

A Concise Introduction to LLM-based Multi-agent Systems

Yang Chen
Shanghai AI Lab
Shanghai, China
chenyang4@pjlab.org.cn

Shuyue Hu
Shanghai AI Lab
Shanghai, China
hushuyue@pjlab.org.cn

ABSTRACT

Large language models (LLMs) are increasingly used as core components of multi-agent systems, enabling new forms of interaction, coordination, and collective decision-making. This tutorial provides a concise, agent-centric introduction to LLM-based multi-agent systems (MAS), focusing on the interplay between LLMs and classical multi-agent principles. It highlights how multi-agent mechanisms—such as debate, role specialization, coordination protocols, and strategic interaction—can address key limitations of single LLM agents, including brittle reasoning and lack of robustness, as well as how LLMs can serve as powerful infrastructure for communication and coordination in multi-agent settings. Emphasis is placed on connections to core AAMAS concepts such as decision-making, incentives, interaction dynamics, and emergent behavior, offering a unified perspective on this rapidly evolving research area.

KEYWORDS

Multi-agent systems, Large language models

ACM Reference Format:

Yang Chen and Shuyue Hu. 2026. A Concise Introduction to LLM-based Multi-agent Systems. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages.

1 DETAILED TUTORIAL OUTLINE

Preferred length: Half-day tutorial (3.5 hours of technical content)

Part I: Foundations of LLMs and LLM-based Agents (1 hour)

- Agent abstraction: policy, memory, interaction loop
- Single-agent LLMs as decision-making entities
- Roles of LLMs in agents: priors, planners, or reasoning modules
- From single-agent to multi-agent systems

Break (10 minutes)

Part II: Multi-Agent Systems for LLMs (1 hour)

- Limitations of single LLM agents: hallucination, myopia, bias, and brittle reasoning
- Multi-agent debate, deliberation, and critique to improve reasoning reliability
- Redundancy, voting, and consensus mechanisms for robustness

Break (10 minutes)

Part III: LLMs for Multi-Agent Systems (1 hour)

- Advantages of LLMs as infrastructure for MAS

- Language-mediated coordination and communication
- LLMs as coordination and communication layers for MAS
- Connections to mechanism design, learning in games, and collective decision-making

Break (10 minutes)

Part IV: Open Challenges and Future Directions (30 minutes)

- Evaluation of interactive, deliberative, and strategic agent systems
- Scalability in agent number, interaction length, and reasoning depth
- Efficient and effective communications in LLM-based agents
- Open theoretical questions at the intersection of MAS and LLMs

See Fig. 1 for an illustration of the outline in a clearer structure.

2 TARGET AUDIENCE AND PREREQUISITES

The tutorial targets researchers, practitioners, and advanced graduate students in multi-agent systems, machine learning, and artificial intelligence. It is particularly relevant to attendees interested in learning, decision-making, and interaction in agent-based systems.

Prerequisite knowledge:

- Basic familiarity with agents and multi-agent systems
- Introductory knowledge of machine learning or reinforcement learning. No prior experience with large language models is required.

3 RELEVANCE TO THE AAMAS COMMUNITY

LLM-based MAS constitute an emerging research direction that aligns closely with core AAMAS themes, including agent architectures, interaction protocols, coordination, strategic behavior, and incentive design. While LLMs have primarily been studied in NLP and systems venues, their deployment in multi-agent settings introduces new challenges in decision-making under interaction, robustness, emergent behavior, and evaluation of collective outcomes, all of which are central to the AAMAS community. This tutorial bridges recent advances in LLM-based agents with established multi-agent principles, offering a unified perspective that enables AAMAS researchers to analyze, design, and reason about LLM-based MAS using familiar conceptual tools, and, in turn, to understand how these MAS concepts help expand the capability boundary of LLMs.

4 RESUME OF THE PRESENTERS

Presenter Yang Chen

- **Postal Address:** Floor 10, No. 129 Longwen Road, Xuhui District, Shanghai, China
- **Phone Numbers:** +86 18511822753 / +64 0225250016

