

Lingkai KONG

150 Western Ave, Allston, MA 02134

✉ linkaikong@g.harvard.edu ☎ (+1)404-317-0711 🏠 <https://lingkai-kong.com/>

RESEARCH INTERESTS

My research advances **generative AI** with **optimization** and **reinforcement learning** to enable reliable decision-making in high-stakes environments. I work closely with public-sector partners in health and environmental sustainability to translate algorithmic innovations into real-world impact. These domains are characterized by limited data, high uncertainty, and the need to make timely decisions across large and complex action spaces under significant risk.

ACADEMIC POSITIONS

Harvard University

Postdoctoral Fellow

Cambridge, USA

2024 - present

- Advisor: Milind Tambe

EDUCATION

Georgia Institute of Technology

Ph.D. in Computational Science and Engineering

Atlanta, USA

2019 - 2024

- Advisor: Chao Zhang
- Thesis Committee: B. Aditya Prakash, Yao Xie, Bo Dai

Georgia Institute of Technology

Ph.D. in Electrical and Computer Engineering (transferred to CSE)

Atlanta, USA

2017 - 2019

Southeast University

B.E. in Information Engineering

Nanjing, China

2013 - 2017

PREPRINTS

1. Diffusion-DFL: Decision-Focused Diffusion Models for Stochastic Optimization
Zihao Zhao, Christopher Yeh, **Lingkai Kong**, Kai Wang
Under review, 2025

PUBLICATIONS

* *Denotes equal contribution*

1. Generative AI Against Poaching: Latent Composite Flow Matching for Wildlife Conservation
Lingkai Kong*, Haichuan Wang*, Charles A. Emogor, Vincent Börsch-Supan, Lily Xu, Milind Tambe
Innovative Applications of Artificial Intelligence (IAAI). 2026
Oral Presentation
2. Composite Flow Matching for Reinforcement Learning with Shifted-Dynamics Data
Lingkai Kong*, Haichuan Wang*, Tonghan Wang*, Guojun Xiong, Milind Tambe
Advances in Neural Information Processing Systems (NeurIPS), 2025.
Spotlight Presentation (top 3%)
3. Robust Optimization with Diffusion Models for Green Security
Lingkai Kong, Haichuan Wang, Yuqi Pan, Cheol Woo Kim, Mingxiao Song, Alayna Nguyen, Tonghan Wang, Haifeng Xu, Milind Tambe
Conference on Uncertainty in Artificial Intelligence (UAI), 2025.

4. What is the Right Notion of Distance between Predict-then-Optimize Tasks?
Paula Rodriguez-Diaz, **Lingkai Kong**, Kai Wang, David Alvarez-Melis, Milind Tambe
Conference on Uncertainty in Artificial Intelligence (UAI), 2025.
5. DF²: Distribution-Free Decision-Focused Learning
Lingkai Kong, Wenhao Mu, Jiaming Cui, Yuchen Zhuang, B Aditya Prakash, Bo Dai, Chao Zhang
Conference on Uncertainty in Artificial Intelligence (UAI), 2025.
6. PRIORITY2REWARD: Incorporating Healthworker Preferences for Resource Allocation Planning
Shresth Verma, Alayna Nguyen, Niclas Boehmer, **Lingkai Kong**, Milind Tambe
AAAI Conference on Artificial Intelligence (AAAI), 2025.
7. Navigating the Social Welfare Frontier: Portfolios for Multi-objective Reinforcement Learning
Cheol Woo Kim, Jai Moondra, Shresth Verma, Madeleine Pollack, **Lingkai Kong**, Milind Tambe, Swati Gupta
International Conference on Machine Learning (ICML), 2025.
8. LLM-Augmented Chemical Synthesis and Design Decision Programs
Haorui Wang, Jeff Guo, **Lingkai Kong**, Rampi Ramprasad, Philippe Schwaller, Yuanqi Du, Chao Zhang
International Conference on Machine Learning (ICML), 2025.
9. Diffusion Models as Constrained Samplers for Optimization with Unknown Constraints
Lingkai Kong*, Yuanqi Du*, Wenhao Mu*, Kirill Neklyudov, Valentin De Bortol, Haorui Wang, Dongxia Wu, Aaron Ferber, Yi-An Ma, Carla P. Gomes, Chao Zhang
International Conference on Artificial Intelligence and Statistics (AISTATS), 2025.
10. Efficient Evolutionary Search Over Chemical Space with Large Language Models
Haorui Wang, Marta Skreta, Cher-Tian Ser, Wenhao Gao, **Lingkai Kong**, Felix Strieth-Kalthoff, Chenru Duan, Yuchen Zhuang, Yue Yu, Yanqiao Zhu, Yuanqi Du, Alán Aspuru-Guzik, Kirill Neklyudov, Chao Zhang
International Conference on Representation Learning (ICLR), 2025.
11. Aligning Large Language Models with Representation Editing: A Control Perspective
Lingkai Kong*, Haorui Wang*, Wenhao Mu*, Yuanqi Du, Yuchen Zhuang, Yifei Zhou, Yue Song, Rongzhi Zhang, Kai Wang, Chao Zhang
Advances in Neural Information Processing Systems (NeurIPS), 2024.
12. Time-MMD: A New Multi-Domain Multimodal Dataset for Time Series Analysis
Haoxin Liu, Shangqing Xu, Zhiyuan Zhao, **Lingkai Kong**, Harshavardhan Kamarthi, Aditya B Sasanur, Megha Sharma, Jiaming Cui, Qingsong Wen, Chao Zhang, B Aditya Prakash
Advances in Neural Information Processing Systems (NeurIPS), 2024.
13. Time-Series Forecasting for Out-of-Distribution Generalization Using Invariant Learning
Haoxin Liu, Harshavardhan Kamarthi, **Lingkai Kong**, Zhiyuan Zhao, Chao Zhang, B. Aditya Prakash
International Conference on Machine Learning (ICML), 2024.
14. Two Birds with One Stone: Enhancing Uncertainty Quantification and Interpretability with Graph Functional Neural Process
Lingkai Kong*, Haotian Sun*, Yuchen Zhuang, Haorui Wan, Wenhao Mu, Chao Zhang
International Conference on Artificial Intelligence and Statistics (AISTATS), 2024.
15. AdaPlanner: Adaptive Planning from Feedback with Language Models
Haotian Sun, Yuchen Zhuang, **Lingkai Kong**, Bo Dai, Chao Zhang
Advances in Neural Information Processing Systems (NeurIPS), 2023.
16. MUBen: Benchmarking the Uncertainty of Pre-Trained Models for Molecular Property Prediction
Yinghao Li, **Lingkai Kong**, Yuanqi Du, Yue Yu, Yuchen Zhuang, Wenhao Mu, Chao Zhang
Transaction on Machine Learning Research (TMLR).
17. When Rigidity Hurts: Soft Consistency Regularization for Probabilistic Hierarchical Time Series Forecasting
Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodríguez, Chao Zhang, B. Aditya Prakash
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2023.
18. DyGen: Fine-Tuning Language Models with Noisy Labels by Dynamics-Enhanced Generative Modeling
Yuchen Zhuang, Yue Yu, **Lingkai Kong**, Xiang Chen, Chao Zhang
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2023.

