Tianyi Peng (2025/12/01)

Contact Information

tianyi.peng@columbia.edu https://tianyipeng.github.io

Research Interests

Digital Twin, Reinforcement Learning, Efficient Agent Deployment, A/B Testing and Causal Inference

Academic Appointments

Columbia University, Graduate School of Business 2024 - present
Decision, Risk, & Operations Division
Assistant Professor

Academic Degrees

Massachusetts Institute of Technology 2017 - 2023
Ph.D. in Statistics and Aeronautics/Astronautics, GPA: 5.0/5.0
Field in Operations Research
Advisor: Vivek Farias

Tsinghua University 2013 - 2017
Bachelor in Computer Science
- ⋆ Selected for the *Yao Class* (a CS pilot program led by Prof. Andrew C. Yao)
- ⋆ Graduated with Best Thesis Award

Honors and Awards

Winner, INFORMS Junior Faculty Interest Group Paper Competition	2025
Second Place, George Nicholson Student Paper Competition	2025
Finalist, Jeff McGill Student Paper Award	2025
Winner, Jeff McGill Student Paper Award	2022
Winner, Applied Probability Society Best Student Paper Prize	2022
Winner, Daniel H. Wagner Prize for Excellence	2022
Finalist, MSOM Best Student Paper Prize	2022
Best Paper Award, International Conference on Computing, Networking and Communications	2020
1st place, MIT Quantum Hackathon	2020
2nd place (among 2780 teams), IEEE programming competition IEEExtreme 13.0	2019
Best Thesis Award, Tsinghua University	2017
China 12-person team for International Olympiad in Informatics (IOI)	2013
International Gold Prize, Asia-Pacific Informatics Olympiad (APIO)	2012

Preprints

** indicates that the authors are in alphabetical order*

[1] **E-GEO: A Testbed for Generative Engine Optimization in E-Commerce**
* Puneet S Bagga, Vivek F Farias, Tamar Korkotashvili, Tianyi Peng, Yuhang Wu
arXiv preprint arXiv:2511.20867

[2] **A Mega-Study of Digital Twins Reveals Strengths, Weaknesses, and Opportunities for Further Improvement**
Tianyi Peng, George Gui, Daniel J. Merlau, Grace Jiarui Fan, Malek Ben Sliman, Melanie Brucks, Eric J. Johnson, Vicki Morwitz, Abdullah Althenayyan, Silvia Bellezza, Dante Donati, Hortense Fong, Elizabeth Friedman, Ariana Guevara, Mohamed Hussein, Kinshuk Jerath, Bruce Kogut, Akshit Kumar, Kristen Lane, Hannah Li, Patryk Perkowski, Oded Netzer, Olivier Toubia
arXiv: `https://arxiv.org/abs/2509.19088`

[3] **RAISE: Reliable Agent Improvement via Simulated Experience**
Sahar Omidi Shayegan, Joshua Meyer, Victor Shih, Sebastian Sosa, Tianyi Peng, Kostis Kaffes, Eugene Wu, Andi Partovi, Mehdi Jamei
Scaling Environments for Agents Workshop, NeurIPS 2025 Workshop

[4] **Speculative Actions: A Lossless Framework for Faster Agentic Systems**
* Naimeng Ye, Arnav Ahuja, Georgios Liargkovas, Yunan Lu, Kostis Kaffes, Tianyi Peng
arXiv preprint arXiv:2510.04371

[5] **Estimation of Treatment Effects Under Nonstationarity via Truncated Difference-in-Q's**
* Ramesh Johari, Tianyi Peng, Wenqian Xing
arXiv preprint arXiv:2506.05308

[6] **Agents for Web Testing: A Case Study in the Wild**
Naimeng Ye, Xiao Yu, Ruize Xu, Tianyi Peng, Zhou Yu
NeurIPS 2025 Workshop on Bridging Language, Agent, and World Models

[7] **Throughput-Optimal Scheduling Algorithms for LLM Inference and AI Agents**
Yueying Li, Tianze Deng, Jim Dai, Tianyi Peng
arXiv: `https://arxiv.org/abs/2504.07347`

