Proceedings of the 2nd International Conference on Machine Learning and Automation
DOI: 10.54254/2755-2721/109/20241330

Large Language Model Applied in Multi-agent System—A Survey

Kaiwen Dong

Lancaster University Business School

kaiwendong2024@163.com

Abstract. The application of large language models (LLMs) in single-agent systems within complex environments has proven successful, prompting a growing interest in their use within multi-agent systems (MAS). Despite the impressive capabilities of LLMs, it remains unclear how they can be optimally integrated and utilized to empower agents in MAS. Understanding how to effectively leverage the advantages of LLMs to enhance agent performance is crucial. This survey provides a comprehensive overview of the application of LLMs in MAS, focusing on their impact on agent cooperation, reasoning, and adaptive abilities. Finally, we discuss future directions and open questions in this evolving field.

Keywords: Multi-agent system, large language model, reinforcement learning.

1. Introduction

The introduction of Large Language Models (LLMs) into complex single-agent environments, such as robotics, games, and embodied AI scenarios, has been remarkably successful. By leveraging the cooperation and unique skills of agents, multi-agent systems (MAS) enhance the performance of individual LLM agents [1][2][3][4][5]. These systems use the coordination and collaboration of multiple agents to tackle tasks that are too complex for a single agent to manage alone [6]. The strengths of LLMs enable MAS to address more challenging and practical issues. However, the integration of LLMs into MAS is still in its early stages.

Previously, training an agent within MAS was extremely difficult due to the dynamic and complex nature of the environment. Each agent in a multi-agent system has specific duties and abilities, working together to achieve shared goals. Many methods focus on using deep reinforcement learning to train these agents, but most fail to solve intrinsic MAS problems, such as cooperation [7], adaptability [8], and reasoning [9][10]. This paper discusses some of the current challenges faced by MAS and how LLMs can be applied to help address these issues.

We identify three key abilities that distinguish MAS agents from single-agent systems: adaptability, reasoning, and cooperation:

1.1. Adaptive Power

Agents trained in a MAS can adapt to different environments or other agents. MAS can be tailored to the specific requirements of various use cases.

DOI: 10.54254/2755-2721/109/20241330

1.2. Reasoning Power

Promoting rigorous reasoning in a subset of agents through recurring discussions or debates improves interim outcomes. By utilizing reasoning power, agents can adjust their goals or actions to achieve optimal results.

1.3. Cooperative Power

Cooperative power enables agents to allocate tasks more optimally and assist each other to enhance team performance. Different agents can collaborate more effectively and assign tasks more logically in a MAS, allowing the entire system to handle assigned tasks more accurately.

This paper provides a comprehensive investigation into the use and interpretation of LLMs in MAS. It offers a thorough examination of multi-agent systems, offering a survey of the existing works and highlighting the challenges and unresolved issues in the field.

The remaining sections of the paper will cover related work, background, algorithms, future work, and conclusions.

2. Related work

2.1. Adaptive agents

Adaptive agents have the ability to perform transfer learning [11]. In a multi-agent system (MAS) setting, there are two types of transfer learning: inter-agent transfer and intra-agent transfer. Inter-agent transfer focuses on receiving and reusing messages from other agents through communication, while intra-agent transfer focuses on transferring skills across different tasks or scenarios [12]. With the assistance of large language models (LLMs), communication becomes more feasible and efficient. Significant progress has been made in inter-agent transfer through the use of these agents [13][14].

2.2. Reasoning Agents

Previous studies have employed text-based test variations, such as the unexpected transfer task (also known as the Smarties Task) or the unexpected contents task (also referred to as the Maxi Task or Sally–Anne Test), to evaluate large language models' (LLMs) Theory of Mind (ToM) capabilities [15][16]. Findings indicate that leading LLMs can successfully pass over 90% of these tests. However, Ullman [17] found that LLMs struggle with making more sophisticated ToM inferences, particularly those involving communication or second-order beliefs.

