# Yingru Li

(086) 13006106862 szrlee@gmail.com https://richardli.xyz/


**Interests**: Sequential Decision-making, Optimization, Probability, AI & Operations.

## Education

- **The Chinese University of Hong Kong, Shenzhen**, China 2025 (expected)
  ***Ph.D.*** *in Computer Science and Information Engineering.*
  **Supervisor**: *Zhi-Quan (Tom) Luo* https://en.wikipedia.org/wiki/Zhi-Quan_Tom_Luo
  **Committee**: *Jim Dai*, *Xinyun Chen*, *Baoxiang Wang*, *Benjamin Van Roy* (Stanford & Google DeepMind)
- **Huazhong University of Science and Technology**, Wuhan, China 2020
  ***M.S.*** *in Computer Science.*
  Ranked 1st/134 overall, 1st/26 in Computer Theory and Software specialization.
- **Huazhong University of Science and Technology**, Wuhan, China 2017
  ***B.E.*** *in Computer Science (Honors Program). Outstanding Graduate.*

## Professional Experience

- **Tencent AI & Robotics X**, Shenzhen, China 2019, 2020
  *Research Intern* in Agent Center
  Participated in a team project on a high-throughput distributed Actor-Learner system for parallel on-policy rollout and sample-efficient asynchronous reinforcement learning with off-policy data reuse.
- **SenseTime Research**, Peking, China 2018
  *Computer Vision Trainee Researcher*
  Worked on an incremental and continual learning system for multi-label image classification.
- **Department of Computer Science, Cornell University**, Ithaca, NY 2017
  *Undergraduate Research Assistant*
  Worked with John E. Hopcroft on hidden community detection, a novel graph-theoretical concept.
- **Microsoft Research Lab**, Asia 2016
  *Research Intern* in Theory Center
  Worked on influence learning and maximization in social networks.
- **John Hopcroft Center on Computing Science**, Wuhan, China 2015-2017
  *Undergraduate Research Assistant*
  Led a reading group on machine learning and spectral graph theory. Worked on network analysis.

## Awards

- **Best Paper Award**, in the 3rd doctoral and postdoctoral Daoyuan academic forum, 2024.
- **Best Student Paper Award**, in the 13th IEEE Sensor Array and Multichannel Signal Processing Workshop, 2024.
- **SRIBD Ph.D. Fellowship** (Gold Class), by Shenzhen Research Institute of Big Data (SRIBD), 2023.
- **Presidential Ph.D. Fellowship**, by The Chinese University of Hong Kong, Shenzhen, 2019–2023.
- **Tencent AI Ph.D. Fellowship**, by Tencent & The Chinese University of Hong Kong, Shenzhen, 2018.
- **Award of Excellence in Internship**, by Microsoft Research Lab, 2016.
- **Qiming Star Award (Selected as one of 5 recipients out of 7,112 undergraduates.)**, by Huazhong University of Science and Technology, 2016. **Reports**: [1] Newspaper. [2] HUST Online.
- **National Scholarship**, by Huazhong University of Science and Technology, 2018.
- **First Prize**, in Parallel computation and Application Contest (PAC) held by Intel and CCF, 2015.
- **First Prize**, in China National Mathematics Olympiad (Province-level Math League), 2012.

## Selected Oral Presentations

- **GPT-HyperAgent: Adaptive Foundation Models for Online Decisions**
  *Invited talk* in 2024 INFORMS Annual Meeting, Seattle, Oct. 21, 2024.
- **HyperAgent: Advancing Scalable Exploration through Fast Uncertainty Estimation in RL**
  *Invited talk* in International Symposium on Mathematical Programming (ISMP), Montréal, Jul. 25, 2024.
- **HyperAgent: A Simple, Efficient and Scalable RL Framework for Complex Environments**
  a.k.a. "Q* meets Thompson Sampling: Scaling up Exploration via HyperAgent"
  *Invited talk* at RLChina.org, Jun. 25, 2024.
  *Invited talk* at Princeton University, May 2, 2024.
  *Invited talk* in INFORMS Optimization Society (IOS) Conference, Rice University, Mar. 23, 2024.
  *Contributed talk*, in the third doctoral and postdoctoral Daoyuan academic forum, Jan. 13, 2024.
- **No-Regret Learning in Unknown Game with Applications**
  *Invited talk* in RL Theory Student Workshop at Nanjing University, Aug. 23, 2022.
  *Contributed Talk* in the second doctoral and postdoctoral Daoyuan academic forum, Aug. 20, 2022.
- **HyperDQN: A Randomized Exploration Method for Deep Reinforcement Learning**
  *Contributed Talk* in NeurIPS Workshop Ecological Theory of Reinforcement Learning, Dec. 14, 2021

