Jiashuo Liu

Department of Computer Science and Technology, Tsinghua University
East Main Building 9-316, Beijing, China, 100084
Email: liujiashuo77@gmail.com Website: https://ljsthu.github.io Twitter: @liujiashuo77

Education

Department of Computer Science and Technology 2020.09 - 2025.06 (expexted)
Tsinghua University, Beijing, China
Ph.D. Candidate, advised by Prof. Peng Cui and Prof. Bo Li
GPA: 3.97/4.0; Rank: 3/129

Department of Computer Science and Technology 2016.08 - 2020.06
Tsinghua University, Beijing, China
Bachelor of Engineering
GPA: 3.81/4.0; Rank:11/169

Experience

Applied Mathematics and Theoretical Physics 2024.06 - 2024.11
Cambridge University, Cambridge, UK
Advisor: Prof. Mihaela van der Schaar
Visiting Student Researcher

Management Science and Engineering 2023.10 - 2024.04
Stanford University, California, US
Advisor: Prof. Jose Blanchet
Visiting Student Researcher

Decision, Risk, and Operations division 2023.01 - 2023.10
Columbia Business School, New York, US
Advisor: Prof. Hongseok Namkoong
Visiting Student Researcher (remote)

Invited Talks

[T5] Empirical Analysis & Inductive Modeling for Distribution Shifts on Tabular Data
Jiashuo Liu
Invited talk at Department of Management Science and Engineering, Tsinghua, 2024

[T4] Stability Evaluation via Distributional Perturbation Analysis
Jiashuo Liu
Invited presentation at 2024 INFORMS Annual Meeting
Session: Advances in Data-Driven Distributionally Robust Optimization

[T3] Data Heterogeneity Analysis for Out-of-Distribution Generalization
Peng Cui, **Jiashuo Liu**
Tutorial *at Conference on Lifelong Learning Agents (**CoLLAs**), 2024*

[T2] Model the Data Heterogeneity for Out-of-Distribution Generalization
Peng Cui, **Jiashuo Liu**, Bo Li, Renzhe Xu
Tutorial *at SIAM International Conference on Data Mining (**SDM**), 2024*

[T1] Modeling & Exploiting Data Heterogeneity under Distribution Shifts
Website: https://sites.google.com/view/neurips2023-tutorial-datahete
Video: https://neurips.cc/virtual/2023/tutorial/73953
Jiashuo Liu, Tiffany (Tianhui) Cai, Peng Cui, Hongseok Namkoong

Tutorial *at Neural Information Processing Systems (**NeurIPS**), 2023*
Hightlighted as NeurIPS 2023 Favorite Papers/Presentations by Two Sigma (9/3500+)

Publications

My research lies at the interface of machine learning and operations research, including:

- Generalization under Distribution Shifts: Theoretical foundations and tools for understanding real-world distribution shifts;
- Data Heterogeneity: Algorithms to model and exploit data heterogeneity under distribution shifts;
- Distributionally robust optimization: Addressing the over-pessimism problem in real applications.

Recently, I'm focusing on the evaluation of LLMs as well as incorporating LLMs into robust optimization.

($*$ indicates equal contribution; † denotes alphabetical order.)

First Author

[13] Stability Evaluation of Large Language Models via Distributional Perturbation Analysis
Jiashuo Liu, Jiajin Li, Peng Cui, Jose Blanchet
NeurIPS'24 Workshop on Red Teaming GenAI

[12] LLM Embeddings Improve Test-time Adaptation to Tabular Y—X-Shifts
Yibo Zeng*, **Jiashuo Liu***, Henry Lam, Hongseok Namkoong
NeurIPS'24 Workshop on Table Representation Learning

[11] On the Need of a Modeling Language for Distribution Shifts: Illustrations on Tabular Datasets
Jiashuo Liu*, Tianyu Wang*, Peng Cui, Hongseok Namkoong
INFORMS Workshop on Data Science 2024 (full paper)
Under review at Management Science

[10] Stability Evaluation via Distributional Perturbation Analysis
Jose Blanchet†, Peng Cui†, Jiajin Li†, **Jiashuo Liu**†
*In the International Conference on Machine Learning (**ICML**), 2024*