19. Autoregressive Diffusion Model for Graph Generation
Lingkai Kong, Jiaming Cui, Haotian Sun, Yuchen Zhuang, B. Aditya Prakash, Chao Zhang
International Conference on Machine Learning (ICML), 2023.
20. End-to-End Stochastic Optimization with Energy-based Model
Lingkai Kong, Jiaming Cui, Yuchen Zhuang, Rui Feng, B. Aditya Prakash, Chao Zhang
Advances in Neural Information Processing Systems (NeurIPS), 2022.
Oral Presentation (top 1.8%)
21. AcTune: Uncertainty-Aware Active Self-Training for Active Fine-Tuning of Pretrained Language Models
Yue Yu, **Lingkai Kong**, Jieyu Zhang, Rongzhi Zhang, Chao Zhang
Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2022.
22. CAMul: Calibrated and Accurate Multi-view Time-Series Forecasting
Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodríguez, Chao Zhang, B. Aditya Prakash
The Web Conference (WWW), 2022.
23. When in Doubt: Neural Non-Parametric Uncertainty Quantification for Epidemic Forecasting
Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodríguez, Chao Zhang, B. Aditya Prakash
Advances in Neural Information Processing Systems (NeurIPS), 2021.
24. Calibrated Language Model Fine-Tuning for In- and Out-of-Distribution Data
Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lye, Tuo Zhao, Chao Zhang
Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020.
25. SDE-Net: Equipping Deep Neural Networks with Uncertainty Estimates
Lingkai Kong, Jimeng Sun, Chao Zhang
International Conference on Machine Learning, (ICML), 2020.
26. Learning Deep Hidden Nonlinear Dynamics from Aggregate Data
Yisen Wang, Bo Dai, **Lingkai Kong**, Hongyuan Zha
Conference on Uncertainty in Artificial Intelligence (UAI), 2018.

WORKSHOP PUBLICATIONS

1. LLM-Based Agent Simulation for Maternal Health Interventions: Uncertainty Estimation and Decision-Focused Evaluation
Sarah Martinson, **Lingkai Kong**, Cheol Woo Kim, Aparna Taneja, Milind Tambe
AAMAS Workshop on Autonomous Agents for Social Good (AAMAS-AASG), 2025.
2. Learning Graph Structures and Uncertainty for Accurate and Calibrated Time-series Forecasting
Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodríguez, Chao Zhang, B. Aditya Prakash
KDD Workshop on Uncertainty Reasoning and Quantification in Decision Making (KDD-UDM), 2024.

CONFERENCE TUTORIAL

- Uncertainty Quantification in Deep Learning
Lingkai Kong, Harshavardhan Kamarthi, Peng Chen, B. Aditya Prakash, Chao Zhang
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2023.

