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

Taicheng Guo[1] , Xiuying Chen[2] , Yaqi Wang[3*] , Ruidi Chang[4] , Shichao Pei[5] , Nitesh V. Chawla[1] , Olaf Wiest[1] , Xiangliang Zhang[1]
[1]University of Notre Dame
[2]King Abdullah University of Science and Technology
[3]Southern University of Science and Technology
[4]Independent Researcher
[5]University of Massachusetts Boston
{tguo2, ywang84, nchawla, owiest, xzhang33}@nd.edu, xiuying.chen@kaust.edu.sa,
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 their notable capabilities in planning and reasoning, LLMs have been utilized as autonomous agents for the automatic execution of various tasks. Recently, LLM-based agent systems have rapidly evolved from single-agent planning or decision-making to operating as multi-agent systems, enhancing their ability in complex problem-solving and world simulation. To offer an overview of this dynamic field, we present this survey to offer an in-depth discussion on the essential aspects and challenges of LLM-based multi-agent (LLM-MA) systems. Our objective is to provide readers with an in-depth understanding of these key points: the domains and settings where LLM-MA systems operate or simulate; the profiling and communication methods of these agents; and the means by which these agents develop their skills. For those interested in delving into this field, we also summarize the commonly used datasets or benchmarks. To keep researchers updated on the latest studies, we maintain an open-source GitHub repository (github.com/taichengguo/LLM_MultiAgents_Survey_Papers), dedicated to outlining the research of LLM-MA research.

1 Introduction

Large Language Models (LLMs) have recently shown remarkable potential in reaching a level of reasoning and planning capabilities comparable to humans. 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 have systematically summarized the progress of LLM-based agents [Xi *et al.*, 2023; Wang *et al.*, 2023b]. Recently, agent systems based on LLMs have rapidly evolved from using a single planning or decision-making agent to collaborating as multi-agent systems. Compared to systems using a single LLM-powered agent, LLM-based Multi-Agents (LLM-MA) 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 and their latent potential to specialize in specific tasks. Recent research has demonstrated promising results in utilizing LLM-MA 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]. Due to its interdisciplinary nature, this field 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 [Gao *et al.*, 2023b], as shown in Fig. 1, thus broadening the impact of LLM-MA 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-MA. Readers of this survey will gain a comprehensive overview of LLM-MA systems, including fundamental concepts, the latest research trends and applications. We recognize this field is in its early stages and is rapidly evolving with fresh methodologies and applications. We hope our survey will inspire further innovations in this field, as well as applications across a wide

