

Modeling U.S. Attitudes Toward China via an Event-Steered Multi-Agent Simulator

Chenxu Zhu, Hantao Yao, *Member, IEEE*, Wu Liu, *Senior Member, IEEE*, Junbo Guo, and Yongdong Zhang, *Fellow, IEEE*

Abstract—Understanding the dynamic evolution of opinions, such as U.S. public attitudes toward China, is essential for assessing geopolitical risks. However, existing LLM-based multi-agent simulators predominantly rely on static rules and fixed datasets, limiting their ability to capture the dynamic, event-driven nature of macro-level opinion shifts in real-world settings. To address this limitation, we propose an Event-Steered Multi-Agent Simulator (ES-MAS), in which significant events and daily news continuously drive opinion evolution through dynamic interactions among agents. We first construct the China–U.S. Relation Evolution (CURE) dataset, covering 20 quarters from 2021 to 2025, including 258 major events and over 14,000 daily news articles, and providing a comprehensive temporal foundation for modeling opinion dynamics. Building upon the CURE dataset, we propose a Dual-Stream Data Integration Engine (DSDIE) that aligns simulations with historical timelines via macro-level events while enabling personalized information exposure based on individual agent profiles and contextual signals. Furthermore, we design a News-Driven Dynamic Interaction (NDDI) module, which adaptively groups agents with shared news interests into localized interaction contexts, facilitating bottom-up consensus formation while mitigating the risk of isolated information cocoons. Experimental results on the CURE dataset demonstrate that ES-MAS substantially outperforms existing simulators in reproducing real-world historical trends, offering a scalable and effective framework for modeling dynamic opinion evolution.

Index Terms—Multi-Agent system, Opinion Dynamics, Social Simulation.

I. INTRODUCTION

The dynamic evolution of opinions in social networks plays a crucial role in shaping collective behavior and informing policy decisions. Understanding these processes is essential for anticipating societal trends and mitigating potential risks. In this context, China–U.S. relations represent a major geopolitical challenge of the 21st century, with significant implications for global stability [1]. Public attitudes toward this relationship are highly dynamic, influenced by events, media narratives, and social interactions. Ignoring these shifts may lead to misinformed decisions and increased risks of conflict [2]. Therefore, accurately analyzing and forecasting U.S. public attitudes toward China is critical. However, existing approaches mainly rely on surveys and expert analyses [3], which are resource-intensive and insufficient for capturing the rapid, continuous evolution of real-world public sentiment.

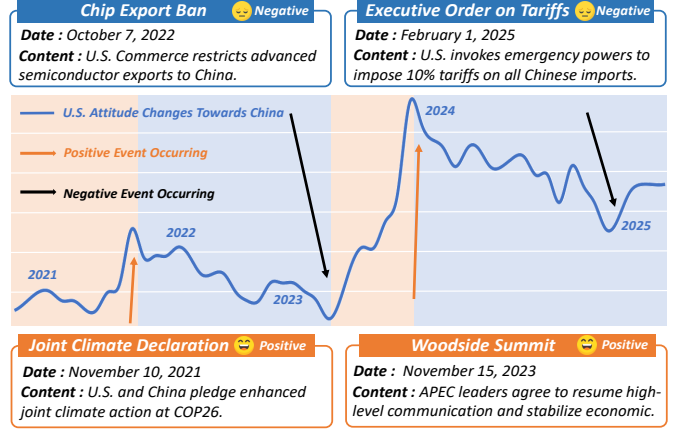


Fig. 1. The trajectory of U.S. attitudes toward China, illustrating how Real-world Geopolitical Events shape macro-level trends

