

# HONGRU WANG

Emails: hongru.carrywang@gmail.com / hrwise98@gmail.com

[Homepage](#) | [Ph.D. Thesis](#) | [Thesis Slides](#) | [Google Scholar](#)

## RESEARCH STATEMENT

---

My research builds the Theory of Agent (ToA) — a unified framework that treats internal reasoning and external acting as two epistemically alternative knowledge-acquisition mechanisms, modeling the parametric internal world and the physical external world respectively. My long-term goal is to resolve the impossible triangle among safety, personalization, and autonomy in language agents that learn from both internal and external interactions.

## EDUCATION

---

<b>University of Edinburgh</b> Postdoc (Research Associate) supervised by Prof. Amos Storkey and Prof. Jeff Z. Pan Self-evolving Agents, Long-horizon Reasoning	<b>2025/09 - Now</b>
<b>The Chinese University of Hong Kong</b> Ph.D. supervised by Prof. Kam-Fai Wong (ACL Fellow) Dialogue System, Large Language Models, Tool Learning Thesis Committee: Prof. Yu Jeffrey Xu, Prof. WAI Hoi To, Prof. Zhiyuan Liu (Tsinghua)	<b>2021/09 - 2025/09</b> GPA: 3.82/4.00
<b>The Chinese University of Hong Kong</b> M.Eng. Computer Science and Engineering Data Mining, Big Data Technology, Machine Learning	<b>2019/09 - 2020/11</b> GPA: 3.60/4.00
<b>Communication University of China</b> B.Eng. Computer Science and Technology Computer Operation System, Computer Networks, Data Structures and Algorithms	<b>2015/09 - 2019/09</b> GPA: 3.27/4.00

## EXPERIENCE

---

Visiting Student <b>BlenderLab</b> (Mentor: Prof. Heng Ji (ACL Fellow))	University of Illinois Urbana-Champaign	<b>2024/12 - 2025/6</b>
Research Intern <b>ByteDance (SZ)</b> (Mentor: Dr. Deng Cai, Dr. Wanjun Zhong)	Seed-LLM-Horizon (Doubao)	<b>2024/08 - 2024/12</b>
Visiting Student <b>EdinburghNLP</b> (Mentor: Prof. Jeff Z. Pan)	University of Edinburgh	<b>2024/04 - 2024/11</b>
Research Assistant <b>CUHK</b> (Mentor: Prof. Kam-Fai Wong)	MoE Key Lab of High Confidence Software Technologies	<b>2019/12 - 2021/07</b>

## (CO-)FIRST/CORRESPONDING AUTHOR CONFERENCES

---

20. Position: Agent Should Invoke External Tools ONLY When Epistemically Necessary  
Hongru Wang, Cheng Qian, Manling Li, et al., Mengdi Wang, Heng Ji, Amos Storkey, Kam-Fai Wong  
**ICML 2026 Position Track** Theory of Agent  
**A unified theoretical framework for agents to guide training and evaluation**
19. Search-R2: Enhancing Search-Integrated Reasoning via Actor-Refiner Collaboration  
Bowe He, Minda Hu, Zenan Xu, Hongru Wang<sup>†</sup>, et al., Pluto Zhou, Irwin King