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

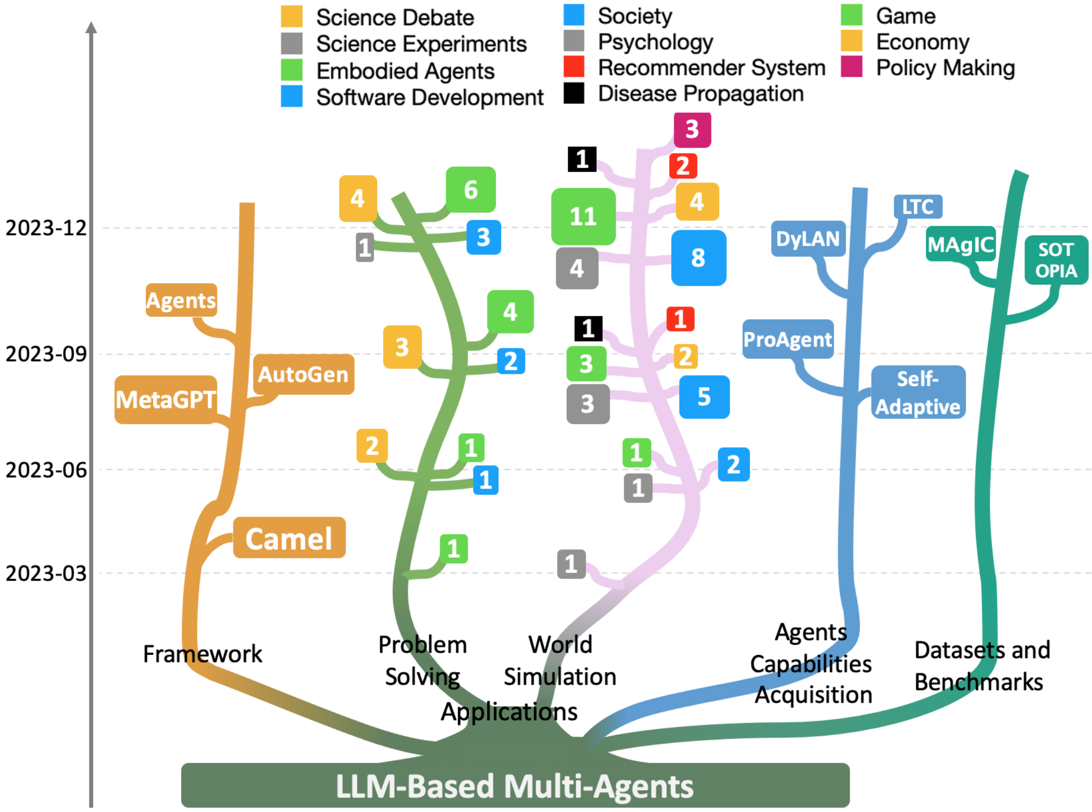

Figure 1: The rising trend of LLM-MA study. The work in *Problem Solving* and *World Simulation* is categorized by topics and papers within each category are counted every 3-month, as shown on leaves.

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 agent, guided by prompts, to break down complex tasks into smaller goals [Khot *et al.*, 2022], think through each part methodically (sometimes exploring multiple paths) [Yao *et al.*, 2023], or learn from feedback [Shinn *et al.*, 2023] to perform better decision-making on complex tasks. This capability enhances the autonomy of 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.*, 2022] as short memory or external vector database [Lewis *et* *al.*, 2020] 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

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(a). 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. Examples of these environments include software development [Hong *et al.*, 2023], gaming [Mao *et al.*, 2023], 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 environment sets the game's framework, including discussion periods, voting mechanics, reward rules, etc. Agents, such as werewolves and the Seer, perform specific actions like killing or checking roles. Then agents receive feedback from the game setting, which guides agents in adjusting their strategies. Referring to the way in which agents interact with the environment as the *Agents-Environment Interface*, we categorize the existing interfaces into three types, *Sandbox, Physcial*, 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], and gaming (game rules as simulated environment) [Mao *et al.*, 2023]. 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 making sandwiches and packing groceries, 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] utilize multiple agents to debate a question to reach a consensus. These applications primarily focus on agents communication and do not rely on the environment.

3.2 Agents Profiling

In LLM-MA systems, agents generally assume distinct roles, each with comprehensive descriptions encompassing characteristics, capabilities, and constraints. For instance, 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 strategies to fulfill their roles effectively. These profiles are crucial for defining the agents' interactions 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.

3.3 Agents Communication

The communication between agents in LLM-MA systems is the critical infrastructure supporting collective intelligence. We dissect agent communication from: 1) *Communication Paradigms*: the styles and methods of interaction between agents; 2) *Communication Structure*: the organization 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. 2(b) shows four typical communication structures in LLM-MA systems. ***Layered*** is structured hierarchically, with agents at each level having distinct roles and primarily interacting within adjacent layers. [Liu *et al.*, 2023] introduces a framework called Dynamic LLM-Agent Network, which organizes agents in a multi-layered feed-forward network. ***Decentralized*** operates

