MeKB-Sim: Personal Knowledge Base-Powered Multi-Agent Simulation

Zhenran Xu, Jifang Wang, Baotian Hu*, Longyue Wang, Min Zhang
Harbin Institute of Technology (Shenzhen), Shenzhen, China
xuzhenran@stu.hit.edu.cn, vincentwang0229@gmail.com, hubaotian@hit.edu.cn

Abstract

Language agents have demonstrated remarkable emergent social behaviors within simulated sandbox environments. However, the characterization of these agents has been constrained by static prompts that outline their profiles, highlighting a gap in achieving simulations that closely mimic real-life interactions. To close this gap, we introduce MeKB-Sim, a multi-agent simulation platform based on a dynamic personal knowledge base, termed MeKB. Each agent's MeKB contains both fixed and variable attributes—such as linguistic style, personality, and memory—crucial for theory-of-mind modeling. These attributes are updated when necessary, in response to events that the agent experiences. Comparisons with human annotators show that the LLM-based attribute updates are reliable. Based on the dynamic nature of MeKB, experiments and case study show that MeKB-Sim enables agents to adapt their planned activities and interactions with other agents effectively. Our platform includes a Unity WebGL game interface for visualization and an interactive monitoring panel that presents the agents' planning, actions, and evolving MeKBs over time[1].

1 Introduction

Agent-based modeling and simulation focus on modeling complex systems by simulating individual agents and their interactions within an environment (Gao et al., 2023). The rapid development of large language models (LLMs) has significantly advanced these simulations, offering more realistic representations of agents' decision-making processes, communication, and adaptation within simulated environments (Shinn et al., 2023; Zhang et al., 2024). The observation of emergent social behaviors in Generative Agents (Park et al., 2023) has spurred a series of multi-agent simulation demonstrations (Wang et al., 2023b; Lin et al., 2023; Chen et al., 2024b; Wu et al., 2023). However, in these demonstrations, each agent is specified by a paragraph of natural language description, detailing the agent's identity, occupation, and relationships with other agents. Such specification is far from true-to-life simulations of human-like agents, thus constraining the potential for simulating more sophisticated human behaviors to test and prototype social systems and theories.

To address the limitations in existing demonstrations, we introduce MeKB-Sim, a multi-agent simulation platform that leverages a dynamic personal knowledge base, denoted as MeKB. The MeKB of each agent incorporates attributes critical for theory-of-mind modeling (Sang et al., 2022). Specifically, the MeKB for each agent is structured into hierarchical layers, comprising the central fixed attributes such as occupation, race, education level, relationships, and linguistic style, surrounded by variable attributes such as personality, long-term and short-term memory, the emotion status. Note that variable attributes are subject to modification in response to experienced events. By comparing LLM-based attribute modifications with human annotations, we find a high acceptance rate for MeKB updates, demonstrating the reliability of our system. This dynamic attribute adjustment mechanism underscores the flexibility and adaptability of MeKB-Sim, making it suitable for simulating complex interactions and mental states in language agents.

Drawing upon the main components of language agents outlined by Xi et al. (2023), the architecture of agents in MeKB-Sim integrates a planning system and a MeKB-empowered characterization with an included memory module. The agent simulation begins by setting a daily goal, generated by LLMs using prior experiences as in-context examples. Subsequently, the planning system performs

* Corresponding author.

[1]For more information, including the open-source code, a video walkthrough and a live demo website, please visit our project page at https://mekb-sim.github.io/.

Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations), pages 393–403

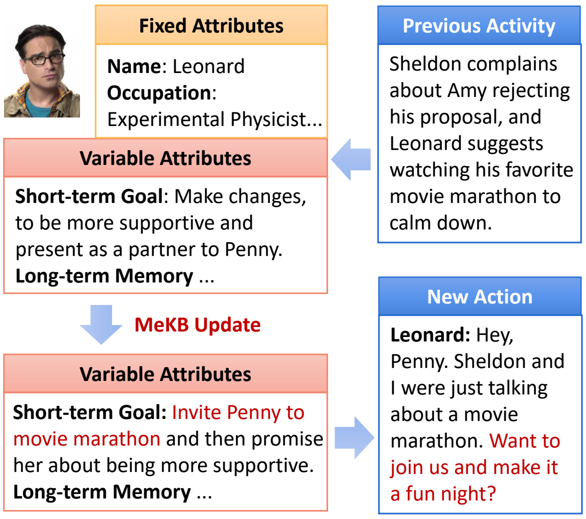

Figure 1: An example in MeKB-Sim demonstrating how prior experiences influence the MeKB, which subsequently impacts the ongoing dialogue. The complete conversation for this instance is detailed in Sec. 6.