2.3. Cooperative Agents

Traditional multi-agent systems predefine communication methods between agents, such as data sharing or state observations. The introduction of large language models enables more flexible, human-understandable communication [18]. As a result, some research has sought to enhance LLM-based multi-agent systems by allowing agents to engage in dialogues or debates, improving their decision-making capabilities [19][20]. This peer-communication technique has also been applied to financial tasks. However, these methods are not ideal for financial tasks with a singular goal of profit maximization, as they often lead to unclear optimization objectives and can incur unnecessary communication costs [21].

3. Background

LLM: Large language models (LLMs) excel in text completion and other Natural Language Processing (NLP) applications due to their training on vast text corpora [22][23]. Recent studies demonstrate their capabilities in reasoning[24][25] and action plan generation [26], particularly when prompt engineering techniques like chain-of-thought are applied. However, some researchers highlight the limitations of these models in creating feasible plans for interacting with real-world objects [27][28]. A growing trend is the application of LLMs to develop intelligent agents.

MAS: The subfield of distributed artificial intelligence known as multi-agent systems (MAS) has seen rapid growth due to its flexibility and intelligence in solving distributed problems [29]. Multi-agent systems can address problems that are challenging for a single agent or system to solve. Their intelligence is primarily reflected in organized, functional, and procedural behaviors, as well as in search algorithms and reinforcement learning. To accomplish tasks, multi-agent systems assign different roles to various agents and achieve complex tasks through agent cooperation. Compared to single-agent systems, there is less interaction with users during the task completion process.

4. Method

This chapter will cover the three primary LLM components of the current multi-agent system, which are adaptive, reasoning and cooperation, as well as any issues that have arisen and their corresponding solutions. We will also highlight any current functional shortcomings that can be perfected.

4.1. Adaptive Agents

Adaptability is a crucial element in multi-agent systems, ensuring that these systems can adjust to various simulation environments, including interactions with both new agents and novel scenarios.

Traditional approaches to solving adaptive problems have included several strategies. For example, Long [30] used an evolutionary population curriculum to address the scale-up problem, enabling agents to adapt to more complex, densely populated environments by leveraging previous experiences. Long [31] introduced the concept of a social gradient to tackle scenarios where agents must adapt to unknown counterparts. They employed score matching for offline learning and used the learned score function to guide agents toward an optimal goal.

Recently, there has been a growing body of work utilizing Large Language Models (LLMs) to address adaptive challenges. Nathalia [32] proposed that effective information exchange through communication is key to achieving self-adaptation. To support this, they integrated LLMs, such as GPT-based technologies, into multi-agent systems. Their approach is based on the MAPE-K model, which is well-known for its robust capabilities in monitoring, analyzing, planning, and executing system modifications in response to changing environments. They demonstrated their approach by implementing and evaluating a simple MAS-based application. This innovative framework significantly advances the state of the art in self-adaptive systems by leveraging LLM capabilities to enhance the adaptability of autonomous systems.

In another study, Meng [33] explored the leader-following consensus problem in heterogeneous multi-agent systems that are vulnerable to actuator and sensor attacks, specifically focusing on false data injection attacks. They introduced a novel adaptive controller designed to ensure consensus tracking with cooperative uniform ultimate boundedness, even in the simultaneous presence of sensor and actuator attacks. Since complete state information was unavailable, the adaptive dynamics were developed using compromised and uncompromised output information. The effectiveness of their approach was demonstrated through simulations, highlighting the potential for future research in resilient consensus of heterogeneous multi-agent systems under actuator dynamics and attacks.

Yuan [12] introduced another innovative framework called EVOAGENT. This framework combines the powerful language understanding capabilities of LLMs with evolutionary methods to generate multiple agents based on initialized agents. EVOAGENT represents a generic approach to multi-agent generation, offering new possibilities for creating adaptable agents.

Together, these papers expand the horizons for developing advanced adaptive agents, providing innovative frameworks and methodologies that enhance the adaptability and resilience of multi-agent systems in various dynamic environments.

4.2. Reasoning Agents

Reasoning ability is crucial for large language models (LLMs) and multi-agent systems for several key reasons:

DOI: 10.54254/2755-2721/109/20241330

Understanding Complex Contexts: Large language models require advanced reasoning capabilities to comprehend and interpret intricate textual information, allowing them to navigate and process complex contexts effectively.