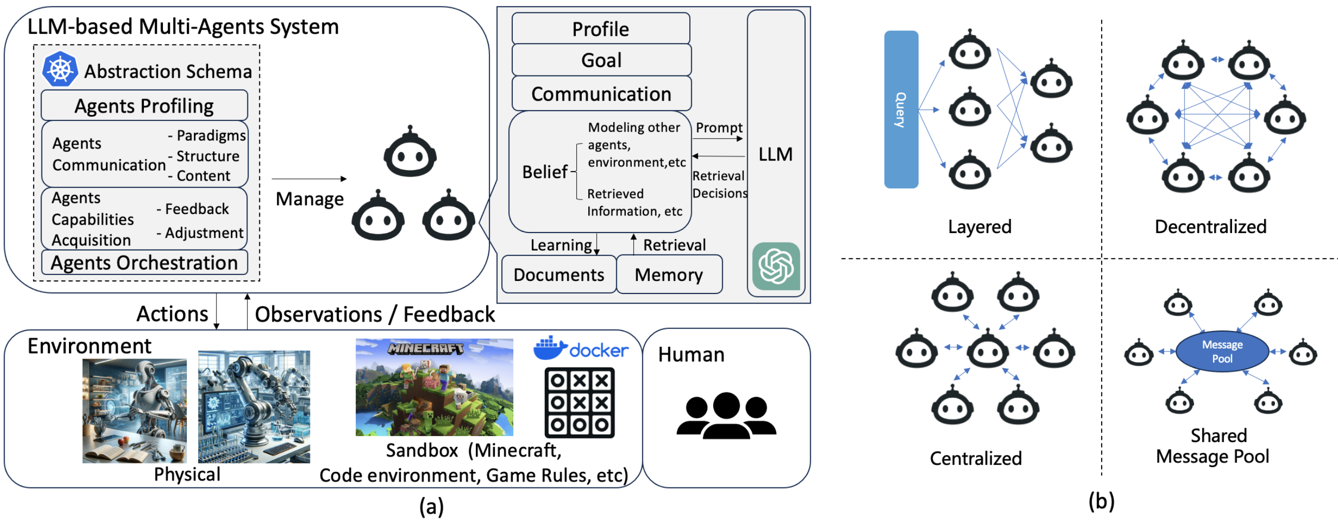

Figure 2: (a) The Architecture of LLM-MA Systems. (b) The Agent Communication Structure.

on a peer-to-peer network, where agents directly communicate with each other. ***Centralized*** involves a central agent 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] which 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

LLM-MA systems enable agents to learn and evolve dynamically for developing their skills. 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 scenarios like science debates where agents learn to critically evaluate and refine the conclusions through communications, as well as in game simulation where agents learn to refine strategies through previous interactions between other agents. **3) Human Feedback** comes directly from humans and is crucial for aligning the multi-agent system with human preferences. This kind of feedback is widely used in most "Human-in-the-loop" applications [Wang *et al.*, 2021]. **4) None**. There are cases with no feedback for the agents. This often happens in world simulation. 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 Retrieval.** 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]. **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, e.g., 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-agent 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. **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 involving a larger number of agents, the escalating complexity of managing agents has been a critical problem. *Agents Orchestration* emerged as a pivotal challenge [Moura, 2023; Dibia, 2023]. We will further discuss this topic in Section 6.4.

4 Applications

We summarize two kinds of applications in Table 1: **Problem Solving** and **World Simulation** as discussed below.

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.

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 workflow of 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. Recent LLM-MA work for software development includes the Role-play [Li *et al.*, 2023b; Dong *et al.*, 2023], Workflow Procedures such as Standardized Operating Procedures (SOPs) [Hong *et al.*, 2023], Waterfall Model [Qian *et al.*, 2023], effective test case generation, execution, and optimization [Huang *et al.*, 2023a].

Embodied Agents Most embodied agents applications inherently utilize multiple robots working together to perform complex real-world planning and manipulation tasks with heterogeneous robot capabilities. Hence, LLM-MA can be used to empower robots with different capabilities, especially planning and communication capabilities, and cooperate to solve real-world physical tasks, such as navigation [Yu *et al.*, 2023; Zhang *et al.*, 2023c], food making, grocery packing, and cabinet arrangement [Mandi *et al.*, 2023]. Furthermore, some work aims to improve the efficiency of robots' cooperation including evaluating the effectiveness of communication structures [Chen *et al.*, 2023d], and building robots' consensus-seeking with negotiation framework [Chen *et al.*, 2023b].

Science Experiments Multiple agents who play as different specialists can also be used to form a science team for science experiments. Unlike previous methods, this scenario emphasizes human oversight due to the high expenses of science experiments and the hallucination of LLM agents. Human experts are at the center of these agents to process information of agents and give feedback to the agents. [Zheng *et al.*, 2023] utilizes multiple agents, each focusing on specific tasks including strategy planning, literature search, robotic operations, etc. All these agents interact with humans to work collaboratively to optimize the synthesis process of complex materials.

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 [Hendrycks *et al.*, 2020], StrategyQA [Geva *et al.*, 2021], medical diagnosis [Tang *et al.*, 2023]. The main idea is that each agent initially offers its analysis of a problem, which is then followed by a joint debating process. Through multiple rounds of debate, the agents converge on a consensus answer. Through debating, the reasoning factuality [Du *et al.*, 2023] and inter-consistency between different LLMs [Xiong *et al.*, 2023] can be improved.