self-reflection, which decomposes the daily goal into a series of questions and answers them on its own, covering who to meet, where, and why. With these plans in place, the agent initiates targeted conversations, proceeding to the designated locations to interact with the intended agents. When applying MeKB in our simulation, as illustrated in Figure 1, prior experiences influence the attributes within the MeKB-based profiling (e.g., short-term goal), and this profiling module subsequently impacts actions in future activities (e.g., the invitation to a movie marathon in the simulated conversation). Experiments demonstrate that the MeKB-based profiling module achieves better goal alignment, timeline coherence, and character consistency compared to static profiles, as MeKB showcases the continuous evolution of characters.

The contributions of this work are threefold:

- We introduce the concept of MeKB, a personal knowledge base designed to comprehensively and dynamically characterize an agent.
- We present MeKB-Sim, a platform that can simulate the social behavior of MeKB-based human-like language agents.
- We visualize MeKB-Sim with an Unity WebGL game interface and present the planning, actions and evolving MeKB of agents over time with an interactive monitoring panel.

2 Related Work

Building Agents with Large Language models (LLMs). LLMs have demonstrated remarkable emergent abilities in reasoning and planning (Wei et al., 2022). These advantages have earned LLMs the designation of sparks for AGI (Bubeck et al., 2023), making them highly desirable for building the "brains" of intelligent agents. The architecture of LLM-based agents typically includes several key modules: profiling, memory, planning, and action (Wang et al., 2024). Following the above overall framework, AutoGPT (Significant Gravitas, 2023), Voyager (Wang et al., 2023a) and AppAgent (Zhang et al., 2023) seek to create task-oriented agents that can autonomously interact with the environment.

Multi-Agent Simulation. In the transition from a single-agent framework into multi-agent simulations, the pioneering research on Generative Agents (Park et al., 2023) has laid the groundwork for the development of "Simulated Society". These societies are conceptualized as dynamic systems where multiple agents engage in intricate interactions within a well-defined environment (Guo et al., 2024). Recent research on simulated societies has followed two primary lines, namely, exploring the boundaries of collective intelligence (Li et al., 2023; Du et al., 2023; Wu et al., 2023; Xu et al., 2023; Qian et al., 2024a; Chen et al., 2024c) and using them to accelerate discoveries in the social sciences (Lin et al., 2023; Wang et al., 2023b; Qian et al., 2024b; Chen et al., 2024b). However, the agent specifications in these studies are oversimplified and *static*, by providing name, age, and a few sentences describing the agent. Inspired by theory-of-mind modeling (Sang et al., 2022), here we present a *dynamic* human-centered knowledge base to enhance the profiling process.

Personal Knowledge Base (KB). The concept of personal KBs has become an emerging solution for managing structured information about individuals (Balog and Kenter, 2019). Recent years have witnessed a growing interest in leveraging personal KBs across personalized applications to align with each user's unique habits and preferences. These applications range from research assistants (Chakraborty et al., 2022), e-learning tutors (Ilkou, 2022), and product recommendation (Yang et al., 2022) to suicidal ideation detection on social media (Cao et al., 2022). In this work, we incorporate personal KBs into multi-agent simulation by constructing MeKB, applying it to planning and action, with continuous updates.

3 MeKB-Sim

In this section, we introduce the technical details of our multi-agent simulation platform, MeKB-Sim. We first describe the construction and updating processes of the MeKB (Sec. 3.1). We use the OpenAI `gpt-4-1106-preview` API for all generations during simulation. Following this, we detail the agent architecture, the simulation process and the methods for integrating MeKBs into their planning and actions (Sec. 3.2).

3.1 MeKB

Construction. The MeKB of each agent includes the attributes crucial for theory-of-mind modeling (Sang et al., 2022). Specifically, 14 attributes of MeKB are organized into three hierarchical layers. As shown in Figure 2, at the core are the central fixed attributes, i.e. name, gender, race, occupation, education level, linguistic style, interpersonal relationships and long-term goal. Surrounding this core, the second layer comprises variable attributes, i.e. personality, long-term and short-term memory, the emotion status and short-term goal. The outermost layer includes all fine-grained experienced events, which may influence the second-layer variable attributes. The more central a layer is, the more stable its attributes are. The initialization of each attribute is predetermined based on relevant documentation about the simulated environment (will be detailed in Sec. 4.1).