ICML 2026 Agent † denotes corresponding author

w/ Tencent-Hunyuan Team

18. From Word to World: Can LLM be Implicit Text-based World Model?  
Yixia Li, Hongru Wang†, Jiahao Qiu, Zhenfei Yin, et al., Guanhua Chen, Heng Ji  
ACL 2026 LLM World Model  
w/ UIUC and MindLab which is a famous start-up on RL infra
17. Rethinking the Role of Entropy in Optimizing Tool-Use Behaviors for Large Language Model Agents  
Zeping Li, Hongru Wang†, Yiwen Zhao, Guanhua Chen, Yixia Li, et al., Zhenfei Yin  
ACL 2026 LLM Tool Learning  
w/ Oxford and Fudan
16. Self-Sum: Teaching Agent Itself to Decide When and What to Summarize  
Hongru Wang, Rui Wang, Jushi Kai, et al., Xiaoteng Ma, Jeff Z. Pan, Amos Storkey  
Findings of ACL 2026 LLM Tool Learning  
w/ EdinburghNLP
15. SafeToolBench: Pioneering a Prospective Benchmark to Evaluating Tool Utilization Safety in LLMs  
Hongfei Xia\*, Hongru Wang\*, Zeming Liu, Qian Yu, Yuhang Guo, and Haifeng Wang  
Findings of EMNLP 2025 LLM Tool Learning \* denotes equal contribution
14. UAlign: Leveraging Uncertainty Estimations for Factuality Alignment on Large Language Models  
Boyang XUE, Fei Mi, Qi Zhu, Hongru Wang†, Rui Wang, et al., Kam-Fai Wong  
ACL 2025 LLM
13. Self-Reasoning Language Models: Unfold Hidden Reasoning Chains with Few Reasoning Catalyst  
Hongru Wang, Deng Cai, Wanjun Zhong, Shijue Huang, Jeff Z. Pan, Zeming Liu, Kam-Fai Wong  
Reasoning and Planning for Large Language Models of ICLR 2025  
Findings of ACL 2025 LLM (data synthesis to self-improve itself for LLMs)  
work during internship at ByteDance Seed
12. Rethinking Stateful Tool Use in Multi-Turn Dialogues: Benchmarks and Challenges  
Hongru Wang, Wenyu Huang, et al., Zeming Liu, Jeff Z. Pan, Kam-Fai Wong  
Findings of ACL 2025 LLM Tool Learning
11. MlingConf: A Comprehensive Study of Multilingual Confidence Estimation on LLMs  
Boyang XUE\*, Hongru Wang\*, Rui Wang, et al., Wenxuan Zhang, Kam-Fai Wong  
Findings of ACL 2025 LLM
10. ToolSpectrum: Towards Personalized Tool Utilization for Large Language Models  
Zihao Cheng\*, Hongru Wang\*, Zeming Liu, Yuhang Guo, et al., Yunhong Wang, Haifeng Wang  
Findings of ACL 2025 LLM Tool Learning  
w/ Baidu
9. TransBench: Breaking Barriers for Transferable GUI Agents in Dynamic Digital Environments  
Yuheng Lu\*, Qian Yu\*, Hongru Wang\*, Zeming Liu, Wei Su, et al., Yunhong Wang, Haifeng Wang  
Findings of ACL 2025 LLM Agent  
w/ Baidu
8. Self-DC: When to Reason and When to Act? Self Divide-and-Conquer for Compositional ..  
Hongru Wang, Boyang Xue, Baohang Zhou, et al., Kam-Fai Wong  
NAACL 2025 LLM , Tool Learning Oral Blog
7. AppBench: Planning of Multiple APIs from Various APPs for Complex User Instruction  
Hongru Wang, Rui Wang, Boyang XUE, et al., Jeff Z. Pan, Kam-Fai Wong

EMNLP 2024 LLM , Tool Learning (Apple Intelligence)