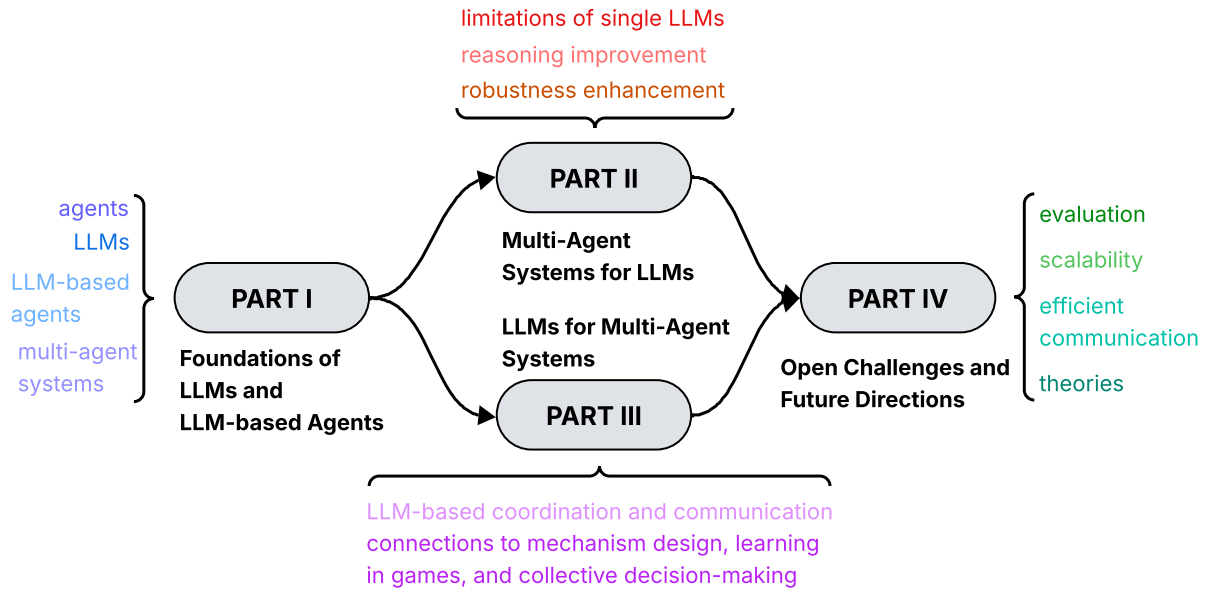


Figure 1: Illustration of the tutorial outline.

- **Email Address:** chenyang4@pjlab.org.cn
- **Website:** <https://www.yangchen.info/>
- **Background:** Yang Chen has expertise in reinforcement learning and multi-agent systems. Recently, his research interest lies in the reinforcement learning for LLMs, in particular for LLM reward modelling and training LLM-based multi-agent systems.
- **Example of Work:** Yang Chen co-authored the tutorial paper *Hands-on LLM-based Agents: A Tutorial for General Audiences* [10], which aligns with the topics of this tutorial. From the perspective of a reader, it serves as a tutorial that systematically introduces the LLM-based agents to the elementary-level audiences.
- **Teaching Experience:** While he was working in The University of Auckland, Yang Chen was the instructor of the post-graduate courses COMPSCI 713 Artificial Intelligence Fundamentals and COMPSCI 761 Advanced Topics in Artificial Intelligence during the period from 2022 to 2024. He participated in designing the course syllabuses, tests and examinations. The slides and examination papers designed by Yang Chen are publicly available at <https://www.yangchen.info/teaching/>.
- **Scholarship:** Yang Chen’s scholarship focuses on LLM-based agents in multi-agent systems, emphasizing decision-making, reward modeling, and strategic interaction. His research develops principled methods for reward inference and optimization to enable stable and interpretable agent behavior in long-horizon, partially observable environments. By bridging (inverse) reinforcement learning, imitation learning, and LLMs, his work addresses incentive alignment and

adaptation in collaborative and competitive settings. Overall, his research advances the theoretical and algorithmic foundations of reliable LLM-based multi-agent intelligence. He has a series of publications on agent(s) incentive/reward design [3–6] and multi-agent reinforcement learning [9, 17] at internationally prestigious venues such as NeurIPS, AAAI, AAMAS, and IJCAI. In parallel, he has also publications on language model reasoning [1] and question answering [7, 8] at flagship natural language processing conferences such as ACL and COLING. He is an active contributor to AAMAS community and has served as the local co-chair of AAMAS 2024 and co-chair for the its AAAI Track. He also serves as regular PC members for leading AI conferences.

• Previous Tutorials:

- *From One to Infinity: New Perspectives and Methods for Inverse Reinforcement Learning*. University of Copenhagen. Two hours. Copenhagen, Denmark. 8 October 2024.
- *Identify and Mitigate Bias-Perpetuating Behaviours: Decoding Animal Behaviour via Many-agent IRL*. AAMAS 2024 Affinity Group Event on Diversity in AI. One hour. Auckland, New Zealand. 07 May 2024.
- *Mean Field Game as a Framework for Many-agent Inverse Reinforcement Learning*. ML and MFG seminar. 1.5 hours. Online. 06 Dec 2022.