Update. After each conversation, the interaction is recorded as an event, which prompts an update to the MeKB based on the event's details. For example, to determine the emotion status expressed by the agent, we use the following prompt:

> *In the following conversation, what emotion does {agent_name} express?*
> *{conversation_history}*
> *Please respond only with one word from this list ["neutral", "disgusted", "afraid", "sad", "surprised", "happy", "angry"].*

The set of emotions consists of the seven emotions listed above according to Ekman (1992). The default emotion status is set to "neutral". The conversation is then summarized and archived into the agent's long-term memory. Concurrently, the short-term memory (i.e. the dialogue history) is cleared, in preparation for the next activity.

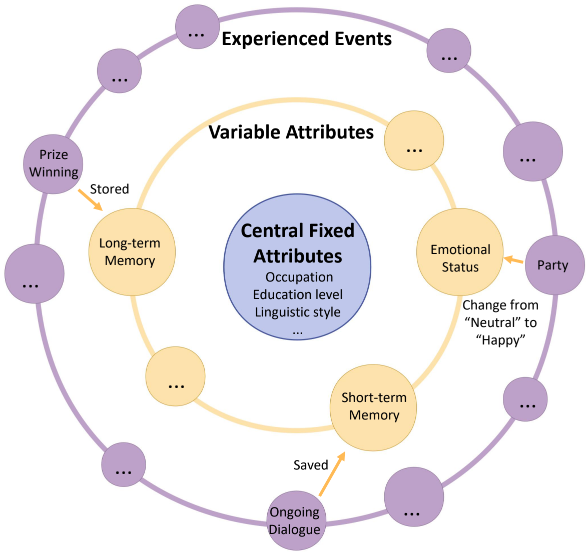

Figure 2: The hierarchical layers of MeKB. The more central a layer is, the more stable its attributes are. The experienced events at the outermost layer may influence the second-layer variable attributes.

3.2 Agent Simulation

Agent Simulation Process. As shown in Figure 3, the agent architecture in MeKB-Sim integrates a short-term and long-term memory system, a planning system, and a MeKB-empowered profiling system. The agent simulation begins by setting an initial short-term goal, generated by LLMs using prior experiences as in-context examples. Subsequently, the planning system performs self-reflection, which decomposes the goal into a series of questions and then answers them independently. These questions are fundamental in the achievement of the goal, including the identities of agents to meet, locations for meetings, and the underlying purposes of such interactions. After devising these plans, the agent proceeds to the designated location and initiates conversations with the intended agents. After each interaction, the MeKB is updated, prompting the planning system to adjust the next planned activity accordingly.

Using MeKB in Agent Simulation. When applying MeKB in our simulation, the linguistic style is illustrated through in-context demonstrations. Long-term memory retrieval is based on the scene's purpose and conversation context, ensuring a coherent and goal-oriented interaction. Other attributes in MeKB are explicitly expressed in the prompt. For example, the prompt framework for the response generation is shown in Appendix A.

Regarding the long-term memory retrieval, we

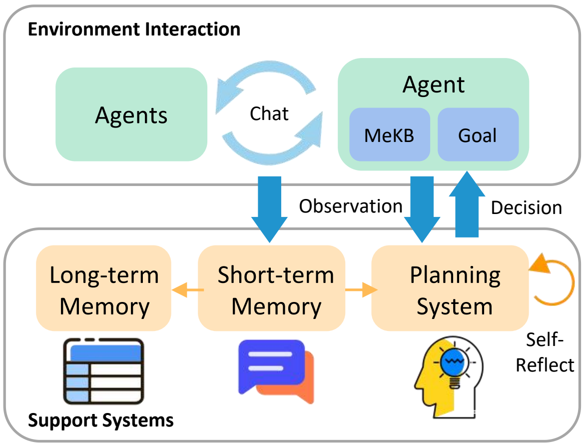

Figure 3: The overview of agent architecture in MeKB-Sim, encompassing a memory module, a planning module and a MeKB-based profiling module.

embed the dialogue by concatenating the dialogue topic with the conversation history to retrieve the most relevant and recent memory. Our retrieval model is based on M3-Embedding (Chen et al., 2024a). We utilize the Faiss library (Douze et al., 2024) as our vector database for embedding storage and similarity search. In this work, we only retrieve one previous memory for generating responses in the subsequent stage considering the length of prompts.

4 Platform Implementation

In this section, we introduce our simulation of the sandbox environment (Sec. 4.1), and describe how the Unity WebGL Game Interface and Interactive Monitoring Panel facilitate user visualization of agent statuses (Sec. 4.2).

