

Large Language Model based Multi-Agents: A Survey of Progress and Challenges

Taicheng Guo¹, Xiuying Chen², Yaqi Wang^{3*}, Ruidi Chang, Shichao Pei⁴,
Nitesh V. Chawla¹, Olaf Wiest¹, Xiangliang Zhang^{1†}

¹University of Notre Dame, ²King Abdullah University of Science and Technology

³Southern University of Science and Technology, ⁴University of Massachusetts Boston

{tguo2, nchawla, owiest, xzhang33}@nd.edu, xiuying.chen@kaust.edu.sa, ywang84@nd.edu,
ruidic@alumni.cmu.edu, shichao.pei@umb.edu

Abstract

Large Language Models (LLMs) have achieved remarkable success across a wide array of tasks. Due to the impressive planning and reasoning abilities of LLMs, they have been used as autonomous agents to do many tasks automatically. Recently, based on the development of using one LLM as a single planning or decision-making agent, LLM-based multi-agent systems have achieved considerable progress in complex problem-solving and world simulation. To provide the community with an overview of this dynamic field, we present this survey to offer an in-depth discussion on the essential aspects of multi-agent systems based on LLMs, as well as the challenges. Our goal is for readers to gain substantial insights on the following questions: What domains and environments do LLM-based multi-agents simulate? How are these agents profiled and how do they communicate? What mechanisms contribute to the growth of agents' capacities? For those interested in delving into this field of study, we also summarize the commonly used datasets or benchmarks for them to have convenient access. To keep researchers updated on the latest studies, we maintain an open-source [GitHub repository](#), dedicated to outlining the research on LLM-based multi-agent systems.

1 Introduction

Large Language Models (LLMs) have recently shown remarkable potential in reaching a level of reasoning and planning capabilities comparable to humans. This ability exactly aligns with the expectations of humans for autonomous agents that can perceive the surroundings, make decisions, and take actions in response [Xi *et al.*, 2023; Wooldridge and Jennings, 1995; Russell and Norvig, 2009; Guo *et al.*, 2023; Liang *et al.*, 2023]. Hence, LLM-based agent has been studied and rapidly developed to understand and generate human-like instructions, facilitating sophisticated interactions and

decision-making in a wide range of contexts [Yao *et al.*, 2023; Shinn *et al.*, 2023; Li *et al.*, 2023d]. Timely survey papers systematically summarize the progress of LLM-based agents, as seen in works [Xi *et al.*, 2023; Wang *et al.*, 2023b].

Based on the inspiring capabilities of the single LLM-based agent, LLM-based Multi-Agents have been proposed to leverage the collective intelligence and specialized profiles and skills of multiple agents. Compared to systems using a single LLM-powered agent, multi-agent systems offer advanced capabilities by 1) specializing LLMs into various distinct agents, each with different capabilities, and 2) enabling interactions among these diverse agents to simulate complex real-world environments effectively. In this context, multiple autonomous agents collaboratively engage in planning, discussions, and decision-making, mirroring the cooperative nature of human group work in problem-solving tasks. This approach capitalizes on the communicative capabilities of LLMs, leveraging their ability to generate text for communication and respond to textual inputs. Furthermore, it exploits LLMs' extensive knowledge across various domains and their latent potential to specialize in specific tasks. Recent research has demonstrated promising results in utilizing LLM-based multi-agents for solving various tasks, such as software development [Hong *et al.*, 2023; Qian *et al.*, 2023], multi-robot systems [Mandi *et al.*, 2023; Zhang *et al.*, 2023c], society simulation [Park *et al.*, 2023; Park *et al.*, 2022], policy simulation [Xiao *et al.*, 2023; Hua *et al.*, 2023], and game simulation [Xu *et al.*, 2023c; Wang *et al.*, 2023c]. Due to the nature of interdisciplinary study in this field, it has attracted a diverse range of researchers, expanding beyond AI experts to include those from social science, psychology, and policy research. The volume of research papers is rapidly increasing, as shown in Fig. 1 (inspired by the design in [Gao *et al.*, 2023b]), thus broadening the impact of LLM-based Multi-Agent research. Nonetheless, earlier efforts were undertaken independently, resulting in an absence of a systematic review to summarize them, establish comprehensive blueprint of this field, and examine future research challenges. This underscores the significance of our work and serves as the motivation behind presenting this survey paper, dedicated to the research on LLM-based multi-agent systems.

We expect that our survey can make significant contributions to both the research and development of LLMs and to

*This work was done when Yaqi was visiting students at the University of Notre Dame.

†Corresponding author.




Figure 1: The rising trend in the research field of LLM-based Multi-Agents. For Problem Solving and World Simulation, we categorize current work into several categories and count the number of papers of different types at 3-month intervals. The number at each leaf node denotes the count of papers within that category.

a wider range of interdisciplinary studies employing LLMs. Readers will gain a comprehensive overview of LLM-based Multi-Agent (LLM-MA) systems, grasp the fundamental concepts involved in establishing multi-agent systems based on LLMs, and catch the latest research trends and applications in this dynamic field. We recognize that this field is in its early stages and is rapidly evolving with fresh methodologies and applications. To provide a sustainable resource complementing our survey paper, we maintain an open-source GitHub repository¹. We hope that our survey will inspire further exploration and innovation in this field, as well as applications across a wide array of research disciplines.

To assist individuals from various backgrounds in understanding LLM-MA techniques and to complement existing surveys by tackling unresolved questions, we have organized our survey paper in the following manner. After laying out the background knowledge in Section 2, we address a pivotal question: *How are LLM-MA systems aligned with the collaborative task-solving environment?* To answer this, we present a comprehensive schema for positioning, differentiating, and connecting various aspects of LLM-MA systems in Section 3. We delve into this question by discussing: 1) the **agents-environment interface**, which details how agents interact with the task environment; 2) **agent profiling**, which explains how an agent is characterized by an LLM to behave

in specific ways; 3) **agent communication**, which examines how agents exchange messages and collaborate; and 4) **agent capability acquisition**, which explores how agents develop their abilities to effectively solve problems. An additional perspective for reviewing studies about LLM-MA is their application. In Section 4, we categorize current **applications** into two primary streams: multi-agents for **problem-solving** and multi-agents for **world simulation**. To guide individuals in identifying appropriate **tools and resources**, we present open-source implementation frameworks for studying LLM-MA, as well as the usable datasets and benchmarks in Section 5. Based on the previous summary, we open the discussion for future research challenges and opportunities in Section 6. The conclusions are summarized in Section 7.

2 Background

2.1 Single-Agent Systems Powered LLMs

We introduce the background by first outlining the capabilities of a single-agent system based on LLMs, following the discussion presented in [Weng, 2023].

Decision-making Thought: This term denotes the capability of LLM-based agents, guided by prompts, to break down complex tasks into smaller subgoals [Khot *et al.*, 2023], think through each part methodically (sometimes exploring multiple paths) [Yao *et al.*, 2023], and learn from past experiences [Shinn *et al.*, 2023] to perform better decision-making

¹<https://github.com/taichengguo/LLM.MultiAgents.Survey.Papers>

on complex tasks. This capability enhances the autonomy of a single LLM-based agent and bolsters its effectiveness in problem-solving.

Tool-use: LLM-based agents’ tool-use capability allows them to leverage external tools and resources to accomplish tasks, enhancing their functional capabilities and operate more effectively in diverse and dynamic environments [Li *et al.*, 2023d; Ruan *et al.*, 2023; Gao *et al.*, 2023b].

Memory: This ability refers to the capability of LLM-based agent for conducting in-context learning [Dong *et al.*, 2023a] as short memory or external vector database [Lewis *et al.*, 2021] as long memory to preserve and retrieve information over prolonged periods [Wang *et al.*, 2023b]. This ability enables a single LLM-based agent to maintain contextual coherence and enhance learning from interactions.

2.2 Single-Agent VS. Multi-Agent Systems

Single-Agent systems empowered by LLMs have shown inspiring cognitive abilities [Sumers *et al.*, 2023]. The construction of such systems concentrates on formulating their internal mechanisms and interactions with the external environment. Conversely, LLM-MA systems emphasize diverse agent profiles, inter-agent interactions, and collective decision-making processes. From this perspective, more dynamic and complex tasks can be tackled by the collaboration of multiple autonomous agents, each of which is equipped with unique strategies and behaviors, and engaged in communication with one another.

3 Dissecting LLM-MA Systems: Interface, Profiling, Communication, and Capabilities

In this section, we delve into the intricacies of LLM-MA systems, where multiple autonomous agents engage in collaborative activities akin to human group dynamics in problem-solving scenarios. A critical inquiry we address is how these LLM-MA systems are aligned to their operational environments and the collective objectives they are designed to achieve. To shed light on this, we present the general architecture of these systems in Fig. 2. Our analysis dissects the operational framework of these systems, focusing on four key aspects: the agents-environment interface, agent profiling, agent communication, and agent capability acquisition.