Knowledge Application and Transfer: Reasoning enables models to infer rules from learned knowledge and apply them to new problems, thereby enhancing the system's generalization ability and adaptability across different scenarios.

Problem-Solving Ability: In both natural language processing tasks and multi-agent environments, reasoning is essential for breaking down tasks methodically and identifying solution paths, ultimately leading to successful problem resolution.

Li [34] highlighted the importance of reasoning in interactive team tasks where agents' mental states fluctuate with each interaction. In such scenarios, Theory of Mind (ToM) evaluations become necessary, as agents update their mental states based on both communication and observations. This complexity makes reasoning more challenging compared to the text-based, static tests used in earlier studies. Their findings, which include evaluations using the Sally-Anne test, reveal that LLMs' ToM capabilities remain limited, particularly in dynamic teamwork situations where belief states are constantly shifting, and communication is frequent.

Liang [35] proposed the MAD framework, where multiple agents engage in a "tit for tat" exchange of arguments, with a judge overseeing the debates to reach a final solution. This framework has shown strong performance by carefully considering each agent's perspective. However, challenges arise when different LLMs are used by the agents, as the judge may struggle to remain impartial due to variations in the configurations and scales of the LLMs, making it difficult to discern the underlying truth.

Current models can sometimes exhibit overconfidence or make improbable leaps in logic, particularly when the quality and accuracy of extracted natural language vary. To address this issue, Du [36] introduced a complementary strategy where multiple instances of a language model engage in rounds of discussion, sharing their answers and thought processes to reach a consensus. Their research demonstrates that this approach significantly improves mathematical and strategic reasoning across various tasks. Additionally, it reduces the likelihood of false positives and hallucinations, thereby enhancing the factual accuracy of the content generated by the models.

These studies illustrate that LLM-powered agents possess reasoning capabilities. However, this reasoning often relies on carefully crafted prompts tailored to specific tasks. As a result, it can be vulnerable to attacks and may falter when applied to different scenarios, highlighting the need for further advancements in generalizing reasoning abilities across diverse contexts.

4.3. Cooperative Agents

In a Multi-Agent System (MAS), cooperation is crucial. MAS consists of multiple independent, autonomous decision-making intelligent entities (i.e., agents) that need to collaborate to solve complex problems and achieve common goals. Through cooperation, agents can share information, work together to improve overall task efficiency, and avoid duplicating efforts.

To accelerate the transition from static, task-specific models to dynamic, agent-based systems capable of functioning effectively across various applications, Durante [7] introduced an Interactive Agent Foundation Model. This model trains AI agents on a variety of tasks, datasets, and domains using a novel multi-task agent training paradigm that unifies different pretraining strategies. The framework's performance is demonstrated across three distinct domains: robotics, gaming AI, and healthcare. The results are impressive, showing that the integrated LLMs enable agents to cooperate effectively with other agents, including humans. However, more training and safety filtering are necessary, particularly in gaming, where overly realistic robots could contribute to social withdrawal, and in healthcare, where self-prescription based on the model's recommendations is strongly discouraged.

In multi-agent systems (MAS), each agent can only access the information of its respective human user. This creates information asymmetry, posing a new challenge when leveraging agents' communication to enhance human cooperation. To address this, Liu [24] proposed a novel MAS paradigm called iAgents, short for Informative Multi-Agent Systems. iAgents tackle information asymmetry by proactively exchanging human information necessary for task resolution, mirroring the human social network within the agent network. To guide agents' communication toward efficient information exchange, iAgents employ InfoNav, an innovative agent reasoning mechanism. By organizing human information in a mixed memory in conjunction with InfoNav, iAgents provide agents with accurate and comprehensive information for exchange. Additionally, they introduced InformativeBench, the first benchmark specifically designed to evaluate the task-solving proficiency of LLM agents in the context of information asymmetry.

