Jiachen Hu

No.5 Yiheyuan Road,
Haidian District, Beijing,
P.R.China 100871

Email: NickH@pku.edu.cn
Tel: (+86) 15911068937
Website: https://nickhclos.github.io/

EDUCATION

School of Computer Science, Peking University **2020.9 - 2025.7**
Ph.D candidate. Advisor: Liwei Wang

Turing Class, School of EECS, Peking University **2016.9 - 2020.6**
Bachelor of Science, **Summa Cum Laude** (Honorable BS)
GPA: 3.815/4

RESEARCH INTERESTS

- Understanding practical deep RL framework: the pros and cons of Transform-based RL from a theoretical perspective, the possibility of foundation model of decision making, the theoretical property of sim-to-real transfer
- Machine learning assisted mathematical discovery: (search-based) ML methods and traditional methods in cutting-edge research-level math problems
- Provable sample efficient RL
  - Efficient exploration with different structures: online exploration in linear/general MDP and structured POMDP
  - The online exploration on quantum computers: the fundamental separation of classical environments and quantum environments in terms of sample complexity of exploration.
  - Offline RL
- The application of RL in economical models: designing empirical data-driven methods with RL in automated market makers, designing provably efficient theoretical algorithms in sequential economical models.
- Classical online learning: multi-armed bandits, linear bandits

SELECTED PROJECTS

Distributed Acceleration of Bandits **2019**
Advisor: Liwei Wang

This project focuses on understanding theoretical acceleration of reinforcement learning in a collaborated distributed system. We hope to figure out two crucial quantity: the optimal speedup in terms of sample complexity, and the minimum communication overhead in order to achieve this optimal speedup. Starting from multi-armed bandits and linear bandits, we showed that the optimal linear speedup is possible. We constructed several sophisticated communication protocols to reduce the communication overhead, and proved that the communication overhead is optimal (excluding logarithmic terms) in multi-armed bandits setting.

The paper was accepted in ICLR 2020, which opened up a new direction of investigating the tradeoff between communication overhead and sample complexity in bandits and RL. It was followed by many theoretical works since then.

Understanding the Theoretical Benefits of Sim-to-Real Transfer 2021-2022
Advisor: Chi Jin, Lihong Li, Liwei Wang

We investigated the theoretical benefits of sim-to-real transfer, a promising methodology to reduce the high cost of obtaining real-world training data. As the first attempt to analyze the theoretical benefits in the literature, we formalize the simulator as a set of models with tunable parameters. We study the sim-to-real gap of two famous algorithms: domain randomization and robust adversarial training. Our results showed that these two algorithms are highly efficient to return a memory-augmented policy under mild structural conditions on the simulator set.

Two papers are published in ICLR 2022/2023, the former of which is accepted as a spotlight. This project provides insight of the crucial component in sim-to-real transfer.

The limitation of Transformers on HMM-like models **2023-2024**
Advisor: Chi Jin

Understanding the efficiency of Transformers in Hidden Markov Models (HMM) and its variants is crucial for developing robust Transformer-based RL agents, given that HMM serves as one of the fundamental sequential models. This study investigates the predictive capabilities of Transformers and Recurrent Neural Networks (RNN) in HMM-like models, focusing on both next observation and hidden state predictions. We designed HMM-like instances with varying levels of complexity and evaluated the performance of standard RNNs and Transformers. While RNNs consistently demonstrated strong performance across all tasks, shallow Transformers exhibited slightly larger prediction errors in simpler and medium complexity HMM-like models, and struggled to adapt to more challenging scenarios. Complementing our empirical findings, we constructed a theoretical framework for shallow Transformers tailored to simple and medium HMM-like models, demonstrating its efficacy in fitting these models accurately.

PREPRINTS

(* denotes equal contributions)

New sphere packings from the antipode construction

- Ruitao Chen, **Jiachen Hu**, Binghui Li, Liwei Wang, Tianyi Wu (alphabetical order)
- Arxiv preprint

On Limitation of Transformer for Learning HMMs

- **Jiachen Hu**, Qinghua Liu, Chi Jin
- Arxiv preprint

PUBLICATIONS

(* denotes equal contributions)

The Sample Complexity of Online Strategic Decision Making with Information Asymmetry and Knowledge Transportability

- **Jiachen Hu**, Rui Ai, Han Zhong, Xiaoyu Chen, Liwei Wang, Zhaoran Wang, Zhuoran Yang
- ICML 2025

Provably Efficient Exploration in Quantum Reinforcement Learning with Logarithmic Worst-Case Regret

- Han Zhong*, **Jiachen Hu***, Yecheng Xue, Tongyang Li, Liwei Wang
- ICML 2024

Quantum Non-Identical Mean Estimation: Efficient Algorithms and Fundamental Limits

- **Jiachen Hu**, Tongyang Li, Xinzhao Wang, Yecheng Xue, Chenyi Zhang, Han Zhong
- TQC 2024

ZeroSwap: Data-driven Optimal Market Making in DeFi

- Viraj Nadkarni, **Jiachen Hu**, Ranvir Rana, Chi Jin, Sanjeev Kulkarni, Pramod Viswanath
- FC 2024

Provable Sim-to-real Transfer in Continuous Domain with Partial Observations

- **Jiachen Hu***, Han Zhong*, Chi Jin, Liwei Wang
- ICLR 2023

Understanding Domain Randomization for Sim-to-real Transfer

- Xiaoyu Chen*, **Jiachen Hu***, Chi Jin, Lihong Li, Liwei Wang
- ICLR 2022 (**spotlight, 6% 176/3328**)

Near-Optimal Reward-Free Exploration for Linear Mixture MDPs with Plug-in Solver

- Xiaoyu Chen, **Jiachen Hu**, Lin F. Yang, Liwei Wang
- ICLR 2022 (**spotlight, 6% 176/3328**)

Near-optimal Representation Learning for Linear Bandits and Linear RL

- **Jiachen Hu***, Xiaoyu Chen*, Chi Jin, Lihong Li, Liwei Wang
- ICML 2021

Efficient Reinforcement Learning in Factored MDPs with Application to Constrained RL

- Xiaoyu Chen, **Jiachen Hu**, Lihong Li, Liwei Wang
- ICLR 2021

Distributed Bandit Learning: Near-Optimal Regret with Efficient Communication

- Yuanhao Wang*, **Jiachen Hu***, Xiaoyu Chen, Liwei Wang
- ICLR 2020

INTERNSHIPS

Algorithm developer at Wizard Quant **2024.7 - 2024.9**
Working as a machine learning algorithm developer for quant trading

Visiting scholar at Princeton University **2023.6 - 2023.11**
Host: Chi Jin

Summer intern at UIUC **2019.7 - 2019.8**
Host: Nan Jiang

SERVICES

Conference reviewer: ICLR (2021,2022,2023), ICML (2022, 2025), NeurIPS (2021,2023), AAAI (2025)

Journal reviewer: AURO (IF = 4.8)

AWARDS

Yuehua Luo Scholarship (2024)

Excellent Academic Research Reward (2022)

Peking University President Scholarship (2020-2021)

Turing Class Summit Scholarship (2019)

Jane Street Electronic Trading Challenge 2018 in Tsinghua University, 3rd Place

Wu Si Scholarship (2018)

Merit Student of Peking University (top 5%, 2017)

Robin Li Scholarship (top 1%, 2017)

Third-class annually New Students Scholarship (2016)

National Olympic of informatics (NOI) 2015, Sliver Medal