4.1 Environment

We have developed a sandbox environment based on “The Big Bang Theory” TV show, with plans to support additional worlds on our platform in the future. We choose this comedy because it contains personalities that are well-known to many (e.g. Sheldon). Specifically, we use character information from `https://the-big-bang-theory.com/characters/`, and prompt `gpt-4-1106-preview` for the attribute initialization of MeKB. For the implementation of long-term memory, we rely on the scripts from the first eight seasons of “The Big Bang Theory”, which are publicly available in Sang et al. (2022). We synthesize summaries of all scenes with `gpt-3.5-turbo-1106` and store these summaries, along with the dialogues, in a knowledge base served as long-term memory.

4.2 Visualization Tools

Unity WebGL Game Interface. Based on Lin et al. (2023), we create an HTML game environment using the Unity WebGL game engine to visualize our simulation results[2]. The front end of MeKB-Sim is shown in Figure 4. On the left side, a panel displays the goals and action flows for each agent. The main screen shows the agents' behaviors as they navigate various locations and initiate conversations.

Interactive Monitoring Panel. Our interactive monitoring panel allows users to observe the status of various agents over time. Through this panel, users can select an agent from the simulated world to view its activity timeline and MeKB. Figure 7 in Appendix B displays the initial short-term goals brainstormed for all characters, along with detailed planning and conversations. Users can refresh to see the latest simulation results by clicking the "refresh data" button. Additionally, users have the option to decide whether the activities displayed on the timeline should be added to long-term memory. Figure 8 in Appendix B illustrates each agent's current state in MeKB, allowing users to explore how each attribute evolves over time and its impact on the activities (see the case study in Sec. 6).

5 Experiments

We conducted two sets of experiments to assess the reliability of LLM-based attribute updates and the impact of MeKB on simulated activities. We first investigate the reliability of LLM-based attribute updates, validated by human annotating the correctness of updated attributes (Sec. 5.1). Then, we study the effects of MeKB on the simulated activity results, regarding the alignment with brainstormed goals, timeline coherence and consistency with the character profiles (Sec. 5.2).

5.1 Comparison with Human Annotations

The accurate updating of variable attributes in MeKB is essential for human-like social simulation. We randomly select 50 instances of MeKB evolvements and let annotators label for the acceptability of changes in variable attributes. Two graduate students were recruited to annotate changes in personality, emotional status, and short-term goals.

[2]`https://docs.unity3d.com/Manual/webgl-building.html`

Figure 4: The Unity WebGL game interface of MeKB-Sim, showing in a pixel game style. The left-side panel concisely displays the goals and action flows of each agent. The main screen shows the agent behaviors, including moving to locations and initiating conversations. The interface of MeKB monitoring panel is shown in Appendix B.

Category	Cohen's κ	Acceptance Rate
Personality	0.696	0.89
Emotion	0.675	0.81
Short-term Goal	0.688	0.74

Table 1: Cohen's Kappa coefficient and acceptance rates for changes in variable attributes.

	Goal	Timeline	Character
Static Profile	3.15	3.25	1.30
MeKB-based	4.25	3.55	2.00
Correlation	0.623	0.601	0.511

Table 2: **Goal** alignment, **timeline** coherence and **character** consistency scores of MeKB-based profiles and static profiles. The Pearson correlation scores between human annotators are also reported.

The mean acceptance rate and the inter-annotator agreement score (i.e. Cohen's Kappa coefficient) are reported in Table 1.

The results show substantial inter-rater agreement (Cohen's $\kappa > 0.6$) across all attributes (Fleiss and Cohen, 1973). Our system achieves strong performance in updating emotions and personality (average acceptance > 0.8) and demonstrates moderately high performance in updating short-term goals, highlighting the reliability of LLM-based attribute updates.

5.2 Effects of MeKB on Activities

Here we examine the effects of MeKB on simulated activities involving six agents across ten different scenarios. The baseline for comparison is a static profile, utilizing the static attributes of MeKB only. We evaluate three aspects: (1) Goal Alignment: Do the simulated activities align with each agent's initial brainstormed goal? (2) Timeline Coherence: Is each agent's activity timeline fluent and coherent? (3) Character Consistency: Are the lines in the simulation consistent with the characters from the TV series? For the first two aspects, the same two annotators from Sec. 5.1 participate, scoring on a scale of 1 to 5. For the third aspect, we recruit two annotators who are enthusiastic about *The Big Bang Theory* and have watched the entire series.