3.1 Agents-Environment Interface

The operational environments defines the specific contexts or settings in which the LLM-MA systems are deployed and interact. For example, these environments can be like software development [Hong *et al.*, 2023], gaming [Mao *et al.*, 2023], and various other domains such as financial markets [Li *et al.*, 2023g] or even social behavior modeling [Park *et al.*, 2023]. The LLM-based agents perceive and act within the environment, which in turn influences their behavior and decision making. For example, in the Werewolf Game simulation, the sandbox environment sets the game’s framework, including transitions from day to night, discussion periods, voting mechanics, and reward rules. Agents, such as werewolves and the Seer, perform specific actions like killing or

checking roles. Following these actions, agents receive feedback from the environment, informing them of the game’s current state. This information guides the agents in adjusting their strategies over time, responding to the evolving gameplay and interactions with other agents. The Agents-Environment Interface refers to the way in which agents interact with and perceive the environment. It’s through this interface that agents understand their surroundings, make decisions, and learn from the outcomes of their actions. We categorize the current interfaces in LLM-MA systems into three types, *Sandbox*, *Physical*, and *None*, as detailed in Table 1. The *Sandbox* refers to a simulated or virtual environment built by human where agents can interact more freely and experiment with various actions and strategies. This kind of interface is widely used in software development (code interpreter as simulated environment) [Hong *et al.*, 2023], gaming (using game rules as simulated environment) [Mao *et al.*, 2023], etc. The *Physical* is a real-world environment where agents interact with physical entities and obey real-world physics and constraints. In physical space, agents normally need to take actions that can have direct physical outcomes. For example, in tasks such as sweeping the floor, making sandwiches, packing groceries, and arranging cabinets, robotic agents are required to perform actions iteratively, observe the physical environment, and continuously refine their actions [Mandi *et al.*, 2023]. Lastly, *None* refers to scenarios where there is no specific external environment, and agents do not interact with any environment. For example, many applications [Du *et al.*, 2023; Xiong *et al.*, 2023; Chan *et al.*, 2023] utilize multiple agents to debate a question to reach a consensus. These applications primarily focus on communication among agents and do not depend on the external environment.

3.2 Agents Profiling

In LLM-MA systems, agents are defined by their traits, actions, and skills, which are tailored to meet specific goals. Across various systems, agents assume distinct roles, each with comprehensive descriptions encompassing characteristics, capabilities, behaviors, and constraints. For instance, in gaming environments, agents might be profiled as players with varying roles and skills, each contributing differently to the game’s objectives. In software development, agents could take on the roles of product managers and engineers, each with responsibilities and expertise that guide the development process. Similarly, in a debating platform, agents might be designated as proponents, opponents, or judges, each with unique functions and strategies to fulfill their roles effectively. These profiles are crucial for defining the agents’ interactions and effectiveness within their respective environments. Table 1 lists the agent *Profiles* in recent LLM-MA works.

Regarding the *Agent Profiling Methods*, we categorized them into three types: *Pre-defined*, *Model-Generated*, and *Data-Derived*. In the *Pre-defined* cases, agent profiles are explicitly defined by the system designers. The *Model-Generated* method creates agent profiles by models, e.g., large language models. The *Data-Derived* method involves constructing agent profiles based on pre-existing datasets.




Figure 2: The Architecture of LLM-MA Systems.

3.3 Agents Communication

The communication between agents in LLM-MA systems is the critical infrastructure supporting collective intelligence. We dissect agent communication from three perspectives: 1) *Communication Paradigms*: the styles and methods of interaction between agents; 2) *Communication Structure*: the organization and architecture of communication networks within the multi-agent system; and 3) *Communication Content* exchanged between agents.

Communication Paradigms: Current LLM-MA systems mainly take three paradigms for communication: *Cooperative*, *Debate*, and *Competitive*. *Cooperative* agents work together towards a shared goal or objectives, typically exchanging information to enhance a collective solution. The *Debate* paradigm is employed when agents engage in argumentative interactions, presenting and defending their own viewpoints or solutions, and critiquing those of others. This paradigm is ideal for reaching a consensus or a more refined solution. *Competitive* agents work towards their own goals that might be in conflict with the goals of other agents.

Communication Structure: Fig. 3 shows four typical communication structures in LLM-MA systems. *Layered* communication is structured hierarchically, with agents at each level having distinct roles and primarily interacting within their layer or with adjacent layers. [Liu *et al.*, 2023] introduces a framework called Dynamic LLM-Agent Network (DyLAN), which organizes agents in a multi-layered feed-forward network. This setup facilitates dynamic interactions, incorporating features like inference-time agent selection and an early-stopping mechanism, which collectively enhance the efficiency of cooperation among agents. *Decentralized* communication operates on a peer-to-peer net-




Figure 3: The Agent Communication Structure.

work, where agents directly communicate with each other, a structure commonly employed in world simulation applications. *Centralized* communication involves a central agent or a group of central agents coordinating the system’s communication, with other agents primarily interacting through this central node. *Shared Message Pool* is proposed by MetaGPT [Hong *et al.*, 2023] to improve the communication efficiency. This communication structure maintains a shared message pool where agents publish messages and subscribe to relevant messages based on their profiles, thereby boosting

communication efficiency.

Communication Content: In LLM-MA systems, the Communication Content typically takes the form of text. The specific content varies widely and depends on the particular application. For example, in software development, agents may communicate with each other about code segments. In simulations of games like Werewolf, agents might discuss their analyses, suspicions, or strategies.

3.4 Agents Capabilities Acquisition

The Agents Capabilities Acquisition is a crucial process in LLM-MA, enabling agents to learn and evolve dynamically. In this context, there are two fundamental concepts: the types of *feedback* from which agents should learn to enhance their capabilities, and the *strategies* for agents to *adjust* themselves to effectively solve complex problems.

Feedback: Feedback involves the critical information that agents receive about the outcome of their actions, helping the agents learn the potential impact of their actions and adapt to complex and dynamic problems. In most studies, the format of feedback provided to agents is textual. Based on the sources from which agents receive this feedback, it can be categorized into four types. **1) Feedback from Environment**, e.g., from either real world environments or virtual environments [Wang *et al.*, 2023b]. It is prevalent in most LLM-MA for problem-solving scenarios, including Software Development (agents obtain feedback from Code Interpreter), and Embodied multi-agents systems (robots obtain feedback from real-world or Simulated environments). **2) Feedback from Agents Interactions** means that the feedback comes from the judgement of other agents or from agents communications. It is common in problem-solving scenarios like science debates, where agents learn to critically evaluate and refine the conclusions through communications. In world simulation scenarios such as Game Simulation, agents learn to refine strategies based on previous interactions between other agents. **3) Human Feedback** comes directly from humans and is crucial for aligning the multi-agent system with human values and preferences. This kind of feedback is widely used in most “Human-in-the-loop” applications [Wang *et al.*, 2021]. Last **4) None**. In some cases, there is no feedback provided to the agents. This often happens for world simulation works focused on analyzing simulated results rather than the planning capabilities of agents. In such scenarios, like propagation simulation, the emphasis is on result analysis, and hence, feedback is not a component of the system.

Agents Adjustment to Complex Problems: To enhance their capabilities, agents in LLM-MA systems can adapt through three main solutions. **1) Memory.** Most LLM-MA systems leverage a memory module for agents to adjust their behavior. Agents store information from previous interactions and feedback in their memory. When performing actions, they can retrieve relevant, valuable memories, particularly those containing successful actions for similar past goals, as highlighted in [Wang *et al.*, 2023b]. This process aids in enhancing their current actions. **2) Self-Evolution.** Instead of only relying on the historical records

to decide subsequent actions as seen in *Memory*-based solutions, agents can dynamically self-evolve by modifying themselves such as altering their initial goals and planning strategies, and training themselves based on feedback or communication logs. [Nascimento *et al.*, 2023] proposes a self-control loop process to allow each agent in the multi-agents systems to be self-managed and self-adaptive to dynamic environments, thereby improving the cooperation efficiency of multiple agents. [Zhang *et al.*, 2023b] introduces ProAgent which anticipates teammates’ decisions and dynamically adjusts each agent’s strategies based on the communication logs between agents, facilitating mutual understanding and improving collaborative planning capability. [Wang *et al.*, 2023a] discusses a *Learning through Communication* (LTC) paradigm, using the communication logs of multi-agents to generate datasets to train or fine-tune LLMs. LTC enables continuous adaptation and improvement of agents through interaction with their environments and other agents, breaking the limits of in-context learning or supervised fine-tuning, which don’t fully utilize the feedback received during interactions with the environment and external tools for continuous training. Self-Evolution enables agents’ autonomous adjustment in their profiles or goals, rather than just learning from historical interactions. **3) Dynamic Generation.** In some scenarios, the system can generate new agents on-the-fly during its operation [Chen *et al.*, 2023a; Chen *et al.*, 2023c]. This capability enables the system to scale and adapt effectively, as it can introduce agents that are specifically designed to address current needs and challenges.