4.2 LLM-MA for World Simulation

Another mainstream application of LLM-MA is world simulation, which spans a diverse range of fields including social sciences, gaming, 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 viewpoints. The environment of world simulation is usually crafted to reflect the specific scenario being simulated, with agents designed in various profiles to match this context. Unlike problem-solving which focuses 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.

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 [Park *et al.*, 2023], more intimate settings, and larger [Park *et al.*, 2022], more intricate ones [Gao *et al.*, 2023a]. Beyond simulation, recent studies [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.

Gaming LLM-MA is well-suited for creating gaming environments since it enables the development of controlled, scalable, and dynamic settings that closely mimic human interactions, making it ideal for testing game theory hypotheses like

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 *et al.*, 2023]	Sandbox	Pre-defined, Model-Generated	CTO, programmer	Cooperative	Layered	Environment, Agent interaction, Human	Memory, Self-Evolution
			[Hong *et al.*, 2023]	Sandbox	Pre-defined	Product Manager, Engineer	Cooperative	Layered, Shared Message Pool	Environment, Agent interaction, Human	Memory, Self-Evolution
			[Dong *et al.*, 2023]	Sandbox	Pre-defined, Model-Generated	Analyst, coder	Cooperative	Layered	Environment, Agent interaction	Memory, Self-Evolution
	Embodied Agents	Multi-robot planning	[Chen *et al.*, 2023d]	Sandbox, Physical	Pre-defined	Robots	Cooperative	Centralized, Decentralized	Environment, Agent interaction	Memory
		Multi-robot collaboration	[Mandi *et al.*, 2023]	Sandbox, Physical	Pre-defined	Robots	Cooperative	Decentralized	Environment, Agent interaction	Memory
		Multi-Agents cooperation	[Zhang *et al.*, 2023c]	Sandbox	Pre-defined	Robots	Cooperative	Decentralized	Environment, Agent interaction	Memory
	Science Experiments	Optimization of MOF	[Zheng *et al.*, 2023]	Physical	Pre-defined	Strategy planers, literature collector, coder	Cooperative	Centralized	Environment, Human	Memory
	Science Debate	Improving Factuality	[Du *et al.*, 2023]	None	Pre-defined	Agents	Debate	Decentralized	Agent interaction	Memory
		Examining, Inter-Consistency	[Xiong *et al.*, 2023]	None	Pre-defined	Proponent, Opponent, Judge	Debate	Centralized, Decentralized	Agent interaction	Memory
		Evaluators for debates	[Chan *et al.*, 2023]	None	Pre-defined	Agents	Debate	Centralized, Decentralized	Agent interaction	Memory
		Multi-Agents for Medication	[Tang *et al.*, 2023]	None	Pre-defined	Cardiology, Surgery	Debate, Cooperative	Centralized, Decentralized	Agent interaction	Memory
World Simulation	Society	Modest Community (25 persons)	[Park *et al.*, 2023]	Sandbox	Model-generated	Pharmacy, shopkeeper	-	-	Environment, Agent interaction	Memory
		Online community (1000 persons)	[Park *et al.*, 2022]	None	Pre-defined, Model-generated	Camping, fishing	-	-	Agent interaction	Dynamic Generation
		Emotion propagation	[Gao *et al.*, 2023a]	None	Pre-defined, Model-generated	Real-world user	-	-	Agent interaction	Memory
		Real-time social interactions	[Kaiya *et al.*, 2023]	Sandbox	Pre-defined	Real-world user	-	-	Environment, Agent interaction	Memory
		Opinion dynamics	[Li *et al.*, 2023a]	None	Pre-defined	NIN, NINL, NIL	-	-	Agent interaction	Memory
	Gaming	WereWolf	[Xu *et al.*, 2023b] [Xu *et al.*, 2023c]	Sandbox	Pre-defined	Seer, werewolf, villager	Cooperative, Debate, Competitive	Decentralized	Environment, Agent interaction	Memory
		Avalon	[Light *et al.*, 2023b] [Wang *et al.*, 2023c]	Sandbox	Pre-defined	Servant, Merlin, Assassin	Cooperative, Debate, Competitive	Decentralized	Environment, Agent interaction	Memory
		Welfare Diplomacy	[Mukobi *et al.*, 2023]	Sandbox	Pre-defined	Countries	Cooperative, Competitive	Decentralized	Environment, Agent interaction	Memory
	Psychology	Human behavior Simulation	[Aher *et al.*, 2023]	Sandbox	Pre-defined	Humans	-	-	Agent interaction	Memory
		Collaboration Exploring	[Zhang *et al.*, 2023d]	None	Pre-defined	Agents	Cooperative, Debate	Decentralized	Agent interaction	Memory
	Economy	Macroeconomic simulation	[Li *et al.*, 2023e]	None	Pre-defined, Model-generated	Labor	Cooperative	Decentralized	Agent interaction	Memory
		Information Marketplaces	[Weiss *et al.*, 2024]	Sandbox	Pre-defined, Data-Derived	Buyer	Cooperative, Competitive	Decentralized	Environment, Agent interaction	Memory
		Improving financial trading	[Li *et al.*, 2023g]	Physical	Pre-defined	Trader	Debate	Decentralized	Environment, Agent interaction	Memory
		Economic theories	[Zhao *et al.*, 2023]	Sandbox	Pre-defined, Model-Generated	Restaurant, Customer	Competitive	Decentralized	Environment, Agent interaction	Memory, Self-Evolution
	Recommender Systems	Simulating user behaviors	[Zhang *et al.*, 2023a]	Sandbox	Data-Derived	Users from MovieLens-1M	-	-	Environment	Memory
		Simulating user-item interactions	[Zhang *et al.*, 2023e]	Sandbox	Pre-defined, Data-Derived	User Agents Item Agents	Cooperative	Decentralized	Environment, Agent interaction	Memory
	Policy Making	Public Administration	[Xiao *et al.*, 2023]	None	Pre-defined	Residents	Cooperative	Decentralized	Agent interaction	Memory
		War Simulation	[Hua *et al.*, 2023]	None	Pre-defined	Countries	Competitive	Decentralized	Agent interaction	Memory
	Disease	Human Behaviors to epidemics	[Ghaffarzadegan *et al.*, 2023]	Sandbox	Pre-defined, Model-Generated	Conformity traits	Cooperative	Decentralized	Environment, Agent interaction	Memory
		Public health	[Williams *et al.*, 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.