Table 2 shows high inter-annotator agreement for the first two evaluation aspects, with MeKB-based profiling achieving better goal alignment

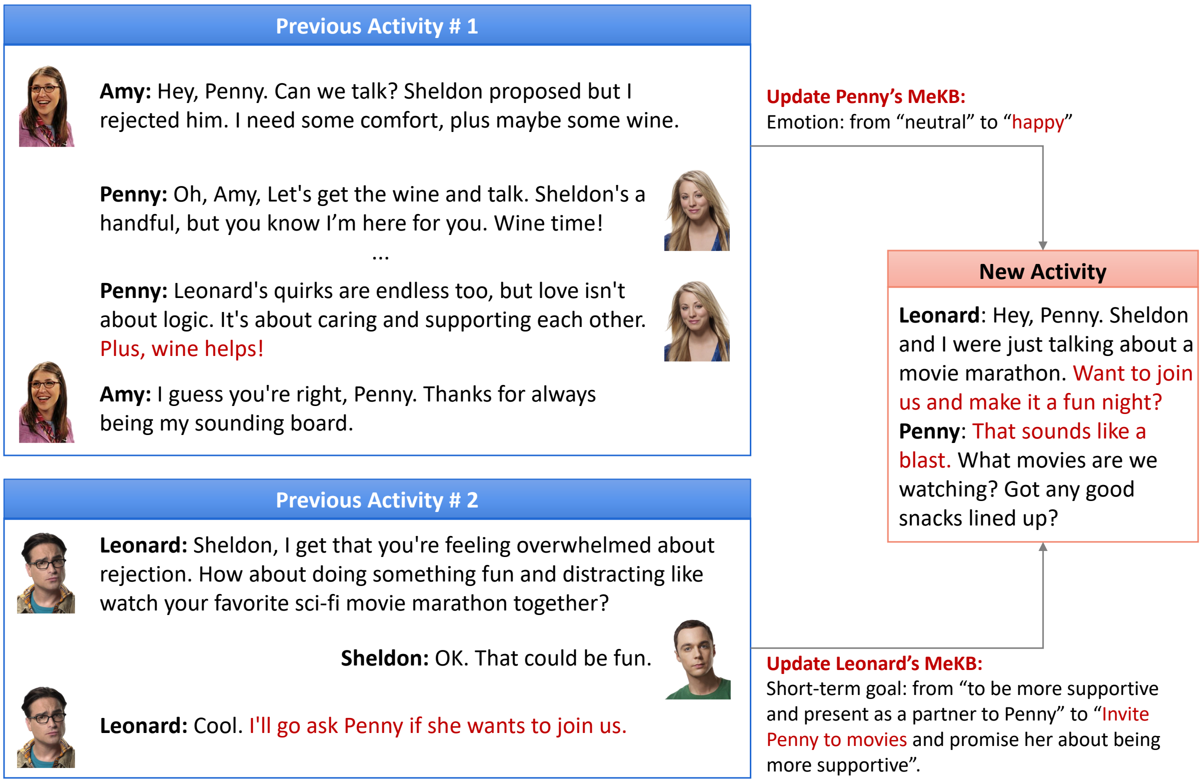

Figure 5: An example of simulated activities and the accompanied MeKB updates in "The Big Bang Theory".

and timeline coherence. This improvement is attributed to the continuous updating of variable attributes, which maintains the consistency of characters' memory and emotions across different activities. Additional cases are discussed in Sec. 6.

For the third evaluation aspect, although MeKB scores higher than static profiles, there remains significant room for improvement in aligning simulated lines with the original characters. Common feedback includes that the lines are not funny enough, diverging from the figures as remembered from the series. Nonetheless, MeKB-Sim occasionally references plots from previous seasons, enhancing the characters' consistency. This finding aligns with cognitive conclusions in Sang et al. (2022) that a character's memory is crucial for humans to construct its Theory of Mind (ToM).

6 Case Study

From the case presented in Figure 5, incorporating MeKB enables character lines to reference their previous experiences in the show, makes the overall timeline coherent, and maintains consistent goals and emotions across different scenes. For example, as we simulate from the end of Season 8, both activities #1 and #2 draw upon a long-term memory of the season's final episode. This memory is summarized as, "Amy rejects Sheldon's proposal and expresses her need for space and time to re-evaluate her relationship with Sheldon". Based on these simulated activities, the MeKBs of the characters are updated to reflect changes, such as Penny's emotions and Leonard's short-term goals. These updated attributes are then incorporated into the next conversation, showing the continuous evolution of characters throughout the simulation process.