With the scaling up LLM-MA with a larger number of agents, the escalating complexity of managing various kinds of agents has been a critical problem. *Agents Orchestration* emerged as a pivotal challenge and began to gain attention in [Moura, 2023; Dibia, 2023]. We will further discuss this topic in Section 6.4.

4 Applications

LLM-MA systems have been used in a wide range of applications. We summarize two kinds of applications in Table 1: **Problem Solving** and **World Simulation**. We elaborate on these applications below. Note that this is a fast growing research field and new applications appear almost everyday. We maintain an [open source repository](#) to report the latest work.

4.1 LLM-MA for Problem Solving

The main motivation of using LLM-MA for problem solving is to harness the collective capabilities of agents with specialized expertise. These agents, each acting as individuals, collaborate to address complex problems effectively, such as software development, embodied agents, science experiments and science debate. These application examples are introduced next.

4.1.1 Software Development

Given that software development is a complex endeavor requiring the collaboration of various roles like product managers, programmers, and testers, LLM-MA systems are typically set to emulate these distinct roles and collaborate to address the intricate challenge. Following the waterfall or

Motivation	Research Domain & Goals		Work	Agents-Env. Interface	Agents Profiling		Agents Communication		Agents Capabilities Acquisition	
					Profiling methods	Profiles (examples)	Paradigms	Structure	Feedback from	Agents Adjustment
Problem Solving	Software development		[Qian <i>et al.</i> , 2023]	Sandbox	Pre-defined, Model-Generated	CTO, programmer	Cooperative	Layered	Environment, Agent interaction, Human	Memory, Self-Evolution
			[Hong <i>et al.</i> , 2023]	Sandbox	Pre-defined	Product Manager, Engineer	Cooperative	Layered, Shared Message Pool	Environment, Agent interaction, Human	Memory, Self-Evolution
			[Dong <i>et al.</i> , 2023b]	Sandbox	Pre-defined, Model-Generated	Analyst, coder	Cooperative	Layered	Environment, Agent interaction	Memory, Self-Evolution
	Embodied Agents	Multi-robot planning	[Chen <i>et al.</i> , 2023d]	Sandbox, Physical	Pre-defined	Robots	Cooperative	Centralized, Decentralized	Environment, Agent interaction	Memory
		Multi-robot collaboration	[Mandi <i>et al.</i> , 2023]	Sandbox, Physical	Pre-defined	Robots	Cooperative	Decentralized	Environment, Agent interaction	Memory
		Multi-Agents cooperation	[Zhang <i>et al.</i> , 2023c]	Sandbox	Pre-defined	Robots	Cooperative	Decentralized	Environment, Agent interaction	Memory
	Science Experiments	Optimization of MOF	[Zheng <i>et al.</i> , 2023]	Physical	Pre-defined	Strategy planers, literature collector, coder	Cooperative	Centralized	Environment, Human	Memory
	Science Debate	Improving Factuality	[Du <i>et al.</i> , 2023]	None	Pre-defined	Agents	Debate	Decentralized	Agent interaction	Memory
		Examining, Inter-Consistency	[Xiong <i>et al.</i> , 2023]	None	Pre-defined	Proponent, Opponent, Judge	Debate	Centralized, Decentralized	Agent interaction	Memory
		Evaluators for debates	[Chan <i>et al.</i> , 2023]	None	Pre-defined	Agents	Debate	Centralized, Decentralized	Agent interaction	Memory
		Multi-Agents for Medication	[Tang <i>et al.</i> , 2023]	None	Pre-defined	Cardiology, Surgery	Debate, Cooperative	Centralized, Decentralized	Agent interaction	Memory
World Simulation	Society	Modest Community (25 persons)	[Park <i>et al.</i> , 2023]	Sandbox	Model-generated	Pharmacy, shopkeeper	-	-	Environment, Agent interaction	Memory
		Online community (1000 persons)	[Park <i>et al.</i> , 2022]	None	Pre-defined, Model-generated	Camping, fishing	-	-	Agent interaction	Dynamic Generation
		Emotion propagation	[Gao <i>et al.</i> , 2023a]	None	Pre-defined, Model-generated	Real-world user	-	-	Agent interaction	Memory
		Real-time social interactions	[Kaiya <i>et al.</i> , 2023]	Sandbox	Pre-defined	Real-world user	-	-	Environment, Agent interaction	Memory
		Opinion dynamics	[Li <i>et al.</i> , 2023a]	None	Pre-defined	NIN, NINL, NIL	-	-	Agent interaction	Memory
	Gaming	WereWolf	[Xu <i>et al.</i> , 2023b] [Xu <i>et al.</i> , 2023c]	Sandbox	Pre-defined	Seer, werewolf, villager	Cooperative, Debate, Competitive	Decentralized	Environment, Agent interaction	Memory
		Avalon	[Light <i>et al.</i> , 2023a] [Wang <i>et al.</i> , 2023c]	Sandbox	Pre-defined	Servant, Merlin, Assassin	Cooperative, Debate, Competitive	Decentralized	Environment, Agent interaction	Memory
		Welfare Diplomacy	[Mukobi <i>et al.</i> , 2023]	Sandbox	Pre-defined	Countries	Cooperative, Competitive	Decentralized	Environment, Agent interaction	Memory
	Psychology	Human behavior Simulation	[Aher <i>et al.</i> , 2023]	Sandbox	Pre-defined	Humans	-	-	Agent interaction	Memory
		Collaboration Exploring	[Zhang <i>et al.</i> , 2023d]	None	Pre-defined	Agents	Cooperative, Debate	Decentralized	Agent interaction	Memory
	Economy	Macroeconomic simulation	[Li <i>et al.</i> , 2023e]	None	Pre-defined, Model-generated	Labor	Cooperative	Decentralized	Agent interaction	Memory
		Information Marketplaces	[Anonymous, 2023]	Sandbox	Pre-defined, Data-Derived	Buyer	Cooperative, Competitive	Decentralized	Environment, Agent interaction	Memory
		Improving financial trading	[Li <i>et al.</i> , 2023g]	Physical	Pre-defined	Trader	Debate	Decentralized	Environment, Agent interaction	Memory
		Economic theories	[Zhao <i>et al.</i> , 2023]	Sandbox	Pre-defined, Model-Generated	Restaurant, Customer	Competitive	Decentralized	Environment, Agent interaction	Memory, Self-Evolution
	Recommender Systems	Simulating user behaviors	[Zhang <i>et al.</i> , 2023a]	Sandbox	Data-Derived	Users from MovieLens-1M	-	-	Environment	Memory
		Simulating user-item interactions	[Zhang <i>et al.</i> , 2023c]	Sandbox	Pre-defined, Data-Derived	User Agents Item Agents	Cooperative	Decentralized	Environment, Agent interaction	Memory
	Policy Making	Public Administration	[Xiao <i>et al.</i> , 2023]	None	Pre-defined	Residents	Cooperative	Decentralized	Agent interaction	Memory
		War Simulation	[Hua <i>et al.</i> , 2023]	None	Pre-defined	Countries	Competitive	Decentralized	Agent interaction	Memory
	Disease	Human Behaviors to epidemics	[Ghaffarzadeegan <i>et al.</i> , 2023]	Sandbox	Pre-defined, Model-Generated	Conformity traits	Cooperative	Decentralized	Environment, Agent interaction	Memory
		Public health	[Williams <i>et al.</i> , 2023]	Sandbox	Pre-defined, Model-Generated	Adults aged 18 to 64	Cooperative	Decentralized	Environment, Agent interaction	Memory, Dynamic Generation

Table 1: Summary of the LLM-MA studies. We categorize current work according to their motivation, research domains and goals, and detail each work from different aspects regarding Agents-Environment Interface, Agents Profiling, Agents Communication and Agents Capability Acquisition. “-” denotes that a particular element is not specifically mentioned in this work.