[9] Geometry-Calibrated DRO: Combating Over-Pessimism with Free Energy Implications
Jiashuo Liu, Jiayun Wu, Tianyu Wang, Hao Zou, Peng Cui
*In the International Conference on Machine Learning (**ICML**), 2024*
*In NeurIPS Information Processing Systems (**NeurIPS**), 2023, Workshop on Distribution Shifts.*

[8] Enhancing Distributional Stability among Sub-Populations
Jiashuo Liu, Jiayun Wu, Jie Peng, Xiaoyu Wu, Yang Zheng, Bo Li, Peng Cui
*In the International Conference on Artificial Intelligence and Statistics (**AISTATS**), 2024*

[7] On the Need for a Language Describing Distribution Shifts: Illustrations on Tabular Datasets
Jiashuo Liu*, Tianyu Wang*, Peng Cui, Hongseok Namkoong
*In Neural Information Processing Systems (**NeurIPS**), 2023, Datasets and Bench-*

marks Track
Hightlighted as NeurIPS 2023 Favorite Papers by Two Sigma (9/3500+)

[6] Measure the Predictive Heterogeneity
Jiashuo Liu, Jiayun Wu, Renjie Pi, Renzhe Xu, Xingxuan Zhang, Bo Li, Peng Cui
*The 11th International Conference on Learning Representations (**ICLR**), 2022*

[5] Distributionally Robust Learning with Stable Adversarial Training
Jiashuo Liu, Zheyan Shen, Peng Cui, Linjun Zhou, Kun Kuang, Bo Li
*In IEEE Transactions on Knowledge and Data Engineering (**TKDE**), 2022*

[4] Distributionally Robust Optimization with Data Geometry
Jiashuo Liu*, Jiayun Wu*, Bo Li, Peng Cui
*In Neural Information Processing Systems (**NeurIPS**), 2022*
[Spotlight, top 3%]

[3] Kernelized Heterogeneous Risk Minimization
Jiashuo Liu*, Zheyuan Hu*, Peng Cui, Bo Li, Zheyan Shen
*In Neural Information Processing Systems (**NeurIPS**), 2021*

[2] Heterogeneous Risk Minimization
Jiashuo Liu, Zheyuan Hu, Peng Cui, Bo Li, Zheyan Shen
*In International Conference on Machine Learning (**ICML**), 2021*
[Short talk, top 21.5%]

[1] Stable Adversarial Learning under Distributional Shifts
Jiashuo Liu, Zheyan Shen, Peng Cui, Linjun Zhou, Kun Kuang, Bo Li, Yishi Lin
*In AAAI Conference on Artificial Intelligence (**AAAI**), 2021.*

Others

[11] Bridging Multicalibration and Out-of-distribution Generalization Beyond Covariate Shift
Jiayun Wu, **Jiashuo Liu**, Peng Cui, Zhiwei Steven Wu
*In Neural Information Processing Systems (**NeurIPS**), 2024*

[10] Domain-wise Data Acquisition to Improve Performance under Distribution Shift
Yue He, Dongbai Li, Pengfei Tian, Han Yu, **Jiashuo Liu**, Hao Zou, Peng Cui
*In the International Conference on Machine Learning (**ICML**), 2024*

[9] Distributionally Generative Augmentation for Fair Facial Attribute Classification
Fengda Zhang, Qianpei He, Kun Kuang, **Jiashuo Liu**, Long Chen, Chao Wu, Jun Xiao, Hanwang Zhang
*In the Conference on Computer Vision and Pattern Recognition (**CVPR**), 2024*

[8] Rethinking the Evaluation Protocol of Domain Generalization
Han Yu, Xingxuan Zhang, Renzhe Xu, **Jiashuo Liu**, Yue He, Peng Cui
*In the Conference on Computer Vision and Pattern Recognition (**CVPR**), 2024*

[7] Towards Robust Out-of-Distribution Generalization Bounds via Sharpness
Yingtian Zou, Kenji Kawaguchi, Yingnan Liu, **Jiashuo Liu**, Mong-Li Lee, Wynne Hsu
*In the International Conference on Learning Representations (**ICLR**), 2024.*
[Spotlight, top 5%]

[6] Offline Policy Evaluation in Large Action Spaces via Outcome-Oriented Action Grouping
Jie Peng, Hao Zou, **Jiashuo Liu**, Shaoming Li, Yibao Jiang, Jian Pei, Peng Cui
*The ACM Web Conference (**WWW**), 2023.*