Although Large Language Models (LLMs) are increasingly used in financial applications due to their ability to perform complex tasks, high-quality sequential financial investment decision-making remains challenging. Yu [19]introduced FINCON, a conceptual verbal reinforcement multi-agent framework designed for various financial tasks, built on LLM technology. FINCON employs a manager-analyst communication hierarchy modeled after successful real-world investment firm structures. The framework's core lies in the Synthesized Manager-Analyst hierarchical communication structure and a dual-level risk control component. Scaling FINCON's framework to manage large portfolios with tens of thousands of assets while maintaining the high decision-making quality observed in smaller portfolios is a promising area for future research.

These agents, integrated with LLMs, demonstrate impressive cooperation skills with other agents through communication. However, this ability is limited when dealing with state information that is difficult to express or understand through natural language, making message exchange in such contexts challenging.

5. Future work

The combination of LLMs and MAS has the potential to address many of the challenges we face today, including solving real-world problems, improving human interaction, and enhancing reasoning abilities. For instance, Sun [37] focused on cooperative tasks where multiple agents work toward a common goal and interact with each other. They also explored scenarios that involve human participation, facilitated by the framework's linguistic component.

The integration of Large Language Models (LLMs) into Multi-Agent Systems (MAS) is still in its infancy, and there are several promising avenues for future research that can significantly enhance the capabilities of MAS:

5.1. Generalizing Reasoning Across Diverse Contexts

While LLMs have shown considerable reasoning abilities in specific scenarios, there is a need to develop more robust methods that enable these models to generalize reasoning across a broader range of tasks and environments. This includes improving prompt engineering techniques and exploring new architectures that can handle dynamic, real-world scenarios more effectively.

5.2. Improving Adaptive Capabilities

Enhancing the adaptability of MAS remains a critical challenge. Future research should focus on developing more sophisticated adaptive mechanisms that allow agents to seamlessly adjust to new environments and collaborate with unfamiliar agents. This could involve the integration of more advanced reinforcement learning techniques and the development of adaptive frameworks that can learn from minimal human intervention.

5.3. Enhancing Cooperative Strategies

The cooperative potential of MAS can be further explored by investigating how agents can better understand and leverage the intentions and capabilities of their peers. Future work could explore new models of communication that go beyond natural language, potentially incorporating multi-modal inputs such as visual or sensory data to improve the accuracy and efficiency of cooperation.

DOI: 10.54254/2755-2721/109/20241330

5.4. Addressing Information Asymmetry

Although applications like iAgents have made strides in overcoming information asymmetry, future research should explore more effective ways to manage and distribute information within MAS. This could include the development of more refined information-sharing protocols or the introduction of decentralized data structures that can better support the dynamic needs of agents.

5.5. Scalability and Efficiency

As MAS continue to grow in complexity, scalability remains a significant challenge. Future work should focus on optimizing the computational efficiency of LLM-powered agents, ensuring that they can operate at scale without compromising performance. This may involve the exploration of distributed computing techniques or the development of more efficient algorithms for agent interaction.

5.6. Ethical and Safety Considerations

The increasing sophistication of MAS, particularly those powered by LLMs, raises important ethical and safety concerns. Future research should address these issues by developing robust safety protocols and ethical guidelines that ensure responsible deployment. This includes mitigating risks related to over-reliance on AI agents, potential biases in decision-making, and the security of sensitive information.

6. Conclusion

The integration of Large Language Models (LLMs) into Multi-Agent Systems (MAS) presents a transformative opportunity to enhance the capabilities of intelligent agents in complex environments. This survey has provided a comprehensive overview of the current state of LLMs in MAS, highlighting their impact on agent cooperation, reasoning, and adaptability. While significant progress has been made, numerous challenges and open questions remain, particularly in the areas of adaptation, cooperation, and scalability.

Future research efforts should focus on addressing these challenges, with a particular emphasis on developing more generalized reasoning abilities, improving adaptive mechanisms, enhancing cooperative strategies, and optimizing the scalability of MAS. Additionally, ethical and safety considerations must be prioritized to ensure the responsible deployment of LLM-powered agents in real-world applications.