## Selected Research Publications

### Preprints

1. **Yingru Li**. *Simple, Unified Analysis of Johnson-Lindenstrauss with Applications*. Preprint. Presentation at ICML 2024 Workshop "High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning". arXiv: 2402.10232 [stat.ML].

2. **Yingru Li**, Liangqi Liu, Hao Liang, Wenqiang Pu, and Zhi-Quan Luo. *Optimistic Thompson Sampling for No-Regret Learning in Unknown Games*. Preprint. Presentation at ICML 2023 Workshop "The Many Facets of Preference-Based Learning". arXiv: 2402.09456 [cs.LG].

3. **Yingru Li**, Jiawei Xu, and Zhi-Quan Luo. *Adaptive Foundation Models for Online Decisions: HyperAgent with Fast Incremental Uncertainty Estimation*. Preprint. Presentation at ICML 2024 Workshops: (1) "Aligning Reinforcement Learning Experimentalists and Theorists"; (2) "Automated Reinforcement Learning: Exploring Meta-Learning, AutoML, and LLMs".

4. **Yingru Li**. *Probability Tools for Sequential Random Projection*. Preprint. Presentation at ICML 2024 Workshop "High-dimensional Learning Dynamics 2024: The Emergence of Structure and Reasoning". 2024. arXiv: 2402.14026 [math.ST].

### Conference Proceedings

5. **Yingru Li** and Zhi-Quan Luo. "Prior-dependent analysis of posterior sampling reinforcement learning with function approximation". In: *The 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2024. arXiv: 2403.11175 [stat.ML].

6. **Yingru Li**, Jiawei Xu, Lei Han, and Zhi-Quan Luo. "Q-Star Meets Scalable Posterior Sampling: Bridging Theory and Practice via HyperAgent". In: *The 41st International Conference on Machine Learning (ICML)*. Proceedings of Machine Learning Research. 2024. arXiv: 2402.10228 [cs.LG].

7. Liangqi Liu, Wenqiang Pu, **Yingru Li**, Bo Jiu, and Zhi-Quan Luo. "Radar Anti-jamming Strategy Learning via Domain-knowledge Enhanced Online Convex Optimization". In: *2024 IEEE 13th Sensor Array and Multichannel Signal Processing Workshop (SAM)*. IEEE. 2024. arXiv: 2402.16274 [eess.SP].

8. Ziniu Li, **Yingru Li**, Yushun Zhang, Tong Zhang, and Zhi-Quan Luo. "HyperDQN: A Randomized Exploration Method for Deep Reinforcement Learning". In: *International Conference on Learning Representations (ICLR)*. 2022. URL: https://openreview.net/pdf?id=X0nrKAXu7g-.

9. Qing Wang, **Yingru Li**, Jiechao Xiong, and Tong Zhang. "Divergence-Augmented Policy Optimization". In: *Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 32. Curran Associates, Inc., 2019. URL: https://openreview.net/pdf?id=rylacVSeIS.

### Journal Articles

10. Kun He, **Yingru Li**, Sucheta Soundarajan, and John E. Hopcroft. "Hidden community detection in social networks". In: *Information Sciences* 425 (2018), pp. 92–106. ISSN: 0020-0255. DOI: https://doi.org/10.1016/j.ins.2017.10.019.

### Under Preparation

11. **Yingru Li**, Xiangbo Wu, Yanchuan Tang, Xiang Wan, Benyou Wang, and Zhi-Quan Luo. "Uncertainty-aware Vision-Language Agents for Multi-turn Medical Decision-making". 2024.