persuasion, leadership, and so on [Xu *et al.*, 2023b]. Most game simulations rely on language communication, offering a sandbox environment within game settings. Currently, some work focuses on designing new decision-making models by simulating strategy games such as Avalon [Wang *et al.*, 2023c] and Werewolf [Xu *et al.*, 2023c]. Some work emphasizes exploring the agents' Theory of Mind [Li *et al.*, 2023c; Fan *et al.*, 2023] or psychology theories by simulating psychological games like the iterated Prisoner's Dilemma [Akata *et al.*, 2023]. Furthermore, some work [Xu *et al.*, 2023b; Mao *et al.*, 2023; Light *et al.*, 2023a; Mukobi *et al.*, 2023] focuses on building a comprehensive game simulation framework for further developing advanced simulation.

Psychology In psychology, one approach that aligns with societal simulations involves using multiple agents to simulate humans with various traits. Psychological theories are then applied to analyze the emergent behavioral patterns, providing insights into how individual psychological traits influence collective actions. Another approach involves directly applying psychological experiments to these agents, and analyzing their behaviors statistically. Each agent operates independently, without interacting with others. Through psychology simulation, current work focuses on understanding how agents with different traits can acquire, demonstrate, and evolve social skills such as joint attention, communication, and cultural or knowledge learning [Kovač *et al.*, 2023; Zhang *et al.*, 2023d; Aher *et al.*, 2023].

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

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

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, information, or set with pre-defined preferences, allowing for an exploration of their actions in 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 [Horton, 2023]. Several studies have demonstrated the diverse applications of LLM-MA in simulating economic scenarios, encompassing macroeconomic activities [Li *et al.*, 2023e], information marketplaces [Weiss *et al.*, 2024], financial trading [Li *et al.*, 2023g], and virtual restaurant operation simulation [Zhao *et al.*, 2023].

Recommender Systems The use of 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 agents within diverse traits and conduct statistics of the preferences of different agents to explore insights into phenomena like the filter bubble effect [Zhang *et al.*, 2023a]. Another way is to treat both users and items as agents and the user-item communication as interactions, simulating the preference propagation to explore the collaborative filtering essence [Zhang *et al.*, 2023e].