Presenter Shuyue Hu

- **Postal Address:** Floor 10, No. 129 Longwen Road, Xuhui District, Shanghai, China
- **Phone Numbers:** +86 15919338791
- **Email Address:** hushuyue@pjlab.org.cn
- **Personal Website:** <https://shuyuehu.github.io>

- **Google Scholar Homepage:** https://scholar.google.com.hk/citations?user=JLdYe_IAAAA&hl=zh-CN
- **Background:** Shuyue Hu is a research scientist at Shanghai Artificial Intelligence Laboratory and an Adjunct Ph.D. Advisor at Shanghai Jiao Tong University. Her research broadly focuses on MAS, LLMs, and game theory, with recent efforts aimed at enhancing LLM capabilities via multi-agent and collective intelligence, and applying LLMs to tackle fundamental challenges in MAS.
- **Example of Work:** Shuyue Hu co-authored the tutorial paper *Hands-on LLM-based Agents: A Tutorial for General Audiences* [10]. Moreover, her recent work spans multiple topics central to this tutorial, including multi-agent debate [2, 16], multi-LLM routing [14, 18, 19], emergent behaviors in LLM-based multi-agent systems [11, 12, 15], and multi-agent reinforcement learning for LLMs [13].
- **Teaching Experience:** Since 2022, Shuyue Hu has been a guest lecturer at several universities in China, including Fudan University, Tianjin University, and Northwestern Polytechnical University, delivering half-day or one-day lectures each semester. Courses she taught include *Introduction to LLMs and LLM-based Agents*, *LLM-based Multi-Agent Systems*, and *Theoretical Foundations of Multi-Agent Learning in Games*. Prior to this, she served as a teaching assistant in the Department of Computer Science and Engineering at The Chinese University of Hong Kong from 2015 to 2019, where she was responsible for delivering weekly one-hour tutorial sessions.
- **Scholarship:** Shuyue Hu has published over 40 papers in leading venues, including AAMAS, AIJ, PNAS, IJCAI, AAAI, NeurIPS, and ICML. She is the recipient of the DAI 2025 Best Paper Award and the 2023 Shanghai Overseas Young Talent Program. She is an active contributor to the AAMAS community, serving as a regular PC member for the AAMAS conferences, a senior PC member in the area of generative and agentic AI for AAMAS 2026, an area chair for reinforcement learning at AAMAS 2024, and a scholarship chair for AAMAS 2024. In addition, she serves as a guest editor for the AAMAS journal special issue titled “When Foundation Models Meet Multi-Agent Systems” and for the Neurocomputing special issue titled “Neural Networks for Game-Theoretic Intelligent Decision Making”.
- **Previous Tutorials:**
 - *Decision Intelligence of Large Language Models*. The 1st Chinese Conference on Evolutionary Computation. December 2025.
 - *When Large Language Models meet Multi-Agent Systems*. The 3rd Chinese Conference on Multi-Agent Systems. June 2025.
 - *Toward Cooperative AI in the Foundation Model Era*. The 3rd International Conference on Evolutionary Game Theory and Artificial Intelligence. July 2025.

5 CORRESPONDING PRESENTER

Yang Chen.

Email: chenyang4@pjlab.org.cn / yang.chen.csphd@gmail.com

ACKNOWLEDGMENTS

This work was supported by a locally commissioned task from the Shanghai Municipal Government.

