fu, Zheng Wang, Cyrus Shahabi, Jieping Ye, Yan Liu. Multi-task Representation Learning for Travel Time Estimation, *ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (KDD), 2018 (Research Track)

**YaGuang Li**, Rose Yu, Cyrus Shahabi, Yan Liu. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting, *International Conference on Learning Representations* (ICLR), 2018

Yijun Lin, Nikhit Mago, Yu Gao, **YaGuang Li**, Yao-Yi Chiang, Cyrus Shahabi, Jos Luis Ambite. Exploiting Spatiotemporal Patterns for Accurate Air Quality Forecasting using Deep Learning. *International Conference on Advances in Geographic Information Systems* (ACM SIGSPATIAL GIS), 2018

**YaGuang Li\***, Rose Yu\*, Ugur Demiryurek, Cyrus Shahabi, Yan Liu (\*Equal Contribution). Deep Learning: A Generic Approach for Extreme Condition Traffic Forecasting. *Proceedings of the Seventeenth SIAM International Conference on Data Mining (SDM)*, 2017 **Best Research Paper Nomination**

**YaGuang Li**, Han Su, Ugur Demiryurek, Bolong Zheng, Tieke He, Cyrus Shahabi. PaRE: A System for Personalized Route Guidance. *Proceedings of the 26th International Conference on World Wide Web (WWW)*, 2017

**YaGuang Li**, Han Su, Ugur Demiryurek, Bolong Zheng, Kai Zeng, Cyrus Shahabi. PerNav: A Route Summarization Framework for Personalized Navigation. *ACM International Conference on Management of Data (demonstration) (SIGMOD)*, 2016

Mohammad Asghari, Dingxiong Deng, Cyrus Shahabi, Ugur Demiryurek, **YaGuang Li**. Price-aware Real-time Ridesharing at Scale - An Auction-based Approach. *International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL GIS)*, 2016

**YaGuang Li**, Dingxiong Deng, Ugur Demiryurek, Cyrus Shahabi, Siva Ravada. Towards Fast and Accurate Solutions to Vehicle Routing in a Large-Scale and Dynamic Environment. *14th International Symposium on Spatial and Temporal Databases (SSTD)*, 2015

**YaGuang Li**, Chengfei Liu, Kuien Liu, Jiajie Xu, Fengcheng He, Zhiming Ding. On Efficient Map-matching According to Intersections You Pass By. *International Conference on Database and Expert Systems Applications (DEXA)*, 2013

Kuien Liu, Bin Yang, Shuo Shang, **YaGuang Li**, Zhiming Ding. MOIR/uots: Trip recommendation with user oriented trajectory search. *International Conference on Mobile Data Management (demonstration) (MDM)*, 2013

Kuien Liu, **YaGuang Li**, Fengcheng He, Jiajie Xu, Zhiming Ding. Effective Map-matching on the Most Simplified Road Network. *International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL GIS) 2012, Contest Paper*

JOURNAL  
PUBLICATIONS

**YaGuang Li**, Cyrus Shahabi. A Brief Overview of Machine Learning Methods for Short-term Traffic Forecasting and Future Directions. *ACM SIGSPATIAL Special*, 2018

Zhiming Ding, Bin Yang, Ralf Hartmut Güting, **YaGuang Li**. Network-Matched Trajectory-Based Moving-Object Database: Models and Applications. *IEEE Transactions on Intelligent Transportation Systems (TITS)*, 2015

PATENT

Rajat Sen, Shuxin NIE, **YaGuang LI**, Abhimanyu Das, Nicolas LOEFF, Ananda Theertha Suresh, Pranjali AWASTHI, Biswajit PARIJA. Regression and Time Series Forecasting Patent, WO2022251857A1, Issued 2022

Shuai Hao, Subhabrata Sen, Emir Halepovic, Zahaib Akhtar, Ramesh Govindan, **YaGuang Li**. Adaptive Bitrate Video Cache. US Patent, US20210185368A1, Issued 2021

Haoming Guo, Zhiming Ding, Kuien Liu, Jiajie Xu, Huaiye Xu, **YaGuang Li**, Tianwei Zhang. Parallel Data Processing Method Based on Distributed Structure. CN Patent, CN103412897B, Issued 2017

WORKSHOP  
PUBLICATIONS

**YaGuang Li**, Chuizheng Meng, Cyrus Shahabi, Yan Liu. Structure-informed Graph Auto-encoder for Relational Inference and Simulation, *International Conference on Machine Learning (ICML)*

*Workshops on Learning and Reasoning with Graph-Structured Representations*, 2019

**YaGuang Li**, Rose Yu, Cyrus Shahabi, Yan Liu. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. *The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS) Time Series Workshop*, 2017, **Oral Presentation**

Rose Yu, **YaGuang Li**, Cyrus Shahabi, Ugur Demiryurek, Yan Liu. Extreme Traffic Forecasting: A Deep Learning Approach. Poster accepted to *ACM SIGKDD Conference on Knowledge Discovery and Data Mining(KDD)* workshop on Mining and Learning from Time Series, 2016

WORK  
EXPERIENCE

**Google DeepMind**, Mountain View, CA, US

*Senior Staff Research Engineer*

Aug. 2019 - Present

- Co-led the finetuning effort of Gemini 1.5 and Gemini 1.0 for Gemini Advanced.
- Lead the serving and efficiency work of Bard/LaMDA.
- Automated machine learning (AutoML) for time-series forecasting.