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 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 and [Hua *et al.*, 2023] introduces WarAgent to simulate key historical conflicts.

Disease Propagation Simulation Leveraging the societal simulation capabilities of LLM-MA can also be used to simulate disease propagation. Recent study [Williams *et al.*, 2023; Navid *et al.*, 2023] delves into the use of LLM-MA in simulating disease spread, showing how these agents can accurately emulate human decision-making processes to disease outbreaks.

5 Implementation Tools and Resources

5.1 Multi-Agents Frameworks

MetaGPT [Hong *et al.*, 2023], CAMEL [Li *et al.*, 2023b], and Autogen [Wu *et al.*, 2023] are the three most popular open-source multi-agent frameworks. MetaGPT emphasizes embedding human workflow processes into the operation of language model agents and using an assembly line approach to assign specific roles to different agents. CAMEL focuses on the role-playing and "mind" exploration. It serves as a tool for generating conversational data that help to understand how communicative agents behave and interact. AutoGen focuses on the concept of "customizable and conversable agents". It enables developers to define/program agents using both natural language and code, making it applicable 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 frameworks for emphasizing the agents' 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. For *Problem Solving*, most datasets are used to evaluate the planning and reasoning capabilities by multiple agents' cooperation or debate. In *World Simulation*, datasets 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 and economic modeling, there is still a need for comprehensive benchmarks. The development of such benchmarks would greatly benefit gauging the applicability of LLM-MA in complex and dynamic fields.

6 Challenges and Opportunities

Studies of LLM-MA 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 LLM-MA research focused on text-based settings, excelling in processing and generating text. However, there is a notable lack in multi-modal settings, where agents need to handle and respond to various inputs like images, audio, and video. Integrating LLMs into multi-modal settings presents additional challenges in processing diverse data types and equipping agents to interpret and react beyond text.

6.2 Addressing Hallucination

The hallucination problem, where LLMs produce factually incorrect text, is a significant challenge [Huang *et al.*, 2023b]. In LLM-MA, this challenge is amplified as one agent's hallucination can have a cascading effect. Addressing this in LLM-MA is critical and complex, requiring not only the correction of individual errors but also control over information flow between agents to stop the spread of inaccuracies.

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 as outlined in Section 3. This learning style requires a reliable interactive environment, which is challenging for LLM-MA scenarios. In addition, current research mainly uses *memory retrieval* and *self-evolution* for individual agent adjustment. However, they overlook the synergistic effects that can emerge from coordinated multi-agent interactions. Hence, jointly adjusting multiple agents and achieving optimal collective intelligence remains a critical challenge in LLM-MA.

6.4 Scaling Up LLM-MA Systems

Scaling up LLM-MA systems poses challenges in computational demand and memory, as each agent requires significant resources. Additionally, as the number of agents in an LLM-MA system increases, additional complexities and research opportunities emerge such as efficient agent coordination, communication, and exploring the scaling laws of multi-agents that govern the behavior and efficiency of LLM-MA systems as they grow larger. As highlighted in [Dibia, 2023], designing advanced *agents orchestration* methodologies is increasingly important. These aim to optimize workflows, task assignments tailored to different agents, and communication patterns across agents such as communication constraints, facilitating harmonious operation among agents, and minimizing conflicts and redundancies. 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 identify two major challenges in benchmarking LLM-MA. Firstly, as noted in [Xu *et al.*, 2023a], most existing research evaluates individual agents' capabilities in narrow scenarios, often overlooking the broader and complex emergent behaviors crucial in LLM-MA systems. Secondly, as outlined in Section 5, there is a notable lack of comprehensive benchmarks across several domains, hindering the accurate assessment of LLM-MA's full potential in these important fields.

6.6 Applications and Beyond

LLM-MA systems can address complex issues and simulate real-world aspects. While current LLMs may have limits, technological progress points to a promising future with more sophisticated methods, applications, datasets, and benchmarks for various fields. Furthermore, exploring LLM-MA systems through lenses like cognitive science [Sumers *et al.*, 2023], symbolic AI, cybernetics, complex systems, and collective intelligence could lead to a deeper understanding and innovative applications in this fast-evolving area.

7 Conclusion

LLM-MA has shown inspiring collective intelligence, attracting growing research interest. In this survey, we systematically review LLM-MA systems by examining different aspects of its operational frameworks. We also summarized its applications in 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-MA.