REFERENCES

- [1] Qiming Bao, Alex Yuxuan Peng, Zhenyun Deng, Wanjun Zhong, Gael Gendron, Timothy Pistotti, Neşet Tan, Nathan Young, Yang Chen, Yonghua Zhu, et al. 2024. Abstract Meaning Representation-based logic-driven data augmentation for logical reasoning. In *Findings of the Association for Computational Linguistics: ACL 2024*. 5914–5934.
- [2] Jianhao Chen, Zishuo Xun, Bocheng Zhou, Han Qi, Hangfan Zhang, Qiaosheng Zhang, Yang Chen, Wei Hu, Yuzhong Qu, Wanli Ouyang, et al. 2026. Do We Truly Need So Many Samples? Multi-LLM Repeated Sampling Efficiently Scales Test-Time Compute. *AAAI* (2026).
- [3] Yang Chen, Xiao Lin, Bo Yan, Libo Zhang, Jiamou Liu, Neset Özkan Tan, and Michael Witbrock. 2024. Meta-inverse reinforcement learning for mean field games via probabilistic context variables. In *AAAI*, Vol. 38. 11407–11415.
- [4] Yang Chen, Libo Zhang, Jiamou Liu, and Shuyue Hu. 2022. Individual-Level Inverse Reinforcement Learning for Mean Field Games. In *AAMAS*. 253–262.
- [5] Yang Chen, Libo Zhang, Jiamou Liu, and Michael Witbrock. 2023. Adversarial Inverse Reinforcement Learning for Mean Field Games. In *AAMAS*. 1088–1096.
- [6] Yang Chen, Menglin Zou, Jiaqi Zhang, Yitan Zhang, Junyi Yang, Gael Gendron, Libo Zhang, Jiamou Liu, and Michael J Witbrock. 2025. Trust Region Reward Optimization and Proximal Inverse Reward Optimization Algorithm. *NeurIPS* (2025).
- [7] Zhenyun Deng, Yonghua Zhu, Yang Chen, Qianqian Qi, Michael J Witbrock, and Patricia Riddle. 2022. Prompt-based conservation learning for multi-hop question answering. In *Proceedings of the 29th International Conference on Computational Linguistics*. 1791–1800.
- [8] Zhenyun Deng, Yonghua Zhu, Yang Chen, Michael Witbrock, and Patricia Riddle. 2022. Interpretable AMR-based question decomposition for multi-hop question answering. *arXiv preprint arXiv:2206.08486* (2022).
- [9] Gaël Gendron, Yang Chen, Mitchell Rogers, Yiping Liu, Mihailo Azhar, Shahrokh Heidari, David Arturo Soriano Valdez, Kobe Knowles, Padriac O’Leary, Simon Eyre, et al. 2024. Behaviour Modelling of Social Animals via Causal Structure Discovery and Graph Neural Networks. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 2276–2278.
- [10] Shuyue Hu, Siyue Ren, Yang Chen, Chunjiang Mu, Jinyi Liu, Zhiyao Cui, Yiqun Zhang, Hao Li, Dongzhan Zhou, Jia Xu, et al. 2025. Hands-on LLM-based Agents: A Tutorial for General Audiences. (2025).
- [11] Siyue Ren, Zhiyao Cui, Ruiqi Song, Zhen Wang, and Shuyue Hu. 2024. Emergence of social norms in generative agent societies: principles and architecture. *IJCAI* (2024).
- [12] Siyue Ren, Wanli Fu, Xinkun Zou, Chen Shen, Yi Cai, Chen Chu, Zhen Wang, and Shuyue Hu. 2026. Beyond the Tragedy of the Commons: Building A Reputation System for Generative Multi-agent Systems. *AAMAS* (2026).
- [13] Ziyu Wan, Yunxiang Li, Xiaoyu Wen, Yan Song, Hanjing Wang, Linyi Yang, Mark Schmidt, Jun Wang, Weinan Zhang, Shuyue Hu, et al. 2025. Rema: Learning to meta-think for llms with multi-agent reinforcement learning. *NeurIPS* (2025).
- [14] Chenxu Wang, Hao Li, Yiqun Zhang, Linyao Chen, Jianhao Chen, Ping Jian, Peng Ye, Qiaosheng Zhang, and Shuyue Hu. 2026. ICL-Router: In-Context Learned Model Representations for LLM Routing. *AAAI* (2026).
- [15] Ziyi Yang, Zaibin Zhang, Zirui Zheng, Yuxian Jiang, Ziyue Gan, Zhiyu Wang, Zijian Ling, Jinsong Chen, Martz Ma, Bowen Dong, et al. 2024. Oasis: Open agent social interaction simulations with one million agents. *arXiv preprint arXiv:2411.11581* (2024).
- [16] Hangfan Zhang, Zhiyao Cui, Xinrun Wang, Qiaosheng Zhang, Zhen Wang, Dinghao Wu, and Shuyue Hu. 2025. If multi-agent debate is the answer, what is the question. *arXiv preprint arXiv:2502.08788* (2025).
- [17] Libo Zhang, Yang Chen, Toru Takisaka, Bakh Khousainov, Michael Witbrock, and Jiamou Liu. 2023. Learning Density-Based Correlated Equilibria for Markov Games. In *AAMAS*. 652–660.
- [18] Yiqun Zhang, Hao Li, Jianhao Chen, Hangfan Zhang, Peng Ye, Lei Bai, and Shuyue Hu. 2025. Beyond gpt-5: Making llms cheaper and better via performance-efficiency optimized routing. *DAI* (2025).
- [19] Yiqun Zhang, Hao Li, Chenxu Wang, Linyao Chen, Qiaosheng Zhang, Peng Ye, Shi Feng, Daling Wang, Zhen Wang, Xinrun Wang, et al. 2026. The Avengers: A Simple Recipe for Uniting Smaller Language Models to Challenge Proprietary Giants. *AAAI* (2026).