w/ EdinburghNLP

6. Enhancing Large Language Models Against Inductive Instructions with Dual-critique Prompting  
Rui Wang\*, Hongru Wang\*, Fei Mi, Boyang XUE, Yi Chen, Kam-Fai Wong, Ruifeng Xu  
NAACL 2024 LLM , Safety
5. UniRetriever: Multi-task Candidates Selection for Various Context-Adaptive Conversational Retrieval  
Hongru Wang, Boyang Xue, Baohang Zhou, et al., Kam-Fai Wong  
LREC-COLING 2024 Dialogue System , Tool Learning
4. M3Sum: A Novel Unsupervised Language-guided Video Summarization  
Hongru Wang, Baohang Zhou, Zhengkun Zhang, Yiming Du, David Ho, Kam-Fai Wong  
ICASSP 2024 LLM , MM
3. Large Language Models as Source Planner for Personalized Knowledge-grounded Dialogue  
Hongru Wang, Minda Hu, Yang Deng, et al., Irwin King, Kam-Fai Wong  
Findings of EMNLP 2023 Dialogue System , LLM **Best Paper Award 😊 @Doctoral Forum**
2. Cue-CoT: Chain-of-thought Prompting for Responding to In-depth Dialogue Questions with LLMs  
Hongru Wang\*, Rui Wang\*, Fei Mi, et al., Ruifeng Xu and Kam-Fai Wong  
Findings of EMNLP 2023 Dialogue System , LLM [Blog](#)  
**120+ citations, followed by Google DeepMind, ...**
1. Integrating Pretrained Language Model for Dialogue Policy Learning  
Hongru Wang, Huimin Wang, Zezhong Wang and Kam-Fai Wong  
ICASSP 2022 Dialogue System , RLHF

#### (CO-)FIRST/CORRESPONDING AUTHOR JOURNAL

---

4. AdaCtrl: Towards Adaptive and Controllable Reasoning via Difficulty-Aware Budgeting  
Shijue Huang\*, Hongru Wang\*, Wanjun Zhong\*, et al., Bowen Cao, Yi R. (May) Fung  
Transactions on Machine Learning Research 2026  
**Teaching model to reason efficiently**
3. A Survey of Self-Evolving Agents: What, When, How, and Where to Evolve on the Path to Artificial Super Intelligence  
Huan-ang Gao et al., Hongru Wang\*†, et al., Qingyun Wu, Heng Ji, Mengdi Wang†  
Transactions on Machine Learning Research 2026 Self-Evolving Agents  
w/ Princeton and lots of top universities at the world  
**First comprehensive survey on self-evolving agents; served as one of two corresponding authors.**
2. KddRES: A Multi-level Knowledge-driven Dialogue .. Towards Customized Dialogue System  
Hongru Wang, Wai-Chung Kwan, Min Li, Zimo Zhou and Kam-Fai Wong  
Computer Speech and Language 2024 (Tsinghua B, JCR Q2, IF: 4.3) Dialogue System  
**First Cantonese-based task-oriented dialogue system**
1. A Survey on Recent Advances and Challenges in Reinforcement Learning Methods ..  
Wai-Chung Kwan\*, Hongru Wang\*, Huimin Wang and Kam-Fai Wong  
Machine Intelligence Research 2023 (JCR Q1, IF: 6.4) Dialogue System , RLAIIF

#### NON-FIRST AUTHOR CONFERENCES/JOURNALS

---

15. On Path to Multimodal Historical Reasoning: HistBench and HistAgent  
Jiahao Qiu, et al., Hongru Wang, Kaixuan Huang, Xun Jiang, Xi Gao, Mengdi Wang