Acknowledgments

This work was supported by the National Science Foundation (CHE–2202693) through the NSF Center for Computer-Assisted Synthesis (C-CAS).

References

[Agashe *et al.*, 2023] Saaket Agashe, Yue Fan, and Xin Eric Wang. Evaluating multi-agent coordination abilities in large language models. *arXiv preprint arXiv:2310.03903*, 2023.

[Aher *et al.*, 2023] Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR, 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.

[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. *arXiv preprint arXiv:2308.07201*, 2023.

[Chen *et al.*, 2023a] Guangyao Chen, Siwei Dong, et al. Autoagents: A framework for automatic agent generation. *arXiv preprint arXiv:2309.17288*, 2023.

[Chen *et al.*, 2023b] Huaben Chen, Wenkang Ji, et al. Multi-agent consensus seeking via large language models. *arXiv preprint arXiv:2310.20151*, 2023.

[Chen *et al.*, 2023c] Weize Chen, Yusheng Su, 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, Arkin, et al. Scalable multi-robot collaboration with large language models: Centralized or decentralized systems? *arXiv preprint arXiv:2309.15943*, 2023.

[Dibia, 2023] Victor Dibia. Multi-agent llm applications — a review of current research, tools, and challenges, 2023.

[Dong *et al.*, 2022] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*, 2022.

[Dong *et al.*, 2023] Yihong Dong, Xue Jiang, Zhi Jin, and Ge Li. Self-collaboration code generation via chatgpt. *arXiv preprint arXiv:2304.07590*, 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 multiagent debate. *arXiv preprint arXiv:2305.14325*, 2023.

[Fan *et al.*, 2023] Caoyun Fan, Jindou Chen, et al. 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. Doyne Farmer and Robert L. Axtell. Agent-Based Modeling in Economics and Finance: Past, Present, and Future. INET Oxford Working Papers 2022-10, June 2022.

[Gao *et al.*, 2023a] Chen Gao, Xiaochong Lan, et al. S^3: 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. *Transactions of the Association for Computational Linguistics*, 9:346–361, 2021.

[Hendrycks *et al.*, 2020] Dan Hendrycks, Collin Burns, et al. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.

[Hong *et al.*, 2023] Sirui Hong, Xiawu Zheng, 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. *arXiv preprint arXiv:2311.17227*, 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. *arXiv preprint arXiv:2312.13010*, 2023.

[Huang *et al.*, 2023b] Lei Huang, Weijiang Yu, 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, et al. Lyfe agents: Generative agents for low-cost real-time social interactions. *arXiv preprint arXiv:2310.02172*, 2023.

[Khot *et al.*, 2022] Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. Decomposed prompting: A modular approach for solving complex tasks. *arXiv preprint arXiv:2210.02406*, 2022.

[Kovač *et al.*, 2023] Kovač, Portelas, et al. The socialai school: Insights from developmental psychology towards artificial socio-cultural agents. *arXiv preprint arXiv:2307.07871*, 2023.

[Lewis *et al.*, 2020] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.

[Lex and Schedl, 2022] Elisabeth Lex and Markus Schedl. Psychology-informed recommender systems: A human-centric perspective on recommender systems. CHIIR '22. ACM, 2022.

[Li *et al.*, 2023a] Chao Li, Xing Su, et al. Quantifying the impact of large language models on collective opinion dynamics. *arXiv preprint arXiv:2308.03313*, 2023.

[Li *et al.*, 2023b] Guohao Li, Hammoud, et al. 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. *arXiv preprint arXiv:2310.10701*, 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. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.

[Li *et al.*, 2023e] Nian Li, Chen Gao, Yong Li, and Qingmin Liao. Large language model-empowered agents for simulating macroeconomic activities. *arXiv preprint arXiv:2310.10436*, 2023.

[Li *et al.*, 2023f] Siyu Li, Jin Yang, et al. 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. *arXiv preprint arXiv:2309.03736*, 2023.

[Li *et al.*, 2023h] Yuan Li, Yixuan Zhang, et al. Metaagents: Simulating interactions of human behaviors for llm-based task-oriented coordination via collaborative generative agents. *arXiv preprint arXiv:2310.06500*, 2023.