Unlike traditional approaches that rely on human expertise for analysis [4], there has been growing interest in leveraging machine learning and artificial intelligence to study the evolution of U.S. attitudes toward China [5]. In particular, Agent-Based Models (ABMs) have been widely adopted to simulate macro-level social dynamics through the collective interactions of individuals [6], [7]. However, these models typically rely on simplified heuristic rules, failing to capture the heterogeneity and cognitive complexity of human decision-making [8]. Recent advances in large language models (LLMs) have enabled the development of agent-based simulators with enhanced reasoning and decision-making capabilities, allowing agents to interpret complex geopolitical contexts [9], [10], providing new opportunities for modeling individual-level perceptions of China–U.S. relations. Nevertheless, existing LLM-based simulators often adopt static interaction protocols and depend on pre-collected datasets to model public sentiment and its evolution [11], [12]. Such designs lack continuous external inputs, limiting their ability to capture the dynamic and event-driven nature of bilateral relations. Moreover, some recent work incorporates streaming daily news to infer societal trends by aggregating independently generated opinions from large populations of agents [13]. However, by overlooking peer-to-peer interactions, these approaches fail to model the social diffusion of opinions and tend to produce “information cocoons”, where agents remain confined to isolated perspectives. In conclusion, existing simulators remain inadequate for accurately reproducing the systemic and dynamic evolution of China–U.S. relations. Therefore, it is essential to consider both real-time data stream integration and the evolution of dynamic

Chenxu Zhu, Hantao Yao, Wu Liu, and Yongdong Zhang are with the University of Science and Technology of China (USTC), Hefei 230026, China (e-mail: cx_zhu@mail.ustc.edu.cn; yaohantao@ustc.edu.cn; liuwu@ustc.edu.cn; zhyd73@ustc.edu.cn).

Junbo Guo is with People.cn, Beijing, China (e-mail: guojunbo@people.cn). Corresponding author: Yongdong Zhang.

interactions to accurately simulate U.S. attitudes toward China.

A compelling motivation is to simulate macro-level bilateral trends through the lens of micro-level individual opinions, employing the integration of meaningful real-world data between China and the United States alongside dynamic multi-agent interactions [14], [15]. The simulation framework is grounded in dual information streams that incorporate both significant macro-level events and micro-level daily news. Unlike traditional simulation algorithms that depend exclusively on daily news data, we contend that substantial events influencing societal dynamics between China and the United States are essential for accurately modeling U.S. attitudes toward China. For example, the escalation of the Chip Export Ban in 2022 led to a significant deterioration in bilateral relations, while the leadership meeting at the Woodside Summit in 2023 contributed to a reduction in tensions, as illustrated in Figure 1. Within the simulation process, dynamic interactions among agents—driven by real-time news—effectively translate micro-level individual opinions into macro-level U.S. attitudes toward China. Research on small-world networks suggests that individuals are more likely to interact with peers who share similar opinions [16]. Furthermore, the dissemination of opinions can be achieved through the interconnectedness of numerous small-world networks on a global scale [17]. Through the analysis of news content, individuals with analogous views on China-U.S. relations can be clustered into distinct groups, allowing for dynamic interactions that facilitate the evolution of cognitive perspectives. Simultaneously, interactions between these groups can foster the spread of opinions on a broader level, enabling the deduction of the overall macro situation.

In light of the aforementioned motivation, we introduce a novel Event-Steered Multi-Agent Simulator (ES-MAS) designed to model U.S. attitudes toward China by utilizing significant real-time events and dynamic interactions between agents. To ensure that our simulations reflect real-world dynamics, we first develop the **China-U.S. Relation Evolution (CURE)** dataset, which comprises significant events and daily news related to China-U.S. interactions over the past five years. Next, we introduce a Dual-Stream Data Integration Engine to incorporate both significant events and daily news into the simulator. At the macro level, significant events establish the overarching historical context of China-U.S. relations. At the micro level, a Personalized Autonomous Ingestion Mechanism (PAIM) is proposed to select daily news based on agent preferences and current events, promoting cognitive diversity. Additionally, we propose a News-Driven Dynamic Interaction (NDDI) module that uses news topics as signals to dynamically cluster agents who are interested in the same China-U.S. news. This clustering process creates locally relevant interaction contexts, allowing agents to engage in group-level discussions. Overall, this strategy effectively demonstrates how micro-level interactions can influence the macro-evolution of China-U.S. relations from the bottom up.