ICML 2026 Multimodal Reasoning

w/ Princeton AI Lab

14. WebAggregator: Enhancing Compositional Reasoning Capabilities of Deep Research Agent Foundation Models  
Rui Wang, Ce Zhang, Jun-Yu Ma, Jianshu Zhang, Hongru Wang, et al., Dong Yu, Kam-Fai Wong  
ACL 2026 Deep Research Agent  
w/ Tencent AI Lab
13. Perception-aware Policy Optimization for Multimodal Reasoning  
Zhenhailong Wang, et al., Hongru Wang, et al., Fei Huang, Heng Ji  
ICLR 2026 Multimodal Reasoning  
w/ Alibaba Tongyi Lab
12. RM-R1: Reward Modeling as Reasoning  
Xiusi Chen, et al., Hongru Wang, et al., Tong Zhang, Hanghang Tong, Heng Ji  
ICLR 2026 Reward modeling as reasoning [Blog](#)
11. ToolRL: Reward is All Tool Learning Needs  
Cheng Qian, et al, Hongru Wang, Xiusi Chen, Dilek Hakkani-Tür, Gokhan Tur, Heng Ji  
NeurIPS 2025 A comprehensive analysis of reward design for tool learning [Blog](#)  
200+ citations, followed by Google, Nvidia, Meta, ....
10. The Landscape of Agentic Reinforcement Learning for LLMs: A Survey  
Guibin Zhang, et al., Hongru Wang, et al., Jun Wang, Shuicheng YAN, Philip Torr, Lei Bai  
Transactions on Machine Learning Research 2026 Agentic Reinforcement Learning
9. SMART: Self-Aware Agent for Tool Overuse Mitigation  
Cheng Qian, Emre Can Acikgoz, Hongru Wang#, et al., Dilek Hakkani-Tür, Gokhan Tur, Heng Ji  
Findings of ACL 2025 LLM Tool Learning # denotes mentorship  
w/ UIUC
8. Steering Knowledge Selection Behaviours in LLMs via SAE-Based Representation Engineering  
Yu Zhao, et al., Hongru Wang, Xuanli He, Kam-Fai Wong, Pasquale Minervini  
NAACL 2025 LLM (tool / knowledge conflicts, detections and control)  
Oral, w/ EdinburghNLP
7. AutoPSV: Automated Process-Supervised Verifier  
Jianqiao Lu, Zhiyang Dou, Hongru Wang, et al., Zhijiang Guo  
NeurIPS 2024 LLM (Meta-reasoning theory)
6. Knowledge Conflicts for LLMs: A Survey  
Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, Wei Xu  
EMNLP 2024 LLM [Blog](#)  
w/ Tsinghua and Westlake University
5. Less is More: Making Smaller Language Models .. Subgraph Retrievers for Multi-hop KGQA  
Wenyu Huang, Guancheng Zhou, Hongru Wang, et al., Mirella Lapata, Jeff Z. Pan  
Findings of EMNLP 2024 LLM (Differentiable Search Index (DSI))  
w/ EdinburghNLP
4. Enhancing Biomedical Knowledge RAG with Self-Rewarding Tree Search and PPO  
Minda Hu, Licheng Zong, Hongru Wang, et al., Kam-Fai Wong, Yu Li, Irwin King  
Findings of EMNLP 2024 LLM
3. REGA: Role Prompting Guided Multi-Domain Adaptation for Large Language Models

- Rui Wang, Fei Mi, Yi Chen, Boyang XUE, Hongru Wang, Qi Zhu, Kam-Fai Wong, Ruifeng Xu  
Findings of **NAACL 2024 LLM** , **Role Playing**
2. SELF-GUARD: Empower the LLM to Safeguard Itself  
Zezhong Wang, Fangkai Yang, Lu Wang, Pu Zhao, Hongru Wang, et al., Kam-Fai Wong  
**NAACL 2024 LLM** , **Safety**
  1. Retrieval-free Knowledge Injection through Multi-Document Traversal for Dialogue Models  
Rui Wang, Jianzhu Bao, Fei Mi, Yi Chen, Hongru Wang, et al., Kam-Fai Wong, Ruifeng Xu  
**ACL 2023 Dialogue System**