Standardized Operating Procedures (SOPs) workflow of the software development, the communication structure among agents is usually layered. Agents generally interact with the code interpreter, other agents or human to iteratively refine the generated code. [Li *et al.*, 2023b] first proposes a simple role-play agent framework, which utilizes the interplay of two roles to realize autonomous programming based on one-sentence user instruction. It provides insights into the “cognitive” processes of communicative agents. [Dong *et al.*, 2023b] makes LLMs work as distinct “experts” for sub-tasks

in software development, autonomously collaborating to generate code. Moreover, [Qian *et al.*, 2023] presents an end-to-end framework for software development, utilizing multiple agents for software development without incorporating advanced human teamwork experience. [Hong *et al.*, 2023] first incorporates human workflow insights for more controlled and validated performance. It encodes SOPs into prompts to enhance structured coordination. [Huang *et al.*, 2023a] delves deeper into multi-agent based programming by solving the problem of balancing code snippet generation with effective

test case generation, execution, and optimization.

4.1.2 Embodied Agents

Most embodied agents applications inherently utilize multiple robots working together to perform complex real-world planning and manipulation tasks such as warehouse management with heterogeneous robot capabilities. Hence, LLM-MA can be used to model robots with different capabilities and cooperate with each other to solve real-world physical tasks. [Dasgupta *et al.*, 2023] first explores the potential to use LLM as an action planner for embedded agents. [Mandi *et al.*, 2023] introduces RoCo, a novel approach for multi-robot collaboration that uses LLMs for high-level communication and low-level path planning. Each robotic arm is equipped with an LLM, cooperating with inverse kinematics and collision checking. Experimental results demonstrate the adaptability and success of RoCo in collaborative tasks. [Zhang *et al.*, 2023c] presents CoELA, a Cooperative Embodied Language Agent, managing discussions and task planning in an LLM-MA setting. This challenging setting is featured with decentralized control, complex partial observation, costly communication, and multi-objective long-horizon tasks. [Chen *et al.*, 2023d] investigates communication challenges in scenarios involving a large number of robots, as assigning each robot an LLM will be costly and impractical due to the long context. The study compares four communication frameworks, centralized, decentralized, and two hybrid models, to evaluate their effectiveness in coordinating complex multi-agent tasks. [Yu *et al.*, 2023] proposes CoNavGPT for multi-robot cooperative visual target navigation, integrating LLM as a global planner to assign frontier goals to each robot. [Chen *et al.*, 2023b] proposes an LLM-based consensus-seeking framework, which can be applied as a cooperative planner to a multi-robot aggregation task.

4.1.3 Science Experiments

Like multiple agents play as different specialists and cooperate to solve the Software Development and Embodied Agents problem, multiple agents can also be used to form a science team to conduct science experiments. One important difference from previous applications lies in the crucial role of human oversight, due to the high expenses of the science experiments and the hallucination of the LLM agents. Human experts are at the center of these agents to process the information of agents and give feedback to the agents. [Zheng *et al.*, 2023] utilizes multiple LLM-based agents, each focusing on specific tasks for the science experiments including strategy planning, literature search, coding, robotic operations, and labware design. All these agents interact with humans to work collaboratively to optimize the synthesis process of complex materials.

4.1.4 Science Debate

LLM-MA can be set for science debating scenarios, where agents debate with each other to enhance the collective reasoning capabilities in tasks such as Massive Multitask Language Understanding (MMLU) [Hendrycks *et al.*, 2020], Math problems [Cobbe *et al.*, 2021], and StrategyQA [Geva *et al.*, 2021]. The main idea is that each agent initially offers its own analysis of a problem, which is then followed

by a joint debating process. Through multiple rounds of debate, the agents converge on a single, consensus answer. [Du *et al.*, 2023] leverages the multi-agents debate process on a set of six different reasoning and factual accuracy tasks and demonstrates that LLM-MA debating can improve factuality. [Xiong *et al.*, 2023] focuses on the commonsense reasoning tasks and formulates a three-stage debate to align with real-world scenarios including fair debate, mismatched debate, and roundtable debate. The paper also analyzes the inter-consistency between different LLMs and claims that debating can improve the inter-consistency. [Tang *et al.*, 2023] also utilizes multiple LLM-based agents as distinct domain experts to do the collaborative discussion on a medical report to reach a consensus for medical diagnosis.

4.2 LLM-MA for World Simulation

Another mainstream application scenario of LLM-MA is the world simulation. Research in this area is rapidly growing and spans a diverse range of fields including social sciences, gaming, psychology, economics, policy-making, etc. The key reason for employing LLM-MA in world simulations lies in their exceptional role-playing abilities, which are crucial for realistically depicting various roles and viewpoints in a simulated world. The environment of world simulation projects is usually crafted to reflect the specific scenario being simulated, with agents designed in various profiles to match this context. Unlike the problem solving systems that focus on agent cooperation, world simulation systems involve diverse methods of agent management and communication, reflecting the complexity and variety of real-world interactions. Next, we explore simulations conducted in diverse fields.

4.2.1 Societal Simulation

In societal simulation, LLM-MA models are used to simulate social behaviors, aiming to explore the potential social dynamics and propagation, test social science theories, and populate virtual spaces and communities with realistic social phenomena [Park *et al.*, 2023]. Leveraging LLM’s capabilities, agents with unique profiles engage in extensive communication, generating rich behavioral data for in-depth social science analysis.

The scale of societal simulation has expanded over time, beginning with smaller, more intimate settings and progressing to larger, more intricate ones. Initial work by [Park *et al.*, 2023] introduces generative agents within an interactive sandbox environment reminiscent of the sims, allowing end users to engage with a modest community of 25 agents through natural language. At the same time, [Park *et al.*, 2022] develops Social Simulacra, which constructs a simulated community of 1,000 personas. This system takes a designer’s vision for a community—its goals, rules, and member personas—and simulates it, generating behaviors like posting, replying, and even anti-social actions. Building on this, [Gao *et al.*, 2023a] takes the concept further by constructing vast networks comprising 8,563 and 17,945 agents, respectively, designed to simulate social networks focused on the topics of Gender Discrimination and Nuclear Energy. This evolution showcases the increasing complexity and size of simulated environments in recent research. Recent studies such as [Chen *et al.*, 2023b;

Kaiya *et al.*, 2023; Li *et al.*, 2023a; Li *et al.*, 2023f; Ziems *et al.*, 2023] highlight the evolving complexity in multi-agent systems, LLM impacts on social networks, and their integration into social science research.

4.2.2 Gaming

LLM-MA is well-suited for creating simulated gaming environments, allowing agents to assume various roles within games. This technology enables the development of controlled, scalable, and dynamic settings that closely mimic human interactions, making it ideal for testing a range of game theory hypotheses [Mao *et al.*, 2023; Xu *et al.*, 2023b; Gong *et al.*, 2023]. Most games simulated by LLM-MA rely heavily on natural language communication, offering a sandbox environment within different game settings for exploring or testing game theory hypotheses including reasoning, cooperation, persuasion, deception, leadership, etc.

[Akata *et al.*, 2023] leverages behavioral game theory to examine LLMs' behavior in interactive social settings, particularly their performance in games like the iterated Prisoner's Dilemma and Battle of the Sexes. Furthermore, [Xu *et al.*, 2023b] proposes a framework using ChatArena library [Wu *et al.*, 2023b] for engaging LLMs in communication games like Werewolf, using retrieval and reflection on past communications for improvement, as well as the Chain-of-Thought mechanism [Wei *et al.*, 2022]. [Light *et al.*, 2023b] explores the potential of LLM agents in playing Resistance Avalon, introducing AVALONBENCH, a comprehensive game environment and benchmark for further developing advanced LLMs and multi-agent frameworks. [Wang *et al.*, 2023c] also focuses on the capabilities of LLM Agents in dealing with misinformation in the Avalon game, proposing the Recursive Contemplation (ReCon) framework to enhance LLMs' ability to discern and counteract deceptive information. [Xu *et al.*, 2023c] introduces a framework combining LLMs with reinforcement learning (RL) to develop strategic language agents for the Werewolf game. It introduces a new approach to use RL policy in the case that the action and state sets are not pre-defined but in the natural language setting. [Mukobi *et al.*, 2023] designs the "Welfare Diplomacy", a general-sum variant of the zero-sum board game Diplomacy, where players must balance military conquest and domestic welfare. It also offers an open-source benchmark, aiming to help improve the cooperation ability of multi-agent AI systems. On top of that, there is a work [Li *et al.*, 2023c] in a multi-agent cooperative text game testing the agents' Theory of Mind (ToM), the ability to reason about the concealed mental states of others and is fundamental to human social interactions, collaborations, and communications. [Fan *et al.*, 2023] comprehensively assesses the capability of LLMs as rational players, and identifies the weaknesses of LLM-based Agents that even in the explicit game process, agents may still overlook or modify refined beliefs when taking actions.