Experiments conducted on the **CURE** dataset show that ES-MAS achieves superior performance in simulating U.S. Attitude Changes Toward China compared to existing social simulators. In summary, the major contributions can be summarized as follows:

- We propose a Dual-Stream Data Integration Engine (DS-DIE) to leverage significant events and daily news to effectively bridge macroscopic trends with micro-individual cognition.
- We propose a News-driven Dynamic Interaction (NDDI) module to dynamically cluster the agents into several groups for conducting the intra-group interaction.
- We construct a novel **China-U.S. Relation Evolution (CURE)** dataset comprising significant events and daily news related to China-U.S. interactions over the past five years.

II. RELATED WORK

Opinion Dynamics Modeling: Opinion dynamics research links micro-level interactions to macro-level phenomena via frameworks like selective exposure [18] and social identity theory [19]. Traditional surveys offer empirical grounding but remain labor-intensive and static [8]. To capture dynamic processes, modeling approaches have evolved from probabilistic methods [15] to traditional Agent-Based Models (ABMs); however, these ABMs typically simulate social dynamics using heuristic rules and often fail to capture human cognitive complexity [20], [21]. In contrast, current studies demonstrate that Large Language Models (LLMs) can reproduce complex cognitive behaviors, including human-like polarization on political issues [22], in-group favoritism [23], and cognitive dissonance [24]. Furthermore, the role of external information environments, such as mass media and critical exogenous events, is pivotal in shaping these cognitive trajectories [25], our work employs dynamic social interactions within a population of agents to simulate macroscopic trends in international perceptions.

LLM-based Social Simulators: LLMs empower agents to exhibit human-like behaviors, self-awareness [26], [27] and dispute deliberation and collective decision-making [28]. For example, Generative Agents [29] simulate the daily interactions among 25 agents in a virtual town, demonstrating that LLM-based agents can generate social behaviors indistinguishable from human-created content. Recent simulators have drastically scaled these societies; for instance, frameworks like OASIS [30] and AgentSociety [31] achieve massive-scale simulations to study emergent social phenomena such as emotional contagion, and misinformation spread. Building on this foundation, some multi-agent simulators utilize static data and fixed topologies for sociological analysis [11], [32], while recent graph-based frameworks such as [33] and [34] further explore large-scale dynamic social networks and evolving interaction structures. In the domain of opinion dynamics, other works incorporate time-varying data streams to track opinion changes, focusing on updating individual agent beliefs based on external information [35], [36]. Another category of research targets macro-level political outcomes, such as elections or the evolution of international relations [37], by modeling collective opinion dynamics to forecast systemic developments. In this work, ES-MAS introduces a dual-stream data integration engine and dynamic interactions among agents to synchronize the simulated society with historical reality.




Fig. 2. The content and structure of the CURE dataset about the Significant Events (SE) and Daily News (DN)

III. CHINA-U.S. RELATION EVOLUTION DATASET

To accurately simulate U.S. attitudes change toward China, we collect a novel **China-U.S. Relation Evolution (CURE)** dataset consisting of the significant events and daily news related to China-U.S. interactions over the past five years.

Significant Events (2021–2025): Significant events play a crucial role in shaping U.S. attitudes toward China. To better understand this, we collected significant events from both the Biden and Trump administrations, sourced from *China Briefing*¹, a policy intelligence platform. This documentation includes important occurrences such as presidential summits, trade sanctions, technology sector restrictions, and diplomatic conflicts. In total, we collected 258 significant events from 2021 to 2025. For example, one event is: “May 12, 2025: China and the U.S. agree to reduce reciprocal tariffs to 10%.” This event illustrates the complexities of the relationship between China and the United States.