[5] Towards the ultimate PMT waveform analysis for neutrino and dark matter experiments
Dacheng Xu, Benda Xu, Erjin Bao, Yiyang Wu, Aiqiang Zhang, Yuyi Wang, Geliang Zhang, Yu Xu, Ziyi Guo, Jihui Pei, Hanyang Mao, **Jiashuo Liu**, Zhe Wang, Shaomin Chen
*In Journal of Instrumentation (**JINST**), 2022.*

[4] Invariant Preference Learning for General Debiasing in Recommendation
Zimu Wang, Yue He, **Jiashuo Liu**, Wenchao Zou, Philip Yu, Peng Cui
*In SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD**), 2022.*

[3] Triple Generative Adversarial Networks.
Chongxuan Li, Kun Xu, Jun Zhu, **Jiashuo Liu**, Bo Zhang
*In Transactions on Pattern Analysis and Machine Intelligence (**TPAMI**), 2021.*

[2] Signed Graph Neural Network with Latent Groups
Haoxin Liu, Ziwei Zhang, Peng Cui, Yafeng Zhang, Qiang Cui, **Jiashuo Liu**, Wenwu Zhu
*In SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD**), 2021.*

[1] Stable Learning via Differentiated Variable Decorrelation
Zheyan Shen, Peng Cui, **Jiashuo Liu**, Tong Zhang, Bo Li, Zhitang Chen
*In SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD**), 2020.*

Preprints

(∗ indicates equal contribution.)

[P3] Topology-Aware Dynamic Reweighting for Distribution Shifts on Graph
Weihuang Zheng*, **Jiashuo Liu***, Jiaxing Li, Jiayun Wu, Peng Cui, Youyong Kong
Under Review at NeurIPS 2024

[P2] Predictive Heterogeneity: Measures and Applications
Jiashuo Liu, Jiayun Wu, Bo Li, Peng Cui
Revision & Resubmit at JMLR

[P1] Towards Out-of-Distribution Generalization: A Survey
Jiashuo Liu*, Zheyan Shen*, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, Peng Cui
under review

Patents

Invariant learning method and device based on heterogeneity hybrid data. Issued 2023
Peng Cui, **Jiashuo Liu**
CN Patent: CN 113205184 B, filed April 28, 2021, and issued January 31, 2023.

Distribution robustness adversarial learning method. Issued 2022
Peng Cui, **Jiashuo Liu**
CN Patent: CN 112085194 B, filed August 30, 2020, and issued December 13, 2022.

Selected Awards

INFORMS Data Science Student Scholarship 2024
Tsinghua Excellent Comprehensive Scholarship 2022, 2023, 2024
Two Sigma Favorite Paper at NeurIPS [9/3500+] 2023

Apple Scholar in AI/ML Nomination [top 2 at Tsinghua]	2021
China National Scholarship [nationwide, top 1%]	2021.10
Tsinghua Excellent Undergraduate [top 10% at Tsinghua]	2020.06
Tsinghua TP-Link Scholarship	2019.11
Tsinghua-Toyota Scholarship	2018.11
Tsinghua Excellent Comprehensive Scholarship	2017.11
Tsinghua Second-Class Freshmen Scholarship [top 2 in Nei Mongol]	2016.08

Services

Reviewer for:
Journal: Operations Research, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), IEEE Transactions on Knowledge and Data Engineering (TKDE), IEEE Transactions on Multimedia (TMM)
Conference: ICLR: 2024, 2025; ICML: 2022, 2023, 2024; NeurIPS: 2023, 2024; UAI: 2022,2023,2024; AISTATS: 2021,2023,2024, 2025; CVPR: 2022,2023,2024; ECCV: 2024; ICCV: 2023; AAAI: 2022; IJCAI: 2022,2023; ICDM: 2024;

Teaching

Machine Learning Summer School TA, Cambridge University, instructed by Prof. Mihaela van der Schaar
Summer 2024

Software Engineering TA, Tsinghua University, instructed by Prof. Xiaoying Bai, Dan Pei, Jianyong Wang
Fall 2019, 2020, 2021, 2022, Spring 2022, 2023

Object-oriented Programming TA, Tsinghua University, instructed by Prof. Jingtao Fan
Summer 2022

Computer Skills

Languages: C, C++, Python
Mathematical Computation: Matlab, Mosek, Gurobi
Operating Systems: Linux, Mac OSX, Windows.