12. **Yingru Li**, Jiawei Xu, Xiangbo Wu, Anningzhe Gao, Baoxiang Wang, Benyou Wang, and Zhi-Quan Luo. "Controlled Decoding via Q-Star on Outcome Feedback for Language Models". 2024.

13. Liangqi Liu, Wenqiang Pu, **Yingru Li**, Bo Jiu, and Zhi-Quan Luo. "Learning an Opponent-aware Anti-jamming Strategy via Online Convex Optimization". Extended version of "Radar Anti-jamming Strategy Learning via Domain-knowledge Enhanced Online Convex Optimization". 2024.

## Research Highlight

I focus on algorithms and theory for interactive agents that operate in complex and uncertain environments. This work necessitates advancements in methods for knowledge and uncertainty representation, exploration, adaptation, and decision-making. To achieve these goals, I use and develop fundamental tools in probability, optimization, game theory, and information theory. My methods have been applied to human-AI alignment and reliable & strategic operations. The significance of my work has been recognized through invitations to speak at *prestigious forums*, including ICML, NeurIPS, ICLR, AISTATS, ISMP and INFORMS Annual Meetings, and through awards, such as the *Best Paper Award* at the 2024 Daoyuan forum and the *Best Student Paper Award* at the 2024 IEEE SAM. Brief statement can be found at https://richardli.xyz/#research.

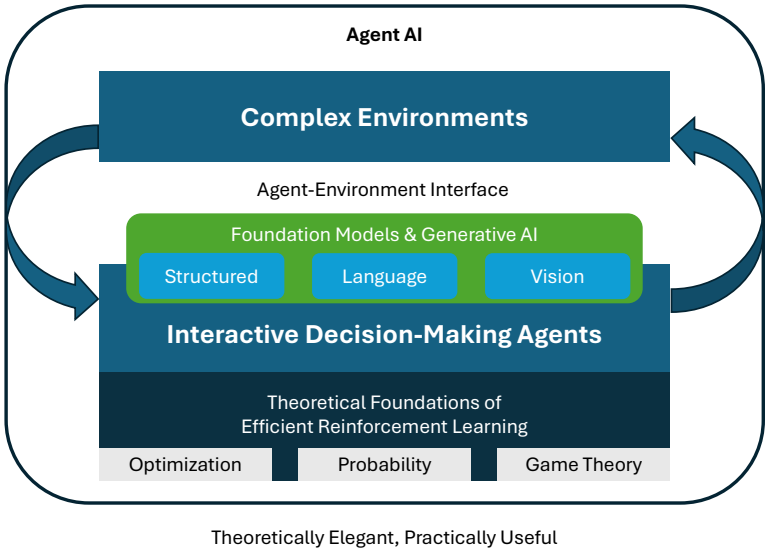


## Academic Service

- **Reviewer** for Conference on Neural Information Processing Systems (NeurIPS) [12 papers], International Conference on Learning Representations (ICLR) [2 papers]; ICLR Workshop "Bridging the Gap Between Practice and Theory in Deep Learning" (2 papers), ICML Workshop "Aligning Reinforcement Learning Experimentalists and Theorists" (2 papers).

## Academic Service (continued)

- **Chair** for RL Seminar in The Chinese University of Hong Kong, Shenzhen, China (Spring 2019, Summer 2020, Fall 2020, Spring 2021, Summer 2021, Fall 2021, Spring 2022, Fall 2022.); for 2 sessions in INFORMS Annual Meeting 2024 on "Integrating Generative AI with Sequential Decision-making".

## Teaching

- **Stochastic Processes** (STA/DDA4001) Fall 2018
- **Optimization II** (MAT3220) Spring 2019
- **Distributed and Parallel Computing** (CSC4005) Fall 2019
- **Reinforcement Learning** (DDA6105/CIE6023) Fall 2020
- **Matrix Analysis** (CIE6002) Spring 2021
- **Deep Learning and Their Applications** (MDS6224) Spring 2022

My teaching duties include delivering weekly tutorials, correcting assignments, and running laboratory sessions when required, **all in English**.

## Beyond Academia

I enjoy photography. I often play tennis and swim, and occasionally play golf. These activities allow me to live in the moment and help me find physical and spiritual freedom.