Daily News (2021–2025): Besides the significant events, we also collect many daily news items related to China-U.S.. Specifically, we utilized *public media API interfaces*² to collect daily news spanning the periods of the Biden administration and the re-emergence of the Trump administration. To ensure comprehensiveness, we retrieved articles containing the keywords “China” and “U.S.” across distinct categories, including politics, economy, technology, and environment. Finally, we collected over 30,000 articles of daily news from 2021 to 2025. However, these daily news items are not directly relevant to shifts in the U.S. attitude towards China. To remove the noise from all articles, we employed the Qwen3-8B to justify whether the article contributes to the evolution of China-U.S. relations. For example, one news reads: “May 12, 2025: U.S. Senate passes Taiwan UN resolution, sparks China debate.”. Following the pre-processing, over 14,000 relevant articles remained.

By combining the collected significant events and daily news, we obtain the **China-U.S. Relation Evolution (CURE)** dataset. Each term of the CURE dataset consists of the *Data*, *Significant Event*, and *Daily News*. The detailed content and

TABLE I
DETAILED QUARTERLY DISTRIBUTION OF SIGNIFICANT EVENTS (SE)
AND DAILY NEWS (DN).

Year	Q1		Q2		Q3		Q4		Total	
	SE	DN	SE	DN	SE	DN	SE	DN	SE	DN
2021	8	527	15	1,040	7	482	17	1,190	47	3,239
2022	7	732	9	937	7	707	5	516	28	2,892
2023	7	328	19	891	14	640	19	928	59	2,787
2024	10	742	7	470	9	696	8	572	34	2,480
2025	19	675	32	1,100	22	798	17	679	90	3,252
Sum	51	3,004	82	4,438	59	3,323	66	3,885	258	14,650

integration structure of the CURE dataset are illustrated in Figure 2.

Dataset Statistics: Spanning 20 quarters from 2021 to 2025, the CURE dataset contains a total of 258 Significant Events (SE) and 14,650 Daily News (DN) articles. As detailed in Table I, the data distribution reflects the real-world intensity of U.S.-China relations. For example, the frequency of significant events peaks sharply in 2025 (90 events), aligning with the policy shifts and geopolitical friction during the U.S. administration transition. Meanwhile, the daily news, averaging over 2,900 articles annually, provides a dense background narrative.

Ground Truth Attitudes: To effectively evaluate whether the evolution of U.S. attitude towards China trajectories generated by our simulator align with historical facts, we give a ground truth of U.S. attitude towards China based on the *Database of Relations Between China and Major Powers (1950–2025)*³ released by the Institute of International Relations at Tsinghua University (TUIIR). Note that the TUIIR database offers monthly updates, enabling it to capture subtle shifts in China-U.S. relations, such as the transitions between confrontation and detente during the Biden and Trump administrations. Finally, we extracted TUIIR scores from 2021 to 2025, ensuring temporal alignment for our simulation with reality.

¹<https://www.china-briefing.com/>

²<https://open-platform.theguardian.com/>

³<https://www.tuiir.tsinghua.edu.cn/>




Fig. 3. The architecture of ES-MAS framework. Simulating U.S. Attitude Changes Toward China by Dual-Stream Data Integration Engine and News-Driven Interaction Module. The engine integrates significant events and daily news to generate individual attitudes and daily news-responsive actions. Subsequently, these actions facilitate the formation of news-interest groups, where dynamic interactions occur to influence the attitude evolution for the next step.

IV. EVENT-STEERED MULTI-AGENT SIMULATOR

Based on the CURE dataset, we introduce a novel Event-Steered Multi-Agent Simulator (ES-MAS), in which the opinions of multiple agents evolve under the influence of significant events and daily news. As shown in Figure 3, ES-MAS comprises two components: a Dual-Stream Data Integration Engine and a News-Driven Dynamic Interaction Module. The Dual-Stream Data Integration Engine synchronizes simulations with historical timelines through macro-level significant events while also allowing individualized information intake at the agent level. The News-Driven Dynamic Interaction (NDDI) Module dynamically groups agents according to news content and models the evolution of collective attitudes. Its strength lies in effectively demonstrating how micro-level local interactions can drive the macro-level evolution of bilateral relations from the bottom up.