HONORS AND AWARDS

- | | |
|---|------|
| • NeurIPS Spotlight (top 3%) | 2025 |
| • ICML Travel Award | 2023 |
| • NeurIPS Oral (top 1.8%) | 2022 |
| • NeurIPS Travel Award | 2022 |
| • Georgia Tech Otto & Jenny Krauss Fellowship | 2017 |

- Outstanding Undergraduate Thesis, Southeast University (Top 5%) 2017
- First Prize of China Undergraduate Mathematical Contest of Modeling (CUMCM) (Top5%) 2015

INVITED TALKS

- **INFORMS Annual Meeting** Oct. 2025
Robust Optimization with Diffusion Models for Green Security
- **Algorithmic Game Theory Workshop, TTIC** July 2025
Generative AI vs Poaching: Diffusion Models for Game-Theoretic Anti-poaching Patrols
- **NSF AI Institute for Societal Decision Making (AI-SDM) Workshop for High School Educators** July 2025
Generative AI for Wildlife Conservation
- **Harvard CS Colloquium Series, Harvard University** Feb. 2025
Robust Optimization with Diffusion Models for Green Security
- **AI4Science Seminar, University of Michigan** Nov. 2024
Uncertainty Quantification in Deep Learning
- **Advanced Topics in Machine Learning, Georgia Tech (Guest Lecture)** Oct. 2023
Advancing Data-driven Decision Making: An Uncertainty-based Predict-and-Optimize Approach
- **NSF AI Institute for Advances in Optimization (AI4OPT) Seminar, Georgia Tech** Mar. 2023
End-to-End Stochastic Optimization with Energy-based Model

WORKING EXPERIENCE

- Amazon** Seattle, USA
Applied Science Intern, Product Graph Team May 2021 - Nov 2021
- Mentor: Xiang He, Chenwei Zhang; Manager: Luna Xin Dong
 - Worked on multi-lingual attribute extraction
- IQVIA** Cambridge, USA
Research Intern, Analytics Center of Excellence Jun 2020 - Aug 2020
- Mentor: Danica (Cao), Xiao
 - Worked on uncertainty quantification for healthcare data

ACADEMIC SERVICE

Program Committee / Reviewer

NeurIPS	Conference on Neural Information Processing Systems	2023–2025
ICLR	International Conference on Learning Representations	2023–2025
ICML	International Conference on Machine Learning	2024–2025
JMLR	Journal of Machine Learning Research	2025
AAMAS	Autonomous Agents and Multi-Agent Systems	2025
AISTATS	Artificial Intelligence and Statistics	2024
UAI	Conference on Uncertainty in Artificial Intelligence	2024
AAAI	AAAI Conference on Artificial Intelligence	2024–2025
KDD	SIGKDD Conf. on Knowledge Discovery and Data Mining	2021–2024
ACL	Annual Meeting of the Association for Computational Linguistics	2021–2023
EMNLP	Empirical Methods in Natural Language Processing	2020–2023
NAACL	North American Chapter of the ACL	2021–2022

GRANT AND FUNDING WRITING EXPERIENCES

I have significantly contributed to the writing of the following grants and fundings:

- NSF III Medium: ‘Data-efficient Decision-Focused Learning’, PI: Bo Dai, Co-PI: Chao Zhang, B. Aditya Prakash, Kai Wang.
- NSF III Medium: ‘Principled Uncertainty Quantification in Deep Learning Models for Time Series Analysis’, PI: Chao Zhang, Co-PI: B. Aditya Prakash.
- NSF Career Award: ‘Accelerating Spatial Network Design: An Uncertainty-Driven Predict-and-Optimize Learning Framework’, PI: Chao Zhang.
- Amazon Research Award: ‘Enabling Pre-Trained Language Models for Open, Low-Resource Information Extraction’, PI: Chao Zhang.

STUDENT ADVISED

At Harvard University

Anagha Satish	2025–present	
Haichuan Wang	2024–present	
Mingxiao Song	2024–present	
Alayna Nguyen	2024–present	
Sarah Martinson	2024–2025	Next step: Ph.D. student at Harvard University

At Georgia Institute of Technology

Shangqing Xu	2023–2024	Next step: Ph.D. student at Georgia Tech
Wenhao Mu	2023–2024	Next step: Ph.D. student at University of Michigan
Haotian Sun	2022–2023	Next step: Ph.D. student at Georgia Tech
Yuchen Zhuang	2019–2021	Next step: Ph.D. student at Georgia Tech
Wendi Ren	2019–2021	Next step: Ph.D. student at CUHK-Shenzhen

TEACHING

Georgia Institute of Technology	Atlanta, USA
• Teaching Assistant, CX4240: Introduction to Computational Data Analysis	Spring, 2024
• Teaching Assistant, CSE8803 Deep Learning for Text Data	Fall, 2020
• Teaching Assistant, CSE8803 Deep Learning for Text Data	Fall, 2019
• Teaching Assistant, CS7641 Machine Learning (Online)	Spring, 2019
• Teaching Assistant, CS7641 Machine Learning (Online)	Fall, 2018