7 Conclusion

In this paper, we introduce MeKB-Sim, a platform specifically designed to simulate the behavior of human-like language agents. The simulation is based on MeKB, a dynamic personal knowledge base aimed at providing a comprehensive theory-of-mind modeling of an agent. To enhance interaction with MeKB-Sim, we employ a Unity WebGL game interface, enabling users to visually engage with the simulation. Additionally, we offer an interactive monitoring panel that details the planning, actions, and the evolution of the MeKB over time.

Limitations

Our system currently supports only six agents in a sandbox environment. Future work may extend this capacity and integrate more user-defined characters, more diverse profiles and underlying LLMs (Qian et al., 2025). Additionally, future advancements may involve the creation of story videos, derived from plots generated by multi-agent simulations (Maas et al., 2023; Xu et al., 2024).

It is important to note that any imperfections in the LLMs will be inherited by the language agents (Park et al., 2023). While the agent frameworks may alleviate some of these issues, addressing them fundamentally requires improving the LLMs and aligning their values.

Ethics Statement

Our system may assist researchers such as computational social scientists to simulate human behavior prior to conducting costly real-world studies. Although our system aims to make agents behave more human-like, it is crucial for users to understand that they are interacting with LLMs that do not perfectly replicate real-world human behavior.

Acknowledgments

We want to thank Zhilin Wang for valuable discussion and feedback, and thank several anonymous reviewers for helpful comments. This work is jointly supported by grants: Natural Science Foundation of China (No. 62376067, No. 62422603), and Shenzhen Science and Technology Program (Grant No. ZDSYS20230626091203008).

References

Krisztian Balog and Tom Kenter. 2019. Personal knowledge graphs: A research agenda. In *Proceedings of the 2019 ACM SIGIR International Conference on Theory of Information Retrieval*, ICTIR '19, page 217–220, New York, NY, USA. Association for Computing Machinery.

Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4.

Lei Cao, Huijun Zhang, and Ling Feng. 2022. Building and using personal knowledge graph to improve suicidal ideation detection on social media. *IEEE Transactions on Multimedia*, 24:87–102.

Prantika Chakraborty, Sudakshina Dutta, and Debarshi Kumar Sanyal. 2022. Personal research knowledge graphs. In *Companion Proceedings of the Web Conference 2022*, WWW '22, page 763–768, New York, NY, USA. Association for Computing Machinery.

Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024a. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation.

Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2024b. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *The Twelfth International Conference on Learning Representations*.

Weize Chen, Ziming You, Ran Li, Yitong Guan, Chen Qian, Chenyang Zhao, Cheng Yang, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2024c. Internet of agents: Weaving a web of heterogeneous agents for collaborative intelligence.

Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. 2024. The faiss library.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multiagent debate.

Paul Ekman. 1992. An argument for basic emotions. *Cognition & Emotion*, 6:169–200.

Joseph L. Fleiss and Jacob Cohen. 1973. The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability. *Educational and Psychological Measurement*, 33:613 – 619.

Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu, and Yong Li. 2023. Large language models empowered agent-based modeling and simulation: A survey and perspectives.

Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges.

Eleni Ilkou. 2022. Personal knowledge graphs: Use cases in e-learning platforms. In *Companion Proceedings of the Web Conference 2022*, WWW '22, page 344–348, New York, NY, USA. Association for Computing Machinery.

Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for "mind" exploration of large language model society. In *Advances in Neural Information Processing Systems*, volume 36, pages 51991–52008. Curran Associates, Inc.

Jiaju Lin, Haoran Zhao, Aochi Zhang, Yiting Wu, Huqiuyue Ping, and Qin Chen. 2023. Agentsims: An open-source sandbox for large language model evaluation.

Philipp Maas, Frank Carey, Chris Wheeler, Edward Saatchi, Pete Billington, and Jessica Yaffa. 2023. To infinity and beyond: Show-1 and showrunner agents in multi-agent simulations.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA. Association for Computing Machinery.

Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2024a. ChatDev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186, Bangkok, Thailand. Association for Computational Linguistics.

Chen Qian, Zihao Xie, Yifei Wang, Wei Liu, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2024b. Scaling large-language-model-based multi-agent collaboration.

Chen Qian, Zihao Xie, YiFei Wang, Wei Liu, Kunlun Zhu, Hanchen Xia, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2025. Scaling large language model-based multi-agent collaboration. In *The Thirteenth International Conference on Learning Representations*.