Formally, ES-MAS is modeled as a multi-agent system with N agents, denoted as $\mathbb{S} = \{\mathbb{A}_1, \mathbb{A}_2, \dots, \mathbb{A}_N\}$. Each agent configuration $\mathbb{A}_i$ consists of the agent a_i , the persona profile $\mathcal{P}_i$, and the memory $\mathcal{M}_i$; that is, $\mathbb{A}_i = \{a_i, \mathcal{P}_i, \mathcal{M}_i\}$. Note that each agent a_i is initialized with a multi-dimensional persona profile $\mathcal{P}_i$, constructed based on demographic distributions [?], [37], and includes attributes such as name, age, gender, interests, personality, and political views. In addition to the profile, the agent's memory $\mathcal{M}_{i,t}$ initialized at simulation step $t = 0$ serves as another key component for adapting the agent's opinions.

In our simulation from 2021 to 2025, ES-MAS operates on quarterly (three-month) time steps, corresponding to discrete time steps $t \in \{1, 2, \dots, 20\}$. At each step t , each agent a_i receives dual-stream news $\mathcal{D}_{i,t}$ generated by the Dual-Stream Data Integration Engine (DSDIE). Based on its persona $\mathcal{P}_i$ and the received news $\mathcal{D}_{i,t}$, each agent produces two outputs: 1) An attitude score $O_{i,t} \in [-10, 10]$, reflecting sentiment toward China; 2) A news-responsive action $\mathcal{B}_{i,t}$, e.g., posting, retweeting, or replying, representing behavioral engagement

with the news stream. Through the News-Driven Dynamic Interaction (NDDI) module, these responses $\mathcal{B}_{i,t}$ facilitate dynamic inter-agent interactions and promote group formation based on shared opinions. These interactions constitute the micro-foundations underlying the evolution of U.S.–China relations. Finally, by aggregating the individual attitude scores $O_{i,t}$ across all agents, we derive the macro-level trajectory of U.S. attitude changes toward China.

A. Dual-Stream Data Integration Engine

To ensure the simulation aligns with real-world dynamics, we propose the Dual-Stream Data Integration engine to integrate information streams by jointly considering the significant events and daily news. Formally, we construct a comprehensive dual-stream data $\mathcal{D}_{i,t}$,

$$\mathcal{D}_{i,t} = \{\mathcal{S}_t, \mathcal{I}_{i,t}\}, \quad (1)$$

where $\mathcal{S}_t$ is the significant events steam, and $\mathcal{I}_{i,t}$ is the personalized autonomous daily news stream. For the agent a_i , the Significant Events Inject Module (SEIM) is proposed to generate the event data $\mathcal{S}_t$, and the Personalized Autonomous Ingestion Module (PAIM) is proposed to select the personalized daily news $\mathcal{I}_{i,t}$. In the following, we give a detailed description of those two components.

At the macro-level, SEIM exposes all agents $\{a_i\}_{i=1}^N$ to the significant events before each simulation step, effectively synchronizing the simulated society with the historical reality of China-U.S. relations. Formally, for each simulation step, the set of significant events is defined as $\mathcal{S}_t = \{s_{t,1}, s_{t,2}, \dots, s_{t,L}\}$, where L denotes the number of events in the current quarter, and each event contains the event ID, title, date, and content.

Different from SEIM, the Personalized Autonomous Ingestion module (PAIM) constructs a candidate news pool derived from the daily news dataset, enabling agents to autonomously select daily news in this pool that align with their interests. Specifically, based on the current significant events $\mathcal{S}_t$, the agent utilizes its persona profile $\mathcal{P}_i$ and recent memory $\mathcal{M}_{i,t}$

as decision criteria to ingest relevant articles $\mathcal{I}_{i,t}$ from the candidate news pool $\mathcal{W}_t = \{w_1, w_2, \dots, w_N\}$ with the Large Language Model(LLM),

$$\mathcal{I}_{i,t} = LLM(\mathcal{F}_{ingest}(\mathcal{S}_t, \mathcal{W}_t, \mathcal{P}_i, \mathcal{M}_{i,t})), \quad (2)$$

where $\mathcal{F}_{ingest}(\cdot)$ functions as the selection engine retrieve the top- K news items that exhibit the semantic alignment with the agent's persona, memory, and environment. More details are provided in the Appendix.