4.2.3 Psychology

In psychological simulation studies, like in the societal simulation, multiple agents are utilized to simulate humans with various traits and thought processes. However, unlike societal simulations, one approach in psychology involves directly applying psychological experiments to these agents.

This method focuses on observing and analyzing their varied behaviors through statistical methods. Here, each agent operates independently, without interacting with others, essentially representing different individuals. Another approach aligns more closely with societal simulations, where multiple agents interact and communicate with each other. In this scenario, psychological theories are applied to understand and analyze the emergent behavioral patterns. This method facilitates the study of interpersonal dynamics and group behaviors, providing insights into how individual psychological traits influence collective actions. [Ma *et al.*, 2023] explores the psychological implications and outcomes of employing LLM-based conversational agents for mental well-being support. It emphasizes the need for carefully evaluating the use of LLM-based agents in mental health applications from a psychological perspective. [Kovač *et al.*, 2023] introduces a tool named SocialAI school for creating interactive environments simulating social interactions. It draws from developmental psychology to understand how agents can acquire, demonstrate, and evolve social skills such as joint attention, communication, and cultural learning. [Zhang *et al.*, 2023d] explores how LLM agents, with distinct traits and thinking patterns, emulate human-like social behaviors such as conformity and majority rule. This integration of psychology into the understanding of agent collaboration offers a novel lens for examining and enhancing the mechanisms behind LLM-based multi-agents systems. [Aher *et al.*, 2023] introduces Turing Experiments to evaluate the extent to which large language models can simulate different aspects of human behaviors. The Turing Experiments replicate classical experiments and phenomena in psychology, economics, and sociology using a question-answering format to mimic experimental conditions. They also design a prompt that is used to simulate the responses of multiple different individuals by varying the name. By simulating various kinds of individuals via LLM, they show that larger models replicate human behavior more faithfully, but they also reveal a hyper-accuracy distortion, especially in knowledge-based tasks.

4.2.4 Economy

LLM-MA is used to simulate economic and financial trading environments mainly because it can serve as implicit computational models of humans. In these simulations, agents are provided with endowments, and information, and set with pre-defined preferences, allowing for an exploration of their actions in economic and financial contexts. This is similar to the way economists model 'homo economicus', the characterization of man in some economic theories as a rational person who pursues wealth for his own self-interest [Horton, 2023]. There are several studies demonstrate the diverse applications of LLM-MA in simulating economic scenarios, encompassing macroeconomic activities, information marketplaces, financial trading, and virtual town simulations. Agents interact in cooperative or debate, decentralized environments. [Li *et al.*, 2023e] employs LLMs for macroeconomic simulation, featuring prompt-engineering-driven agents that emulate human-like decision-making, thereby enhancing the realism of economic simulations compared to rule-based or other AI agents. [Anonymous, 2023] explores the buyer's inspec-

Motivation	Domain	Datasets and Benchmarks	Used by	Data Link
Problem Solving	Software Development	HumanEval	[Hong <i>et al.</i> , 2023]	Link
		MBPP	[Hong <i>et al.</i> , 2023]	Link
		SoftwareDev	[Hong <i>et al.</i> , 2023]	Link
	Embodied AI	RoCoBench	[Mandi <i>et al.</i> , 2023]	Link
		Communicative Watch-And-Help (C-WAH)	[Zhang <i>et al.</i> , 2023c]	Link
		ThreeDWorld Multi-Agent Transport (TDW-MAT)	[Zhang <i>et al.</i> , 2023c]	Link
		HM3D v0.2	[Yu <i>et al.</i> , 2023]	Link
	Science Debate	MMLU	[Tang <i>et al.</i> , 2023]	Link
		MedQA	[Tang <i>et al.</i> , 2023]	Link
		PubMedQA	[Tang <i>et al.</i> , 2023]	Link
		GSM8K	[Du <i>et al.</i> , 2023]	Link
		StrategyQA	[Xiong <i>et al.</i> , 2023]	Link
		Chess Move Validity	[Du <i>et al.</i> , 2023]	Link
World Simulation	Society	SOTOPIA	[Zhou <i>et al.</i> , 2023b]	/
		Gender Discrimination	[Gao <i>et al.</i> , 2023a]	/
		Nuclear Energy	[Gao <i>et al.</i> , 2023a]	/
	Gaming	Werewolf	[Xu <i>et al.</i> , 2023b]	/
		Avalon	[Light <i>et al.</i> , 2023b]	/
		Welfare Diplomacy	[Mukobi <i>et al.</i> , 2023]	/
		Layout in the Overcooked-AI environment	[Agashe <i>et al.</i> , 2023]	/
		Chameleon	[Xu <i>et al.</i> , 2023a]	Link
		Undercover	[Xu <i>et al.</i> , 2023a]	Link
	Psychology	Ultimatum Game TE	[Aher <i>et al.</i> , 2023]	Link
		Garden Path TE	[Aher <i>et al.</i> , 2023]	Link
		Wisdom of Crowds TE	[Aher <i>et al.</i> , 2023]	Link
	Recommender System	MovieLens-1M	[Zhang <i>et al.</i> , 2023a]	Link
		Amazon review dataset	[Zhang <i>et al.</i> , 2023e]	/
	Policy Making	Board Connectivity Evaluation	[Hua <i>et al.</i> , 2023]	Link

Table 2: Datasets and Benchmarks commonly used in LLM-MA studies. “/” denotes the unavailability of data link.

tion paradox in an information marketplace, reveals improved decision-making and answer quality when agents temporarily access information before purchase. [Li *et al.*, 2023g] presents an LLM-MA framework for financial trading, emphasizing a layered memory system, debate mechanisms, and individualized trading characters, thereby fortifying decision-making robustness. [Zhao *et al.*, 2023] utilizes LLM-based Agents to simulate a virtual town with restaurant and customer agents, yielding insights aligned with sociological and economic theories. These studies collectively illuminate the broad spectrum of applications and advancements in employing LLMs for diverse economic simulation scenarios.

4.2.5 Recommender Systems

The use of the LLM-MA in recommender systems is similar to that in psychology since studies in both fields involve the consideration of extrinsic and intrinsic human factors such as cognitive processes and personality [Lex and Schedl, 2022]. One way to use LLM-MA in recommender systems is to directly introduce items to multiple LLM-based agents within diverse traits and conduct statistics of the preferences of different agents. Another way is to treat both users and items as agents and the user-item communication as interactions, simulating the preference propagation. To bridge the gap between offline metrics and real-world performance in recommendation systems, Agent4Rec [Zhang *et al.*, 2023a] introduces a simulation platform based on LLM-MA. 1000 generative agents are initialized with the MovieLens-1M dataset to simulate complex user interactions in a recommendation environment. Agent4Rec shows that LLM-MA can effectively mimic real user preferences and behaviors, provide insights into phenomena like the filter bubble effect, and help uncover

causal relationships in recommendation tasks. In Agent4Rec work, agents are used to simulate users and they do not communicate with each other. Different from Agent4Rec work, [Zhang *et al.*, 2023e] treats both users and items as agents, optimizing them collectively to reflect and adjust to real-world interaction disparities. This work emphasizes simulating user-item interactions and propagates preferences among agents, capturing the collaborative filtering essence.

4.2.6 Policy Making

Similar to simulations in gaming and economic scenarios, Policy Making requires strong decision-making capabilities to realistic and dynamic complex problems. LLM-MA can be used to simulate the policy making via simulating a virtual government or simulating the impact of various policies on different communities. These simulations provide valuable insights into how policies are formulated and their potential effects, aiding policymakers in understanding and anticipating the consequences of their decisions [Farmer and Axtell, 2022]. The research outlined in [Xiao *et al.*, 2023] is centered on simulating a township water pollution crisis. It simulated a town located on an island including a demographic structure of different agents and township head and advisor. Within the water pollution crisis simulation, this work provides an in-depth analysis of how a virtual government entity might respond to such a public administration challenge and how information transfer in the social network in this crisis. [Hua *et al.*, 2023] introduces WarAgent to simulate key historical conflicts and provides insights for conflict resolution and understanding, with potential applications in preventing future international conflicts.

4.2.7 Disease Propagation Simulation

Leveraging the societal simulation capabilities of LLM-MA can also be used to simulate disease propagation. The most recent study in [Williams *et al.*, 2023] delves into the use of LLM-MA in simulating disease spread. The research showcases through various simulations how these LLM-based agents can accurately emulate human responses to disease outbreaks, including behaviors like self-quarantine and isolation during heightened case numbers. The collective behavior of these agents mirrors the complex patterns of multiple waves typically seen in pandemics, eventually stabilizing into an endemic state. Impressively, their actions contribute to the attenuation of the epidemic curve. [Ghaffarzadegan *et al.*, 2023] also discusses the epidemic propagation simulation and decomposes the simulation into two parts: the Mechanistic Model which represents the information or propagation of the virus and the Decision-Making Model which represents the agents’ decision-making process when facing the virus.