[Light *et al.*, 2023a] Jonathan Light, Min Cai, et al. From text to tactic: Evaluating llms playing the game of avalon. *arXiv preprint arXiv:2310.05036*, 2023.

[Light *et al.*, 2023b] Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. Avalonbench: Evaluating llms playing the game of avalon. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.

[Liu *et al.*, 2023] Zijun Liu, Yanzhe Zhang, et al. Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization. *arXiv preprint arXiv:2310.02170*, 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, et al. 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, et al. 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 *ACSOS-C*, pages 104–109. IEEE, 2023.

[Navid *et al.*, 2023] Navid, Majumdar, et al. Generative agent-based modeling: Unveiling social system dynamics through coupling mechanistic models with generative artificial intelligence. *arXiv preprint arXiv:2309.11456*, 2023.

[Park *et al.*, 2022] Joon Sung Park, Lindsay Popowski, et al. Social simulacra: Creating populated prototypes for social computing systems. In *UIST*, pages 1–18, 2022.

[Park *et al.*, 2023] Joon Sung Park, Joseph C O'Brien, et al. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*, 2023.

[Qian *et al.*, 2023] Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*, 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, et al. Tptu: large language model-based ai agents for task planning and tool usage. *arXiv preprint arXiv:2308.03427*, 2023.

[Shinn *et al.*, 2023] Noah Shinn, Federico Cassano, et al. Reflexion: Language agents with verbal reinforcement learning, 2023.

[Sumers *et al.*, 2023] Theodore R Sumers, Shunyu Yao, et al. Cognitive architectures for language agents. *arXiv preprint arXiv:2309.02427*, 2023.

[Tang *et al.*, 2023] Xiangru Tang, Anni Zou, et al. 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, et al. Adapting llm agents through communication, 2023.

[Wang *et al.*, 2023b] Lei Wang, Chen Ma, et al. A survey on large language model based autonomous agents, 2023.

[Wang *et al.*, 2023c] Shenzhi Wang, Chang Liu, et al. Avalon's game of thoughts: Battle against deception through recursive contemplation. *arXiv preprint arXiv:2310.01320*, 2023.

[Weiss *et al.*, 2024] Martin Weiss, Nasim Rahaman, Manuel Wuthrich, Yoshua Bengio, Li Erran Li, Bernhard Schölkopf, and Christopher Pal. Rethinking the buyer's inspection paradox in information markets with language agents, 2024.

[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 Hosseinichimeh, et al. Epidemic modeling with generative agents. *arXiv preprint arXiv:2307.04986*, 2023.

[Wu *et al.*, 2023] Qingyun Wu, Gagan Bansal, et al. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*, 2023.

[Xi *et al.*, 2023] Zhiheng Xi, Wenxiang Chen, et al. The rise and potential of large language model based agents: A survey, 2023.

[Xiao *et al.*, 2023] Bushi Xiao, Ziyuan Yin, et al. 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, 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, et al. Examining inter-consistency of large language models collaboration: An in-depth analysis via debate, 2023.

[Xu *et al.*, 2023a] Lin Xu, Zhiyuan Hu, et al. Magic: Investigation of large language model powered multi-agent in cognition, adaptability, rationality and collaboration, 2023.

[Xu *et al.*, 2023b] Yuzhuang Xu, Shuo Wang, et al. 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, et al. 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, et al. 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. *arXiv preprint arXiv:2310.10108*, 2023.

[Zhang *et al.*, 2023b] Ceyao Zhang, Kaijie Yang, 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, et al. 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, et al. Agentcf: Collaborative learning with autonomous language agents for recommender systems, 2023.

[Zhao *et al.*, 2023] Qinlin Zhao, Jindong Wang, et al. Competeai: Understanding the competition behaviors in large language model-based agents, 2023.

[Zheng *et al.*, 2023] Zhiling Zheng, Oufan Zhang, et al. Chatgpt research group for optimizing the crystallinity of mofs and cofs. *ACS Central Science*, 9(11):2161–2170, 2023.

[Zhou *et al.*, 2023a] Wangchunshu Zhou, Yuchen Eleanor Jiang, 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, et al. Sotopia: Interactive evaluation for social intelligence in language agents. *arXiv preprint arXiv:2310.11667*, 2023.

[Ziems *et al.*, 2023] Caleb Ziems, Omar Shaikh, et al. Can large language models transform computational social science? *Computational Linguistics*, pages 1–53, 2023.