As this field continues to evolve, the ongoing exploration of LLMs in MAS holds great promise for advancing the state-of-the-art in artificial intelligence, enabling the development of more intelligent, adaptive, and cooperative systems capable of tackling increasingly complex tasks in a wide range of domains.

References

[1] Talebirad, Y. and Nadiri, A. Multi-agent collaboration: Harnessing the power of intelligent llm agents. arXiv preprint arXiv:2306.03314, 2023.

[2] Zhang, J., Xu, X., and Deng, S. Exploring collaboration mechanisms for llm agents: A social psychology view, 2023a.

[3] Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., and Bernstein, M. S. Generative agents: Interactive simulacra of human behavior. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, pp. 1–22, 2023.

[4] Li, G., Hammoud, H. A. A. K., Itani, H., Khizbullin, D., and Ghanem, B. Camel: Communicative agents for" mind" exploration of large-scale language model society. arXiv preprint arXiv: 2303.17760, 2023.

[5] Jinxin, S., Jiabao, Z., Yilei, W., Xingjiao, W., Jiawen, L., and Liang, H. Cgmi: Configurable general multi-agent interaction framework. arXiv preprint arXiv:2308.12503, 2023.

[6] Shanshan, H., Qifan, Z., Yuhang, Y., Weizhao, J., Zhaozhou, X., and Chaoyang, H LLMMulti-Agent Systems: Challenges and Open Problems. arXiv preprint arXiv: /2402.03578v1

DOI: 10.54254/2755-2721/109/20241330

[7] Durante, Z., Sarkar, B., Gong, R., Taori, R., Noda, Y., Tang, P., Adeli, E., Lakshmikanth, S, K., Schulman, K., Milstein, A., Terzopoulos, D., Famoti, A., Kuno, N., Llorens, A., Vo, H., Ikeuchi, K., Fei-Fei, L., Gao, J., Wake, N., Huang, Q. An Interactive Agent Foundation Model. arXiv preprint arXiv: 2402.05929.
[8] Nascimento, N., Alencar, P., Cowan, P. Self-Adaptive Large Language Model (LLM)-Based Multiagent Systems. arXiv preprint arXiv:2307.06187, 2023.
[9] Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., and Mor datch, I. Improving factuality and reasoning in lan guage models through multiagent debate. arXiv preprint arXiv:2305.14325, 2023.
[10] Liang, T., He, Z., Jiao, W., Wang, X., Wang, Y., Wang, R., Yang, Y., Tu, Z., and Shi, S. Encouraging divergent thinking in large language models through multi-agent debate. arXiv preprint arXiv:2305.19118, 2023.
[11] Jung H Lee, Henry J Kvinge, Scott Howland, Zachary New, John Buckheit, Lauren A. Phillips, Elliott Skomski, Jessica Hibler, Courtney D. Corley, Nathan O. Hodas. Adaptive Transfer Learning: a simple but effective transfer learning. arXiv preprint arXiv: 2111.10937.
[12] Siyu Yuan, Kaitao Song, Jiangjie Chen, Xu Tan, Dongsheng Li, Deqing Yang. EVOAGENT: Towards Automatic Multi-Agent Generation via Evolutionary Algorithms. arXiv preprint arXiv: 2406.14228v1.
[13] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In The Twelfth International Conference on Learning Representations, 2024.
[14] Guangyao Chen, Siwei Dong, Yu Shu, Ge Zhang, Jaward Sesay, Börje F Karlsson, Jie Fu, and Yemin Shi. Autoagents: A framework for automatic agent generation. arXiv preprint arXiv: 2309.17288, 2023.
[15] Michal Kosinski. 2023. Theory of mind may have spon taneously emerged in large language models. arXiv preprint arXiv:2302.02083.
[16] Shima Rahimi Moghaddam and Christopher J Honey. 2023. Boosting theory-of-mind performance in large language models via prompting. arXiv preprint arXiv:2304.11490.
[17] Tomer D. Ullman. Large Language Models Fail on Trivial Alterations to Theory-of-Mind Tasks. 2023. arXiv preprint arXiv: 2302.08399v5
[18] Qiao, Jie et al. A Lancet Commission on 70 years of women's reproductive, maternal, newborn, child, and adolescent health in China. The Lancet, Volume 397, Issue 10293, 2497 – 2536.
[19] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 2998–3009, 2023.
[20] Bhargavi Paranjape, Scott Lundberg, Sameer Singh, Hannaneh Hajishirzi, Luke Zettlemoyer, Marco Tulio Ribeiro. ART: Automatic multi-step reasoning and tool-use for large language models. arXiv preprint arXiv: 2303.09014.
[21] Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng1 Yupeng Cao, Zhi Chen, Jordan W. Suchow, Rong Liu, Zhenyu Cui, Denghui Zhang, Zhaozhuo Xu, Koduvayur Subbalakshmi, Guojun Xiong, Yueru He, Jimin Huang, Dong Li, Qianqian Xie. 2024. FINCON: ASynthesized LLM Multi-Agent System with Conceptual Verbal Reinforcement for Enhanced Financial Decision Making. arXiv preprint arXiv: 2407.06567.
[22] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311.

