Feng Chen 陈烽

Ph.D. Student, College of Computing and Data Science, Nanyang Technological University
(+86) 150-5235-8633 • chenfeng20010105@gmail.com • feng012@e.ntu.edu.sg • Singapore
Homepage: chenf-ai.github.io

PROFILE

Ph.D. student at Nanyang Technological University (advised by Prof. Bo An), working on large language model post-training and agent-harness design. Completed both M.Sc. and B.S. at the School of Artificial Intelligence, Nanjing University — a member of its **first undergraduate cohort**, **ranked 1st in the major** (postgraduate-recommendation GPA 4.74/5.00). Research spans reinforcement learning, LLM reasoning, and multi-agent systems, with **20+ publications** at venues including NeurIPS, ICLR, IJCAI, TMLR, and IEEE TNNLS and **480+ citations (h-index 11**, Google Scholar). On the industry side, drove LLM post-training research at ByteDance's Seed team — contributing to the **Seed1.5-Thinking** reasoning model — and now builds agent harnesses at Apodex. Fluent in English (**IELTS 8.0**).

EDUCATION

Nanyang Technological University **Singapore**
Ph.D. Student, **College of Computing and Data Science** **Aug 2025 ~ Present**
- Supervisor: Prof. Bo An
- Research: large language model post-training and agent-harness design

Nanjing University **Nanjing, China**
M.Sc. in School of **Artificial Intelligence** **Sep 2022 ~ Jun 2025**
- Supervisor: Prof. Zong-Zhang Zhang
- Research interest: Reinforcement Learning, Large Language Models, Multi-Agent Systems
- Admitted without entrance examination

B.S. in School of **Artificial Intelligence** **Sep 2018 ~ Jun 2022**
- B.S. in Computer Science and Technology; **ranked 1st in the major** (postgraduate-recommendation GPA **4.74**/5.00)

INDUSTRIAL EXPERIENCE

Research Intern, Agent Harness at **Apodex** **Mar 2026 ~ Present**
- Building the agent harness behind Apodex's self-evolving heavy-duty solver — the orchestration and verification layer that coordinates large teams of sub-agents across long-horizon research and problem-solving tasks.

LLM Research Intern, Post-Training at **ByteDance (Seed)** **Sep 2024 ~ Jun 2025**
- Drove algorithmic research in the post-training stage of large language models — reinforcement learning for reasoning, optimization, and alignment — and contributed to Seed1.5-Thinking, ByteDance Seed's flagship reasoning model.

Reinforcement Learning Research Intern at **Polixir** **Jul 2022 ~ Mar 2023**
- Built an adversarial reinforcement learning framework enabling two agents to co-train against each other in adversarial environments.

Reinforcement Learning Research Intern, Game AI at **ByteDance** **Jul 2021 ~ Oct 2021**
- Researched training methodologies for game-playing agents and algorithms for automated game-scene generation.

ACADEMIC SERVICE & INVITED TALKS

- **Co-organizer**, ICLR 2026 Workshop on Multi-Agent Learning and Its Opportunities in the Era of Generative AI.
- **Invited Speaker**, Workshop on Multi-Agent Reinforcement Learning, DAI 2024 (Distributed Artificial Intelligence Conference).
- **Reviewer** for AAAI, NeurIPS, ICLR, ICML, and TMLR.

PUBLICATIONS

20+ papers · 480+ citations · h-index 11 (Google Scholar).

First Author / Co-first Author

- **Efficient Multi-Agent Cooperation Learning through Teammate Lookahead.** Transactions on Machine Learning Research (TMLR), 2025. *Agents look ahead to teammates' future policies to overcome "teammate delay" and strengthen cooperative learning.*
- **Efficient Multi-Agent Communication via Self-supervised Information Aggregation.** Advances in Neural Information Processing Systems (NeurIPS), 2022. *Disentangles information aggregation from extraction using two self-supervised objectives for efficient communication.*
- **Stable Continual Reinforcement Learning via Diffusion-based Trajectory Replay.** ICLR Workshop GenAI4DM, 2024. *Replays high-return past trajectories with a diffusion model to alleviate catastrophic forgetting in continual RL.*
- **Multi-Agent Policy Transfer via Task Relationship Modeling.** Science China Information Sciences (SCIS), 2024. *Handles variable-length inputs and learns effect-based task representations to transfer policies across tasks.*
- **Communication-Robust Multi-Agent Learning by Adaptable Auxiliary Multi-Agent Adversary Generation.** Frontiers of Computer Science (FCS), 2024. *An adversarial framework for learning robust, effective multi-agent communication policies.*

Second Author

- **Efficient Communication via Self-supervised Information Aggregation for Online and Offline Multi-agent Reinforcement Learning.** IEEE Transactions on Neural Networks and Learning Systems (TNNLS), 2025.
- **One by One, Continual Coordinating with Humans via Hyper-Teammate Identification.** Transactions on Machine Learning Research (TMLR), 2024.

Other Publications

- **Seed1.5-Thinking: Advancing Superb Reasoning Models with Reinforcement Learning.** ByteDance Seed Team, 2025. *A large-scale RL-trained reasoning model reaching state-of-the-art results across math, code, and science benchmarks.*
- **Apodex 1.0: A Self-Evolving Heavy-Duty Solver.** Apodex Team, 2026. *A multi-agent solver scaling to up to 150 sub-agents with independent step-level verification; SOTA on BrowseComp, FrontierScience, and SuperChem.*
- **Learning to Coordinate with Anyone.** Distributed Artificial Intelligence (DAI), 2023. **— Best Paper Award**
- **EvaLearn: Quantifying the Learning Capability and Efficiency of LLMs via Sequential Problem Solving.** Advances in Neural Information Processing Systems (NeurIPS), 2025.
- **Lost in the Context: Insufficient and Distracted Attention to Contexts in Preference Modeling.** Annual Meeting of the Association for Computational Linguistics (ACL), 2025.
- **AgentOCR: Reimagining Agent History via Optical Self-Compression.** Annual Meeting of the Association for Computational Linguistics (ACL), 2026.

PROJECT EXPERIENCE

Cooperative Multi-Agent Deep Reinforcement Learning Based on Knowledge Transfer **Jan 2023 ~ Dec 2026**

NSFC General Program (面上项目), No. 62276126 · PI: Prof. Zong-Zhang Zhang, Nanjing University

- Core student member of this national research program on cooperative multi-agent reinforcement learning; developed knowledge- and policy-transfer methods that generalize cooperation and communication across tasks and teammates, yielding first-/co-first-author publications at NeurIPS, TMLR, and Science China Information Sciences.

Multi-AGV Path Finding Project with **Hikvision** **Nov 2021 ~ Nov 2022**

Advisor: Prof. Zong-Zhang Zhang, School of Artificial Intelligence, Nanjing University

- Designed an efficient, safe reinforcement learning algorithm for multi-agent path finding (MAPF) of automated guided vehicles in warehouse logistics, outperforming baselines with scalable distributed training; published as a first-author student abstract at AAAI 2023.

HONORS & AWARDS

- National Scholarship, China (2020)
- Outstanding Graduate Student, Nanjing University (2022)
- Provincial Merit Student, Jiangsu Province (2020)
- Outstanding Student, Nanjing University (2019)