[8] **How Well Do LLMs Compress Their Own Chain-of-Thought? A Token Complexity Approach**
Ayeong Lee, Ethan Che, Tianyi Peng
arXiv: `https://arxiv.org/abs/2503.01141`

Publications

** indicates that the authors are in alphabetical order*

[1] **Multi-Agent Markov Entanglement**
Shuze Chen, Tianyi Peng
NeurIPS 2025 (Spotlight)
⋆ *Winner*, **INFORMS Junior Faculty Interest Group (JFIG) Paper Competition 2025**
⋆ *Second Place*, **George Nicholson Student Paper Competition 2025**

[2] **Tail-Optimized Caching for LLM Inference**
Wenxin Zhang, Yueying Li, Ciamac C. Moallemi, Tianyi Peng
NeurIPS 2025

[3] **Twin-2K-500: A Dataset for Building Digital Twins of Over 2,000 People Based on Their Answers to Over 500 Questions**
Olivier Toubia, George Z. Gui, Tianyi Peng, Daniel J. Merlau, Ang Li, Haozhe Chen
Marketing Science

[4] **Data Mixture Optimization: A Multi-Fidelity Multi-Scale Bayesian Framework**

Tzu-Ching Yen, Andrew Wei Tung Siah, Haozhe Chen, C. Daniel Guetta, Tianyi Peng, Hongseok Namkoong
NeurIPS 2025

[5] **LLM-Generated Persona is a Promise with a Catch**
Ang Li, Haozhe Chen, Hongseok Namkoong, Tianyi Peng
NeurIPS 2025 Position Paper

[6] **Speeding up Policy Simulation in Supply Chain RL**
* Vivek Farias, Joren Gijsbrechts, Aryan Khojandi, Tianyi Peng, Andrew Zheng
ICML 2025

[7] **Performance of LLMs on Stochastic Modeling Operations Research Problems: From Theory to Practice**
* Akshit Kumar, Tianyi Peng, Yuhang Wu, Assaf Zeevi
Winter Simulation Conference 2024

[8] **Differences-in-Neighbors for Network Interference in Experiments**
* Tianyi Peng, Naimeng Ye, Andrew Zheng
EC 2025
⋆ *Finalist*, **Revenue Management and Pricing 2025 Jeff McGill Student Paper Award**

[9] **QGym: Scalable Simulation and Benchmarking of Queuing Network Controllers**
Haozhe Chen, Ang Li, Ethan Che, Tianyi Peng, Jing Dong, Hongseok Namkoong
NeurIPS 2024 Datasets and Benchmarks Track

[10] **Correcting for interference in experiments: A case study at douyin**
* Vivek Farias, Hao Li, Tianyi Peng, Xinyuyang Ren, Huawei Zhang, Andrew Zheng
Recsys 2023

[11] **Learning Treatment Effects in Panels with General Intervention Patterns**
* Vivek Farias, Andrew Li, Tianyi Peng
NeurIPS 2021 (Oral, top 0.6% of submissions)
Major revision in *Operations Research*
⋆ *Finalist*, **MSOM Best Student Paper Prize 2022**

[12] **Markovian Interference in Experiments**
* Vivek Farias, Andrew Li, Tianyi Peng, Andrew Zheng
NeurIPS 2022 (Oral, top 2% of submissions)
Minor revision in *Management Science*
⋆ *Winner*, **Applied Probability Society Best Student Paper Prize 2022**
⋆ *Winner*, **RMP Jeff McGill Student Paper Award 2022**

[13] **Generalized Synthetic Control for TestOps at ABI**
* Vivek Farias, Tianyi Peng, et al.
INFORMS Journal on Applied Analytics
⋆ *Winner*, **INFORMS Daniel H. Wagner Prize 2022**

[14] **Fixing Inventory Inaccuracies at Scale**
* Vivek Farias, Andrew Li, Tianyi Peng
ICML 2021, MSOM Supply Chain SIG 2022
Manufacturing & Service Operations Management