**DiDi Chuxing AI Labs**, Beijing, China

*Visiting Student*

Dec. 2017 - Feb. 2018, May 2018 - Aug. 2018

- Proposed a multi-task representation learning framework for origin-destination travel time estimation. Experimental results on two large-scale real-world datasets showed that the proposed approach clearly outperformed state-of-the-art methods.
- Work published on the International Conference on Knowledge Discovery and Data Mining (KDD) 2018, and the AAI Conference on Artificial Intelligence (AAAI) 2019.

**Facebook Inc.**, Menlo Park, CA, US

*Software Engineer Intern*

May 2017 - Aug. 2017

- Proposed a unified machine learning model for both place visit segmentation and classification based on Hidden Markov Model and Gradient Boosted Decision Tree. Comparing to the method used in production, the proposed model improved the precision by more than 10% at the same recall. The proposed model has been deployed in the production.

**Google Inc.**, New York, NY, US

*Software Engineer Intern*

May 2016 - Aug. 2016

- Proposed a model to generate high-quality business information (that will be displayed on Google Maps) from various low-quality data sources, e.g., street view image, untrusted third-party websites. The proposed model had generated tens of thousands of high-quality business information.
- Proposed a deep learning model to predict whether a website is the authority website of a small business. This work was presented at a Google-wide machine learning workshop.

**Google Inc.**, Montreal, Quebec, Canada

*Research Intern*

May 2015 - Aug. 2015

- Proposed a video summarization algorithm to select interesting frames based on audio, visual and motion information. Results on more than 500 minutes youtube videos showed that the algorithm consistently outperforms baseline methods, e.g., k-means and uniform sampling based ones, in terms of attractiveness and informativeness.
- Proposed an approach to generate summary video for preview in poor Internet conditions. The resulted videos have bitrate at about 40kbit/s while preserving reasonable good image quality.

**Chinese Academy of Sciences**, Beijing, China

*Research Assistant*

Aug. 2012 - Jul. 2014

- Contributed to many projects funded by National Natural Science Foundation of China (NSFC) on spatial-temporal data management.

- Proposed traffic analysis model based on network matched GPS trajectories that achieved higher accuracy with less storage and communication cost. The result was published in a top-tier journal (IEEE Transactions on Intelligent Transportation Systems).

TEACHING

**Machine Learning for Time Series Analysis** - Statistical Models and Deep Learning  
Summer 2018 in Peking University with Prof. Yan Liu

**Geospatial Information Management (CSCI 587)** Fall 2015, Fall 2016  
Teaching assistant: Advised 35 course projects; delivered several lectures; designed exams and organized weekly Q/A discussion.

ACADEMIC SERVICE **Workflow Chair**, International Conference on Learning Representations (ICLR), 2022

**Web Chair**, ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2020

**Co-organizer**, 7th SIGKDD Workshop on Mining and Learning from Time Series (MiLeTS), 2021

**Co-organizer**, 6th SIGKDD Workshop on Mining and Learning from Time Series (MiLeTS), 2020

**Co-organizer**, 5th SIGKDD Workshop on Mining and Learning from Time Series (MiLeTS), 2019

**Co-organizer**, 4th SIGKDD Workshop on Mining and Learning from Time Series (MiLeTS), 2018

**Web Chair**, International Conference on Web Search and Data Mining (WSDM), 2018

Reviewer, Conference on Neural Information Processing Systems (NeurIPS), 2020, 2021

Reviewer, International Conference on Machine Learning (ICML), 2019 - 2021

Reviewer, Transportation Research Part C (TRC), 2019-2021

Reviewer, IEEE Transactions on Intelligent Transportation Systems (TITS), 2015 - 2021

Reviewer, IEEE Transactions on Intelligent Systems and Technology (TIST), 2019 - 2021

Reviewer, International Joint Conference on Artificial Intelligence (IJCAI), 2021

Reviewer, International Conference on Learning Representations (ICLR), 2021

Reviewer, IEEE Transactions on Knowledge and Data Engineering (TKDE), 2019-2021

Reviewer, GeoInformatica, 2019

Reviewer, Journal of Computer Science and Technology (JCST), 2019

External Reviewer, International Conference on Very Large Databases (VLDB), 2016, 2017

External Reviewer, International Conference on Distributed Computing Systems (ICDCS), 2016

External Reviewer, ACM International Conference on Management of Data (SIGMOD), 2015

SELECTED AWARDS AND HONORS **Annenberg Graduate Fellowship**, University of Southern California, 2014-2018

**Best Research Assistant**, University of Southern California, 2018

**KDD Student Travel Award**, 2018

**ICLR Travel Award**, 2018

**NIPS TSW Travel Award**, 2017

**Outstanding Graduate**, University of Chinese Academy of Sciences, 2014

**Pivot of Merit Student**, University of Chinese Academy of Sciences, 2014

**National Scholarship**, China, 2011, 2014

**Third Place in ACM SIGSPATIAL GIS CUP**, 2012

**Outstanding University Graduate**, Beijing, China, 2011