[23] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Lamda: Language models for dialog applications. arXiv preprint arXiv:2201.08239.
[24] Sébastien Bubeck, Varun Chandrasekaran, Ronen El dan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lund berg, et al. 2023. Sparks of artificial general intelli gence: Early experiments with gpt-4. arXiv preprint arXiv:2303.12712.
[25] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. arXiv preprint arXiv:2201.11903.
[26] Wei Liu, Chenxi Wang, YifeiWang, ZihaoXie, RennaiQiu, YufanDang, Zhuoyun Du, WeizeChen, ChengYang, ChenQian. Autonomous Agents for Collaborative Task under Information Asymmetry. arXiv preprint arXiv: 2406.14928v1
[27] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Haus man, et al. 2022. Do as i can, not as i say: Ground ing language in robotic affordances. arXiv preprint arXiv:2204.01691.
[28] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. Language models as zero-shot planners: Extracting actionable knowledge for em bodied agents. In International Conference on Ma chine Learning, pages 9118–9147. PMLR.
[29] Balaji, P.G., Srinivasan, D. (2010). An Introduction to Multi-Agent Systems. In: Srinivasan, D., Jain, L.C. (eds) Innovations in Multi-Agent Systems and Applications - 1. Studies in Computational Intelligence, vol 310. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-14435-6_1
[30] Long, Qian, Zihan Zhou, Abhibav Gupta, Fei Fang, Yi Wu, and Xiaolong Wang. "Evolutionary population curriculum for scaling multi-agent reinforcement learning." arXiv preprint arXiv: 2003.10423 (2020).
[31] Long, Qian, Fangwei Zhong, Mingdong Wu, Yizhou Wang and Song-Chun Zhu. “SocialGFs: Learning Social Gradient Fields for Multi-Agent Reinforcement Learning.” ArXiv abs/2405.01839 (2024): n. pag.
[32] Nathalia Nascimento, Paulo Alencar, Donald Cowan. Self-Adaptive Large Language Model (LLM)-Based Multiagent Systems. arXiv preprint arXiv: 2307.06187v1.
[33] Min Meng, Gaoxi Xiao, Beibei Li. Adaptive consensus for heterogeneous multi-agent systems under sensor and actuator attacks. https://www.sciencedirect.com/science/article/abs/pii/S0005109820304404
[34] Huao Li, Yu Quan Chong, Simon Stepputtis, Joseph Campbell, Dana Hughes, Michael Lewis, Katia Sycara. 2023. Theory of Mind for Multi-Agent Collaboration via Large Language Models. arXiv preprint arXiv: 2310.10701.
[35] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Rui Wang, Yujiu Yang, Shuming Shi, Zhaopeng Tu. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. 2023. arXiv preprint arXiv: 2305.19118.
[36] Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, Igor Mordatch. Improving Factuality and Reasoning in Language Models through Multiagent Debate. arXiv preprint arXiv: 2305.14325v1
[37] Chuanneng Sun, Songjun Huang, Dario Pompili. LLM-based Multi-Agent Reinforcement Learning: Current and Future Directions. arXiv preprint arXiv: 2405.11106v1.