[15] **Synthetically Controlled Bandits**
* Vivek Farias, Ciamac Moallemi, Tianyi Peng, Andrew Zheng
MSOM Service Management SIG 2022
Major revision in *Management Science*

[16] **The Limits to Learning a Diffusion Model**
* Jackie Baek, Vivek Farias, Andreea Georgescu, Retsef Levi, Tianyi Peng, Deeksha Sinha, Joshua Wilde, Andrew Zheng
EC 2021
Management Science

[17] **Uncertainty Quantification for Low-Rank Matrix Completion with Heterogeneous and Sub-Exponential Noise**
* Vivek Farias, Andrew Li, Tianyi Peng
AISTATS 2022

[18] **Optimal Entanglement Swapping and Distribution**
Wenhan Dai, Tianyi Peng, Moe Win
IEEE Journal on Selected Areas in Communications, vol. 38, pp. 540–556, 2020
⋆ **Best Paper Award**, *International Conference on Computing, Networking and Communications (ICNC 2020)*

[19] **Quantum Queuing Delay**
Wenhan Dai, Tianyi Peng, Moe Win
IEEE Journal on Selected Areas in Communications, vol. 38, pp. 605–618, 2020

[20] **Simulating Large Quantum Circuits on a Small Quantum Computer**
Tianyi Peng, Maris Ozols, Aram Harrow, Xiaodi Wu
Physical Review Letters **125**, 150504 (2020)

[21] **Quantum Uncertainty Relation of Coherence**
Xiao Yuan, Ge Bai, Tianyi Peng, Xiongfeng Ma
Physical Review A **96**(3), 032313

[22] **Tight Detection Efficiency Bounds of Bell Tests in No-Signaling Theories**
Zhu Cao, Tianyi Peng
Physical Review A **94**, 042126

[23] **Efficient and Robust Physical Layer Key Generation**
Tianyi Peng, Wenhan Dai, Moe Win
MILCOM 2019

[24] **Remote State Preparation for Multiple Parties**
Wenhan Dai, Tianyi Peng, Moe Win
IEEE ICASSP 2019, pp. 7983–7987, **Invited Paper**

Teaching Experience

Generative AI: Technical and Social (B9153)
Instructor, Columbia University — Fall 2024, 2025

Business Analytics (B6101)
Instructor, Columbia University — Spring 2025

Hands-on Deep Learning (15.S04)
Teaching Assistant for MBA Students, MIT, Rating 6.9/7.0 — Spring 2022

Quantum Information and Quantum Computation
Lecturer for MIT High School Studies Program (Not Rated) — Summer 2019

Service

Reviewer for *Management Science, Operations Research, Annals of Statistics, Mathematical Programming, Biometrika, ICLR 2026, AAAI 2026, NeurIPS 2025, ICML 2025, NeurIPS 2024, EC2024, AAAI 2023, AISTATS 2022, IEEE Journal on Selected Areas in Communications, Quantum, ACM Transactions on Quantum Computing, New Journal of Physics*

Organizer, Frontiers in AI Symposium at Columbia Business School, Fall 2024
Organizer, MIT LIDS Student Conference 2020

Professional Experience

Cimulate.AI
Develop rerank models and shopping agents in e-commerce.
AI Scientist, Founding team — 2023 - 2024

Anheuser-Busch InBev
Leading the development of TestOps, a pioneering experimentation platform for physical retailers.
Student collaborator — 2020 - 2023

TikTok (ByteDance)
Addressing interference problems in the experimentation platform at Bytedance and developing recommendation algorithms at TikTok.
Student collaborator — 2021 - 2023

Liberty Mutual
Developing novel data-imputation methods for improving insurance pricing.
Student collaborator — 2021 - 2023

Outside Activities

Columbia Business School requires its faculty members to disclose any activities that might present a real or apparent conflict of interest. The list below complies with this requirement

Cimulate.AI — 2024 - current
Consultant