5 Implementation Tools and Resources

5.1 Multi-Agents Framework

We provide a detailed introduction to three open-source multi-agent frameworks: MetaGPT [Hong *et al.*, 2023], CAMEL [Li *et al.*, 2023b], and Autogen [Wu *et al.*, 2023a]. They are all frameworks that utilize language models for complex task-solving with a focus on multi-agent collaboration, but they differ in their approaches and applications.

MetaGPT is designed to embed human workflow processes into the operation of language model agents, thereby reducing the hallucination problem that often arises in complex tasks. It does this by encoding Standard Operating Procedures into the system and using an assembly line approach to assign specific roles to different agents.

CAMEL, or Communicative Agent Framework, is oriented towards facilitating autonomous cooperation among agents. It uses a novel technique called inception prompting to guide conversational agents towards fulfilling tasks that are consistent with human objectives. This framework also serves as a tool for generating and studying conversational data, helping researchers understand how communicative agents behave and interact.

AutoGen is a versatile framework that allows for the creation of applications using language models. It is distinctive for its high level of customization, enabling developers to program agents using both natural language and code to define how these agents interact. This versatility enables its use in diverse fields, from technical areas such as coding and mathematics to consumer-focused sectors like entertainment.

More recently, [Chen *et al.*, 2023c; Chen *et al.*, 2023a] introduce frameworks for dynamic multi-agent collaboration, while [Zhou *et al.*, 2023a; Li *et al.*, 2023h; Xie *et al.*, 2023] present platforms and libraries for building autonomous agents, emphasizing their adaptability in task-solving and social simulations.

5.2 Datasets and Benchmarks

We summarize commonly used datasets or benchmarks for LLM-MA study in Table 2. We observe that different re-

search applications use different datasets and benchmarks. In the Problem solving scenarios, most datasets and benchmarks are used to evaluate the planning and reasoning capabilities by Multiple agents cooperation or debate. In World Simulation scenarios, datasets and benchmarks are used to evaluate the alignment between the simulated world and real-world or analyze the behaviors of different agents. However, in certain research applications like Science Team operations for experiments and economic modeling, there is still a need for comprehensive benchmarks. The development of such benchmarks would greatly enhance the ability to gauge the success and applicability of LLM-MA in these complex and dynamic fields.

6 Challenges and Opportunities

Studies of LLM-MA frameworks and applications are advancing rapidly, giving rise to numerous challenges and opportunities. We identified several critical challenges and potential areas for future study.

6.1 Advancing into Multi-Modal Environment

Most previous work on LLM-MA has been focused on text-based environments, excelling in processing and generating text. However, there is a notable lack in multi-modal settings, where agents would interact with and interpret data from multiple sensory inputs and generate multiple outputs such as images, audio, video, and physical actions. Integrating LLMs into multi-modal environments presents additional challenges, such as processing diverse data types and enabling agents to understand each other and respond to more than just textual information.

6.2 Addressing Hallucination

The hallucination problem is a significant challenge in LLMs and single LLM-based Agent systems. It refers to the phenomenon where the model generates text that is factually incorrect [Huang *et al.*, 2023b]. However, this problem takes on an added layer of complexity in a multi-agent setting. In such scenarios, one agent’s hallucination can have a cascading effect. This is due to the interconnected nature of multi-agent systems, where misinformation from one agent can be accepted and further propagated by others in the network. Therefore, detecting and mitigating hallucinations in LLM-MA is not just a crucial task but also presents a unique set of challenges. It involves not only correcting inaccuracies at the level of individual agents but also managing the flow of information between agents to prevent the spread of these inaccuracies throughout the system.

6.3 Acquiring Collective Intelligence

In traditional multi-agent systems, agents often use reinforcement learning to learn from offline training datasets. However, LLM-MA systems mainly learn from instant feedback, such as interactions with the environment or humans, as we discussed in Section 3. This learning style requires a reliable interactive environment and it would be tricky to design such an interactive environment for many tasks, limiting the scalability of LLM-MA systems. Moreover, the prevailing

approaches in current research involve employing Memory and Self-Evolution techniques to adjust agents based on feedback. While effective for individual agents, these methods do not fully capitalize on the potential collective intelligence of the agent network. They adjust agents in isolation, overlooking the synergistic effects that can emerge from coordinated multi-agent interactions. Hence, jointly adjusting multiple agents and achieving optimal collective intelligence is still a critical challenge for LLM-MA.

6.4 Scaling Up LLM-MA Systems

LLM-MA systems are composed of a number of individual LLM-based agents, posing a significant challenge of scalability regarding the number of agents. From the computational complexity perspective, each LLM-based agent, typically built on large language models like GPT-4, demands substantial computational power and memory. Scaling up the number of these agents in an LLM-MA system significantly increases resource requirements. In scenarios with limited computational resource, it would be challenging to develop these LLM-MA systems.

Additionally, as the number of agents in an LLM-MA system increases, additional complexities and research opportunities emerge, particularly in areas like efficient agent coordination, communication, and understanding the scaling laws of multi-agents. For instance, with more LLM-based agents, the intricacy of ensuring effective coordination and communication rises significantly. As highlighted in [Dibia, 2023], designing advanced Agents Orchestration methodologies is increasingly important. These methodologies aim to optimize agents workflows, task assignments tailored to different agents, and communication patterns across agents such as communication constraints between agents. Effective Agents Orchestration facilitates harmonious operation among agents, minimizing conflicts and redundancies. Additionally, exploring and defining the scaling laws that govern the behavior and efficiency of multi-agent systems as they grow larger remains an important area of research. These aspects highlight the need for innovative solutions to optimize LLM-MA systems, making them both effective and resource-efficient.

6.5 Evaluation and Benchmarks

We have summarized the datasets and benchmarks currently available for LLM-MA in Table 2. This is a starting point, and far from being comprehensive. We identify two significant challenges in evaluating LLM-MA systems and benchmarking their performance against each other. Firstly, as discussed in [Xu *et al.*, 2023a], much of the existing research focuses on evaluating individual agents’ understanding and reasoning within narrowly defined scenarios. This focus tends to overlook the broader and more complex emergent behaviors that are integral to multi-agent systems. Secondly, there is a notable shortfall in the development of comprehensive benchmarks across several research domains, such as Science Team for Experiment Operations, Economic analysis, and Disease propagation simulation. This gap presents an obstacle to accurately assessing and benchmarking the full capabilities of LLM-MA systems in these varied and crucial fields.

6.6 Applications and Beyond

The potential of LLM-MA systems extends far beyond their current applications, holding great promise for advanced computational problem-solving in fields such as finance, education, healthcare, environmental science, urban planning and so on. As we have discussed, LLM-MA systems possess the capability to tackle complex problems and simulate various aspects of the real world. While the current role-playing capabilities of LLMs may have limitations, ongoing advancements in LLM technology suggest a bright future. It is anticipated to have more sophisticated methodologies, applications, datasets, and benchmarks tailored for diverse research fields. Furthermore, there are opportunities to explore LLM-MA systems from various theoretical perspectives, such as Cognitive Science [Sumers *et al.*, 2023], Symbolic Artificial Intelligence, Cybernetics, Complex Systems, and Collective Intelligence. Such a multi-faceted approach could contribute to a more comprehensive understanding and innovative applications in this rapidly evolving field.

7 Conclusion

LLM-based Multi-Agents have shown inspiring collective intelligence and rapidly garnered increasing interest among researchers. In this survey, we first systematically review the development of LLM-MA systems by positioning, differentiating, and connecting them from various aspects, regarding the agents-environment interface, the characterization of agents by LLMs, the strategies for managing agent communication and the paradigms for capability acquisition. We also summarized LLM-MA applications for problem-solving and world simulation. By also highlighting the commonly used datasets and benchmarks and discussing challenges and future opportunities, we hope that this survey can serve as a useful resource for researchers across various research fields, inspiring future research to explore the potential of LLM-based Multi-Agents.

References

- [Agashe *et al.*, 2023] Saaket Agashe, Yue Fan, and Xin Eric Wang. Evaluating multi-agent coordination abilities in large language models, 2023.
- [Aher *et al.*, 2023] Gati Aher, Rosa I. Arriaga, and Adam Tauman Kalai. Using large language models to simulate multiple humans and replicate human subject studies, 2023.
- [Akata *et al.*, 2023] Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. Playing repeated games with large language models. *arXiv preprint arXiv:2305.16867*, 2023.
- [Anonymous, 2023] Anonymous. Rethinking the buyer’s inspection paradox in information markets with language agents. In *Submitted to The Twelfth International Conference on Learning Representations*, 2023. under review.
- [Chan *et al.*, 2023] Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. Chateval: Towards better llm-based evaluators through multi-agent debate, 2023.