## WORKSHOP, TUTORIAL AND OTHERS

---

6. Lifelong Agents: Learning, Aligning, Evolving  
Cheng Qian, Emre Can Acikgoz, Hongru Wang, Zhenfei Yin, Manling Li, Yun-Nung Chen, Mengdi Wang, Caiming Xiong  
**First Workshop on Lifelong Agents at ICLR 2026 Lifelong Agents**  
**Workshop Organizer and Program Chair**  
**w/ Princeton, UIUC, Northeastern University, Salesforce Research, ...**
5. Acting Less is More: Learning to Act Efficiently with Large Language Models  
Hongru Wang, Cheng Qian, et al., Bowen Jin, Mengdi Wang, Kam-Fai Wong, Heng Ji  
**LAW@NeurIPS 2025 LLM** , **Tool Learning**  
**w/ Princeton, UIUC**
4. Physics Supernova: AI Agent Matches Elite Gold Medalists at IPhO 2025  
Jiahao Qiu, Jingzhe Shi, Xinzhe Juan, et al., Hongru Wang, Sanfeng Wu, Mengdi Wang  
**LAW@NeurIPS 2025 LLM** , **Tool Learning**  
**AI4Science, w/ Princeton**
3. PerLTQA: A Personal Long-Term Memory Dataset for Memory Classification, Retrieval, and ..  
Yiming Du, Hongru Wang, Zhengyi Zhao, et al., Kam-Fai Wong  
SIGHAN Workshop of **ACL 2024**.  
**Best Paper Award 😊**
2. OSPC: Detecting Harmful Memes with Large Language Model as a Catalyst  
Jingtao Cao, Zheng Zhang, Hongru Wang, Bin Liang, Hao Wang, Kam-fai Wong  
Online Safety Prize Challenge, **WWW 2024**.  
**Champion Solution 😊**
1. Empowering Large Language Models: Tool Learning for Real-World Interaction  
Hongru Wang, Yujia Qin, Yankai Lin, Jeff Z. Pan, Kam-fai Wong  
Tutorial, **SIGIR 2024 LLM** , **Tool Learning**  
**w/ Tsinghua and RUC**

## GRANTS / FUNDING

---

3. Overseas Research Attachment Programme (50,000 HKD, ORAP 2023-2024)
2. A Knowledge Graph Based Dynamic Video Extractive Summarization System (PRP/054/21FX)
1. "SEVES: Semantic-driven Effective Video Extractive Summarization System" – Technology and Business Development Fund (200,000 HKD, TBF22ENG004)

## TEACHING ASSISTANT

---

SEEM 3450 Engineering Innovation and Entrepreneurship  
SEEM 3490 Information Systems Management  
SEEM 5730 / ECLT 5910 Information Technology Management

## COMPETITIONS & AWARDS

---

Top 1 at Online Safety Prize Challenge, WWW 2024	international
Reaching Out Awards (2022-2023)	school
Distinguished Academic Performance Scholarship (2019-2020), CUHK CSE	school
Meritorious Winner, Mathematical Contest In Modeling (2018)	international
A software copyright of "WeCampus WeChat mini-program" (2018SR562540)	national
China National Radio Scholarship	school

## COMMUNITY CONTRIBUTION

---

Program Chair: Lifelong Agents Workshop@ICLR 2026  
Area Chair: NeurIPS 2025, 2026  
Volunteer: NAACL 2025  
Reviewer: ICLR, ICML, ACL/EMNLP/NAACL, AAAI, IJCAI, ARR, TNNLS, ... (2022-present)  
Organizer: [ToolsMeetLLM@SIGIR2024](#)  
Co-founder of 2nd largest NLP community in China: [NICE](#) (150,000+ followers)

## INVITED TALKS

---

2026: EvoAgentX, MindLab, PolyU-NLP  
2025: UIUC-NLP, HKU-NLP, UCL, University of Manchester, Shanghai AI Lab, Fudan University  
2024: EdinburghNLP, AI Time, TechBeat  
2023: MLNLP, OpenMMLab, NILLI Workshop  
More details can be found in [Invited Talks](#)

## REFERENCES

---

Prof. Kam-Fai Wong, ACL Fellow, Professor, The Chinese University of Hong Kong  
Prof. Heng Ji, ACL Fellow, Professor, University of Illinois Urbana-Champaign  
Prof. Jeff Z. Pan, Turing Fellow, Professor, University of Edinburgh  
Prof. Amos Storkey, ELLIS Fellow, Professor, University of Edinburgh  
Prof. Irwin King, ACM/IEEE/AAAI Fellow, Professor, The Chinese University of Hong Kong