Yisi Sang, Xiangyang Mou, Mo Yu, Shunyu Yao, Jing Li, and Jeffrey Stanton. 2022. TVShowGuess: Character comprehension in stories as speaker guessing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4267–4287, Seattle, United States. Association for Computational Linguistics.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652. Curran Associates, Inc.

Significant Gravitas. 2023. AutoGPT.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023a. Voyager: An open-ended embodied agent with large language models.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji-Rong Wen. 2024. A survey on large language model based autonomous agents.

Zhilin Wang, Yu Ying Chiu, and Yu Cheung Chiu. 2023b. Humanoid agents: Platform for simulating human-like generative agents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 167–176, Singapore. Association for Computational Linguistics.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. Emergent abilities of large language models. *Transactions on Machine Learning Research*. Survey Certification.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. 2023. Autogen: Enabling next-gen llm applications via multi-agent conversation framework.

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. 2023. The rise and potential of large language model based agents: A survey.

Zhenran Xu, Senbao Shi, Baotian Hu, Jindi Yu, Dongfang Li, Min Zhang, and Yuxiang Wu. 2023. Towards reasoning in large language models via multi-agent peer review collaboration.

Zhenran Xu, Jifang Wang, Longyue Wang, Zhouyi Li, Senbao Shi, Baotian Hu, and Min Zhang. 2024. Filmagent: Automating virtual film production through a multi-agent collaborative framework. In *SIGGRAPH Asia 2024 Technical Communications*, SA '24, New York, NY, USA. Association for Computing Machinery.

Yu Yang, Jiangxu Lin, Xiaolian Zhang, and Meng Wang. 2022. Pkg: A personal knowledge graph for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, page 3334–3338, New York, NY, USA. Association for Computing Machinery.

Chi Zhang, Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2023. Appagent: Multimodal agents as smartphone users.

Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. 2024. Exploring collaboration mechanisms for llm agents: A social psychology view.

A Prompt Template

The framework for MeKB-based response generation is shown in Figure 6. It primarily comprises MeKB-based profiles, short-term memory (i.e., the current conversation context), retrieved long-term memory, and in-context demonstrations of the character's linguistic style.

B Interface

Our visualization tools include Unity WebGL game interface and interactive monitoring panel. The game interface has been shown in Figure 4. As for the monitoring panel, Figure 7 and 8 illustrate the activity timeline and MeKB of each agent respectively.

```
You need to play a TV comedy character to chat with another character.


**I will give you the following information:**

1. **Character Profile**:
   - Name: {name}
   - Gender: {gender}
   - Occupation: {occupation}
   - Personality: {personality}
   - Interpersonal Relationships: {interpersonal_relationships}
   - Short Term Goal: {short_term_goal}
   - ...

2. **Current Conversation Information**:
   - The name of whom the character is chatting with: {chatTo}
   - The topic that the character wants to talk about: {chatTopic}
   - The character's Long-Term Memory related to this topic: {chatHistory}
   - The last content from the one you are talking to: {chats}

3. **Demonstration of the character's speaking style**:
   - {dialogue_demonstration}


**You must follow the following criteria:**
   - Maintain humor and mimic the character's speaking style in this conversation.
   - The conversation must be conformed to the long-term memory and the bio of the
     character, and it should reflect the character's personality traits.
   - Your knowledge level should not exceed that of a normal person with the bio of
     the character, unless there are relevant memories in the character's Long-
     Term Memory.
   - You should just tell the sentences you want to speak in the JSON format: {"
     content":"{name} : xxx"}
   - If The last content from the one you are talking to is "None" or nothing, you
     must start a conversation politely about the topic.
   - If The last content from the one you are talking to is not "None" or nothing,
     you must respond appropriately to the other person's words.
   - Your reply should not exceed 30 words.
```

Figure 6: Prompt for MeKB-based response generation, modified from Lin et al. (2023).