- [Chen *et al.*, 2023a] 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.
- [Chen *et al.*, 2023b] Huaben Chen, Wenkang Ji, Lufeng Xu, and Shiyu Zhao. Multi-agent consensus seeking via large language models. *arXiv preprint arXiv:2310.20151*, 2023.
- [Chen *et al.*, 2023c] Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, et al. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors in agents. *arXiv preprint arXiv:2308.10848*, 2023.
- [Chen *et al.*, 2023d] Yongchao Chen, Jacob Arkin, Yang Zhang, Nicholas Roy, and Chuchu Fan. Scalable multi-robot collaboration with large language models: Centralized or decentralized systems? *arXiv preprint arXiv:2309.15943*, 2023.
- [Cobbe *et al.*, 2021] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- [Dasgupta *et al.*, 2023] Ishita Dasgupta, Christine Kaeser-Chen, Kenneth Marino, Arun Ahuja, Sheila Babayan, Felix Hill, and Rob Fergus. Collaborating with language models for embodied reasoning. *arXiv preprint arXiv:2302.00763*, 2023.
- [Dibia, 2023] Victor Dibia. Multi-agent llm applications — a review of current research, tools, and challenges. <https://newsletter.victordibia.com/p/multi-agent-llm-applications-a-review>, 2023.
- [Dong *et al.*, 2023a] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. A survey on in-context learning, 2023.
- [Dong *et al.*, 2023b] Yihong Dong, Xue Jiang, Zhi Jin, and Ge Li. Self-collaboration code generation via chatgpt, 2023.
- [Du *et al.*, 2023] Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multi-agent debate, 2023.
- [Fan *et al.*, 2023] Caoyun Fan, Jindou Chen, Yaohui Jin, and Hao He. Can large language models serve as rational players in game theory? a systematic analysis. *arXiv preprint arXiv:2312.05488*, 2023.
- [Farmer and Axtell, 2022] J. Doynne Farmer and Robert L. Axtell. Agent-Based Modeling in Economics and Finance: Past, Present, and Future. INET Oxford Working Papers 2022-10, Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, June 2022.
- [Gao *et al.*, 2023a] Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. S³: Social-network simulation system with large language model-empowered agents. *arXiv preprint arXiv:2307.14984*, 2023.
- [Gao *et al.*, 2023b] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, 2023.
- [Geva *et al.*, 2021] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies, 2021.
- [Ghaffarzadegan *et al.*, 2023] Navid Ghaffarzadegan, Aritra Majumdar, Ross Williams, and Niyousha Hosseinichimeh. Generative agent-based modeling: Unveiling social system dynamics through coupling mechanistic models with generative artificial intelligence. *arXiv preprint arXiv:2309.11456*, 2023.
- [Gong *et al.*, 2023] Ran Gong, Qiuyuan Huang, Xiaojian Ma, Hoi Vo, Zane Durante, Yusuke Noda, Zilong Zheng, Song-Chun Zhu, Demetri Terzopoulos, Li Fei-Fei, et al. Mindagent: Emergent gaming interaction. *arXiv preprint arXiv:2309.09971*, 2023.
- [Guo *et al.*, 2023] Taicheng Guo, Kehan Guo, Zhengwen Liang, Zhichun Guo, Nitesh V Chawla, Olaf Wiest, Xi-angliang Zhang, et al. What indeed can gpt models do in chemistry? a comprehensive benchmark on eight tasks. *arXiv preprint arXiv:2305.18365*, 2023.
- [Hendrycks *et al.*, 2020] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- [Hong *et al.*, 2023] Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, et al. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*, 2023.
- [Horton, 2023] John J Horton. Large language models as simulated economic agents: What can we learn from homo silicus? Technical report, National Bureau of Economic Research, 2023.
- [Hua *et al.*, 2023] Wenyue Hua, Lizhou Fan, Lingyao Li, Kai Mei, Jianchao Ji, Yingqiang Ge, Libby Hemphill, and Yongfeng Zhang. War and peace (waragent): Large language model-based multi-agent simulation of world wars, 2023.
- [Huang *et al.*, 2023a] Dong Huang, Qingwen Bu, Jie M. Zhang, Michael Luck, and Heming Cui. Agentcoder: Multi-agent-based code generation with iterative testing and optimisation, 2023.
- [Huang *et al.*, 2023b] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al.

- A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*, 2023.
- [Kaiya et al., 2023] Zhao Kaiya, Michelangelo Naim, Jovana Kondic, Manuel Cortes, Jiaxin Ge, Shuying Luo, Guangyu Robert Yang, and Andrew Ahn. Lyfe agents: Generative agents for low-cost real-time social interactions. *arXiv preprint arXiv:2310.02172*, 2023.
- [Khot et al., 2023] Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. Decomposed prompting: A modular approach for solving complex tasks, 2023.
- [Kovač et al., 2023] Grgur Kovač, Rémy Portelas, Peter Ford Dominey, and Pierre-Yves Oudeyer. The socialai school: Insights from developmental psychology towards artificial socio-cultural agents. *arXiv preprint arXiv:2307.07871*, 2023.
- [Lewis et al., 2021] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks, 2021.
- [Lex and Schedl, 2022] Elisabeth Lex and Markus Schedl. Psychology-informed recommender systems: A human-centric perspective on recommender systems. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval, CHIIR '22*, page 367–368, New York, NY, USA, 2022. Association for Computing Machinery.
- [Li et al., 2023a] Chao Li, Xing Su, Chao Fan, Haoying Han, Cong Xue, and Chunmo Zheng. Quantifying the impact of large language models on collective opinion dynamics. *arXiv preprint arXiv:2308.03313*, 2023.
- [Li et al., 2023b] Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. Camel: Communicative agents for” mind” exploration of large scale language model society. *arXiv preprint arXiv:2303.17760*, 2023.
- [Li et al., 2023c] Huao Li, Yu Quan Chong, Simon Stepputtis, Joseph Campbell, Dana Hughes, Michael Lewis, and Katia Sycara. Theory of mind for multi-agent collaboration via large language models, 2023.
- [Li et al., 2023d] Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. Api-bank: A comprehensive benchmark for tool-augmented llms, 2023.
- [Li et al., 2023e] Nian Li, Chen Gao, Yong Li, and Qingmin Liao. Large language model-empowered agents for simulating macroeconomic activities, 2023.
- [Li et al., 2023f] Siyu Li, Jin Yang, and Kui Zhao. Are you in a masquerade? exploring the behavior and impact of large language model driven social bots in online social networks. *arXiv preprint arXiv:2307.10337*, 2023.
- [Li et al., 2023g] Yang Li, Yangyang Yu, Haohang Li, Zhi Chen, and Khaldoun Khashanah. Tradinggpt: Multi-agent system with layered memory and distinct characters for enhanced financial trading performance, 2023.
- [Li et al., 2023h] Yuan Li, Yixuan Zhang, and Lichao Sun. Metaagents: Simulating interactions of human behaviors for llm-based task-oriented coordination via collaborative generative agents. *arXiv preprint arXiv:2310.06500*, 2023.
- [Liang et al., 2023] Zhenwen Liang, Wenhao Yu, Tanmay Rajpurohit, Peter Clark, Xiangliang Zhang, and Ashwin Kaylan. Let gpt be a math tutor: Teaching math word problem solvers with customized exercise generation. *arXiv preprint arXiv:2305.14386*, 2023.
- [Light et al., 2023a] Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. Avalonbench: Evaluating llms playing the game of avalon, 2023.
- [Light et al., 2023b] Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. From text to tactic: Evaluating llms playing the game of avalon. *arXiv preprint arXiv:2310.05036*, 2023.
- [Liu et al., 2023] Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization. *arXiv preprint arXiv:2310.02170*, 2023.
- [Ma et al., 2023] Zilin Ma, Yiyang Mei, and Zhaoyuan Su. Understanding the benefits and challenges of using large language model-based conversational agents for mental well-being support. *arXiv preprint arXiv:2307.15810*, 2023.
- [Mandi et al., 2023] Zhao Mandi, Shreeya Jain, and Shuran Song. Roco: Dialectic multi-robot collaboration with large language models. *arXiv preprint arXiv:2307.04738*, 2023.
- [Mao et al., 2023] Shaoguang Mao, Yuzhe Cai, Yan Xia, Wenshan Wu, Xun Wang, Fengyi Wang, Tao Ge, and Furu Wei. Alympics: Language agents meet game theory. *arXiv preprint arXiv:2311.03220*, 2023.
- [Moura, 2023] João Moura. Crewai. <https://github.com/joaomdmoura/crewAI>, 2023.
- [Mukobi et al., 2023] Gabriel Mukobi, Hannah Erlebach, Niklas Lauffer, Lewis Hammond, Alan Chan, and Jesse Clifton. Welfare diplomacy: Benchmarking language model cooperation. *arXiv preprint arXiv:2310.08901*, 2023.
- [Nascimento et al., 2023] Nathalia Nascimento, Paulo Alencar, and Donald Cowan. Self-adaptive large language model (llm)-based multiagent systems. In *2023 IEEE International Conference on Autonomic Computing and Self-Organizing Systems Companion (ACSOS-C)*, pages 104–109. IEEE, 2023.
- [Park et al., 2022] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18, 2022.