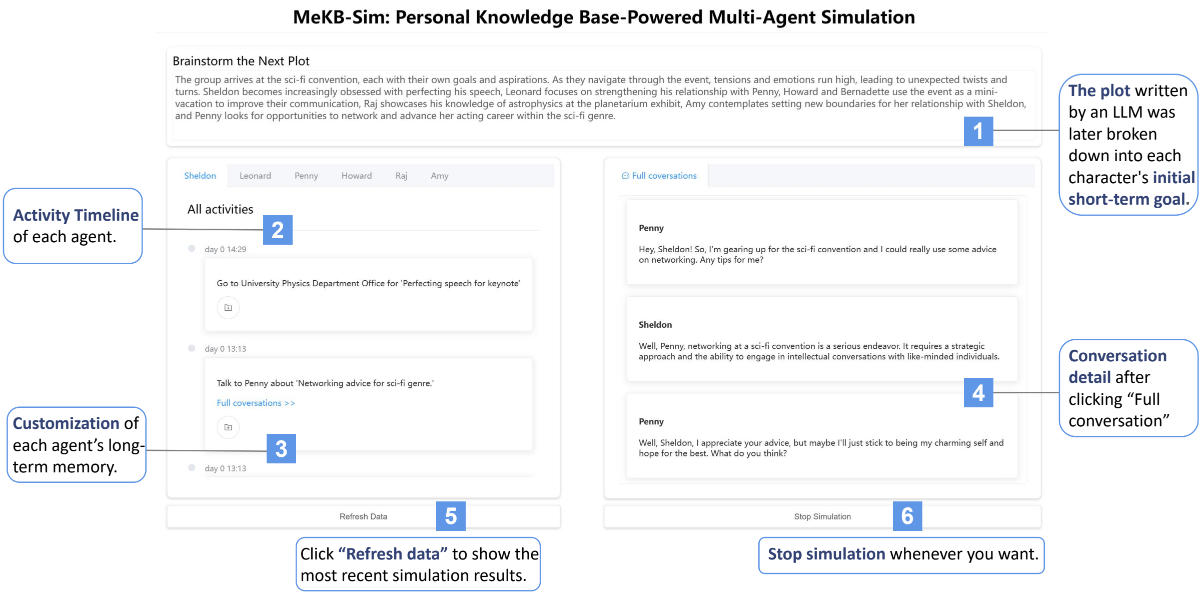

Figure 7: Interactive monitoring panel: the brainstormed goals and activity timelines of each agent.

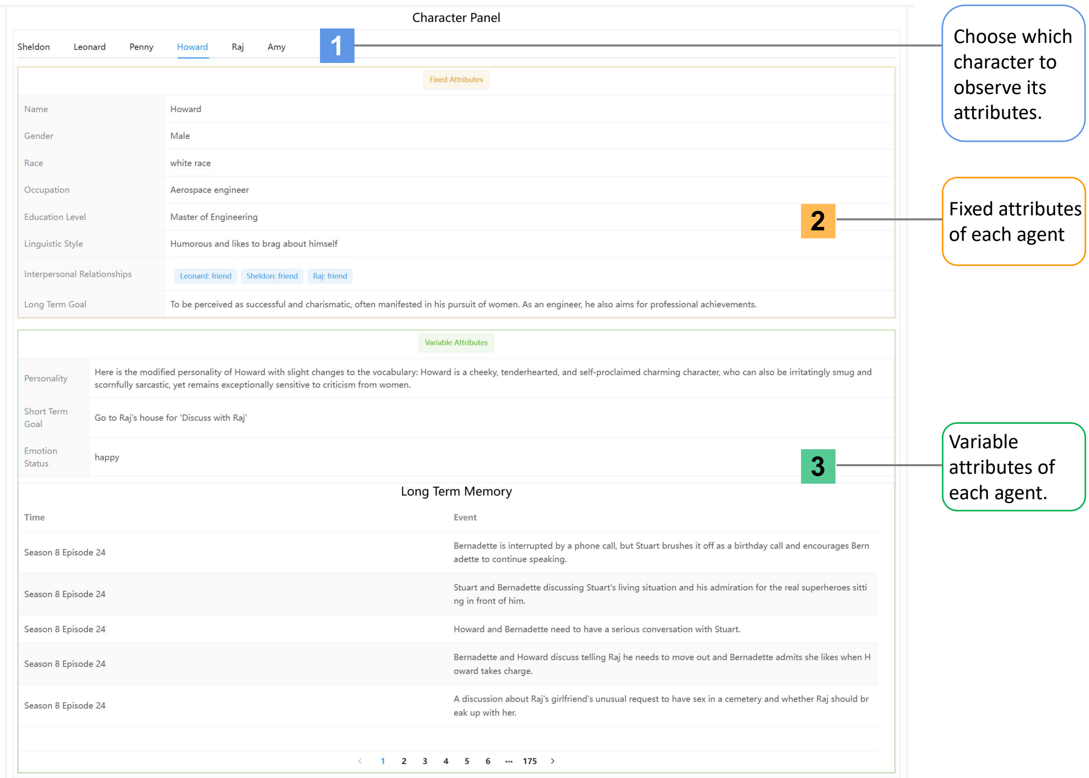

Figure 8: Interactive monitoring panel: the MeKB of each agent.