- [Park *et al.*, 2023] Joon Sung Park, Joseph C O’Brien, Carrie J Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*, 2023.
- [Qian *et al.*, 2023] Chen Qian, Xin Cong, Wei Liu, Cheng Yang, Weize Chen, Yusheng Su, Yufan Dang, Jiahao Li, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. Communicative agents for software development, 2023.
- [Ruan *et al.*, 2023] Jingqing Ruan, Yihong Chen, Bin Zhang, Zhiwei Xu, Tianpeng Bao, Guoqing Du, Shiwei Shi, Hangyu Mao, Ziyue Li, Xingyu Zeng, and Rui Zhao. Tptu: Large language model-based ai agents for task planning and tool usage, 2023.
- [Russell and Norvig, 2009] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall Press, USA, 3rd edition, 2009.
- [Shinn *et al.*, 2023] Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning, 2023.
- [Sumers *et al.*, 2023] Theodore R Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L Griffiths. Cognitive architectures for language agents. *arXiv preprint arXiv:2309.02427*, 2023.
- [Tang *et al.*, 2023] Xiangru Tang, Anni Zou, Zhuosheng Zhang, Yilun Zhao, Xingyao Zhang, Arman Cohan, and Mark Gerstein. Medagents: Large language models as collaborators for zero-shot medical reasoning, 2023.
- [Wang *et al.*, 2021] Zijie J. Wang, Dongjin Choi, Shenyu Xu, and Diyi Yang. Putting humans in the natural language processing loop: A survey, 2021.
- [Wang *et al.*, 2023a] Kuan Wang, Yadong Lu, Michael Santacrose, Yeyun Gong, Chao Zhang, and Yelong Shen. Adapting llm agents through communication, 2023.
- [Wang *et al.*, 2023b] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji-Rong Wen. A survey on large language model based autonomous agents, 2023.
- [Wang *et al.*, 2023c] Shenzhi Wang, Chang Liu, Zilong Zheng, Siyuan Qi, Shuo Chen, Qisen Yang, Andrew Zhao, Chaofei Wang, Shiji Song, and Gao Huang. Avalon’s game of thoughts: Battle against deception through recursive contemplation. *arXiv preprint arXiv:2310.01320*, 2023.
- [Wei *et al.*, 2022] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- [Weng, 2023] Lilian Weng. Llm powered autonomous agents. <https://lilianweng.github.io/posts/2023-06-23-agent/>, 2023.
- [Williams *et al.*, 2023] Ross Williams, Niyousha Hossainichimeh, Aritra Majumdar, and Navid Ghaffarzadegan. Epidemic modeling with generative agents. *arXiv preprint arXiv:2307.04986*, 2023.
- [Wooldridge and Jennings, 1995] Michael Wooldridge and Nicholas R. Jennings. Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10:115 – 152, 1995.
- [Wu *et al.*, 2023a] Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Shaokun Zhang, Erkang Zhu, Beibin Li, Li Jiang, Xiaoyun Zhang, and Chi Wang. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*, 2023.
- [Wu *et al.*, 2023b] Yuxiang Wu, Zhengyao Jiang, Akbik Khan, Yao Fu, Laura Ruis, Edward Grefenstette, and Tim Rocktäschel. Chatarena: Multi-agent language game environments for large language models. *GitHub repository*, 2023.
- [Xi *et al.*, 2023] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu, Xuanjing Huang, and Tao Gui. The rise and potential of large language model based agents: A survey, 2023.
- [Xiao *et al.*, 2023] Bushi Xiao, Ziyuan Yin, and Zixuan Shan. Simulating public administration crisis: A novel generative agent-based simulation system to lower technology barriers in social science research. *arXiv preprint arXiv:2311.06957*, 2023.
- [Xie *et al.*, 2023] Tianbao Xie, Fan Zhou, Zhoujun Cheng, Peng Shi, Luoxuan Weng, Yitao Liu, Toh Jing Hua, Junning Zhao, Qian Liu, Che Liu, et al. Openagents: An open platform for language agents in the wild. *arXiv preprint arXiv:2310.10634*, 2023.
- [Xiong *et al.*, 2023] Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. Examining inter-consistency of large language models collaboration: An in-depth analysis via debate, 2023.
- [Xu *et al.*, 2023a] Lin Xu, Zhiyuan Hu, Daquan Zhou, Hongyu Ren, Zhen Dong, Kurt Keutzer, See Kiong Ng, and Jiashi Feng. Magic: Investigation of large language model powered multi-agent in cognition, adaptability, rationality and collaboration, 2023.
- [Xu *et al.*, 2023b] Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. Exploring large language models for communication games: An empirical study on werewolf. *arXiv preprint arXiv:2309.04658*, 2023.
- [Xu *et al.*, 2023c] Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. Language agents with reinforcement learning for strategic play in the werewolf game. *arXiv preprint arXiv:2310.18940*, 2023.

- [Yao *et al.*, 2023] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models, 2023.
- [Yu *et al.*, 2023] Bangguo Yu, Hamidreza Kasaei, and Ming Cao. Co-navgpt: Multi-robot cooperative visual semantic navigation using large language models, 2023.
- [Zhang *et al.*, 2023a] An Zhang, Leheng Sheng, Yuxin Chen, Hao Li, Yang Deng, Xiang Wang, and Tat-Seng Chua. On generative agents in recommendation, 2023.
- [Zhang *et al.*, 2023b] Ceyao Zhang, Kaijie Yang, Siyi Hu, Zihao Wang, Guanghe Li, Yihang Sun, Cheng Zhang, Zhaowei Zhang, Anji Liu, Song-Chun Zhu, et al. Proagent: Building proactive cooperative ai with large language models. *arXiv preprint arXiv:2308.11339*, 2023.
- [Zhang *et al.*, 2023c] Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B Tenenbaum, Tianmin Shu, and Chuang Gan. Building cooperative embodied agents modularly with large language models. *arXiv preprint arXiv:2307.02485*, 2023.
- [Zhang *et al.*, 2023d] Jintian Zhang, Xin Xu, and Shumin Deng. Exploring collaboration mechanisms for llm agents: A social psychology view, 2023.
- [Zhang *et al.*, 2023e] Junjie Zhang, Yupeng Hou, Ruobing Xie, Wenqi Sun, Julian McAuley, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. Agentcf: Collaborative learning with autonomous language agents for recommender systems, 2023.
- [Zhao *et al.*, 2023] Qinlin Zhao, Jindong Wang, Yixuan Zhang, Yiqiao Jin, Kaijie Zhu, Hao Chen, and Xing Xie. Competeai: Understanding the competition behaviors in large language model-based agents, 2023.
- [Zheng *et al.*, 2023] Zhiling Zheng, Oufan Zhang, Ha L. Nguyen, Nakul Rampal, Ali H. Alawadhi, Zichao Rong, Teresa Head-Gordon, Christian Borgs, Jennifer T. Chayes, and Omar M. Yaghi. Chatgpt research group for optimizing the crystallinity of mofs and cofcs. *ACS Central Science*, 9(11):2161–2170, 2023.
- [Zhou *et al.*, 2023a] Wangchunshu Zhou, Yuchen Eleanor Jiang, Long Li, Jialong Wu, Tiannan Wang, Shi Qiu, Jintian Zhang, Jing Chen, Ruipu Wu, Shuai Wang, et al. Agents: An open-source framework for autonomous language agents. *arXiv preprint arXiv:2309.07870*, 2023.
- [Zhou *et al.*, 2023b] Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, and Maarten Sap. Sotopia: Interactive evaluation for social intelligence in language agents, 2023.
- [Ziems *et al.*, 2023] Caleb Ziems, Omar Shaikh, Zhehao Zhang, William Held, Jiaao Chen, and Diyi Yang. Can large language models transform computational social science? *Computational Linguistics*, pages 1–53, 2023.