Finally, the dual-stream data $\mathcal{D}_{i,t}$ establishes a stable macro-level cognitive foundation for the simulation, facilitating diverse micro-level cognitive processing. This dual-stream architect guarantees that the simulation of U.S. attitude changes toward China evolves dynamically within the boundaries of historical facts.

B. News-Driven Dynamic Interaction Module

In multi-agent simulations, the dynamic interactions between agents play a crucial role in deriving societal behavior at the macro level from individual actions at the micro level [31], [38]. For example, we can infer the global attitude of China and the United States based on individual attitudes toward both countries. Inspired by small-world networks, people communicate primarily with a few individuals they are connected to within the small social network. Therefore, in multi-agent simulations, agents often interact dynamically with several agents belonging to the same group. Moreover, multiple groups can achieve global-level dynamic interactions through shared agents. Unlike existing simulators using a fixed topology or social topology, we construct a local interaction topology by considering agents who focuses the same daily news. To facilitate this, we propose the News-Driven Dynamic Interaction (NDDI) module to simulate interactions among agents based on dual-stream events, capturing the complex interplay between overarching bilateral trends and public interactions.

Given the dual-stream dataset $\mathcal{D}_{i,t} = \{\mathcal{S}_t, \mathcal{I}_{i,t}\}$ at the t -th step, it contains K circulating news items, i.e., $\mathcal{I}_{i,t} = \{\mathcal{I}_{i,t,1}, \mathcal{I}_{i,t,2}, \dots, \mathcal{I}_{i,t,K}\}$, where $\mathcal{I}_{i,t,k}$ is the k -th terms. At each step, daily news $\mathcal{I}_{i,t}$ and significant event $\mathcal{S}_t$ give the agent for selectively engaging with this news by performing actions (POST, RETWEET, or REPLY). Inspired by real social interaction, all agents acting on the same news item are clustered into a corresponding news-driven group. For the given news item $\mathcal{I}_{i,t,k}$, the corresponding group $\mathcal{G}_{t,k}$ comprises the agents who are interested in the same daily news,

$$\mathcal{G}_{t,k} = \{a_i \in \mathcal{A} \mid \mathcal{E}(a_i, \mathcal{I}_{i,t,k}) = 1\}, \quad (3)$$

where $\mathcal{E}(a_i, \mathcal{I}_{i,t,k}) = 1$ denotes that the agent a_i is interesting on the daily news $\mathcal{I}_{i,t,k}$, and $\mathcal{A}$ is the set of all agents. Finally, we group all agents $\mathcal{A} = \{a_1, a_2, \dots, a_N\}$ into K groups $\mathbb{G}_t = \{\mathcal{G}_{t,1}, \mathcal{G}_{t,2}, \dots, \mathcal{G}_{t,K}\}$. More details are provided in the Appendix.

After that, given the daily news $\mathcal{I}_{i,t,k}$, the agent a_i belonging to the group $\mathcal{G}_{t,k}$ conducts the interaction to produce the attitude towards China $O_{i,t}$, behavioral response $\mathcal{B}_{i,t}$, and updated memory $\mathcal{M}_{i,t}$,

$$O_{i,t}, \mathcal{B}_{i,t}, \mathcal{M}_{i,t} = LLM(\mathcal{F}_{act}(\mathcal{D}_{i,t}, \mathcal{P}_i, \mathcal{M}_{i,t-1}, \mathcal{R}_{i,t,k})), \quad (4)$$

Algorithm 1: Event-Steered Multi-Agent Simulator (ES-MAS)

Input : Agent Set $\mathbb{S}$; Steps T ; Significant Events $\mathcal{S}_t$; Candidate News Pool $\mathcal{W}_t$.
Output: Agents' attitude $\{O_{i,t}\}_{t=1}^T$ and behaviors $\{\mathcal{B}_{i,t}\}_{t=1}^T$ for all agents $i \in \{1, \dots, N\}$.

```

1: begin
2:   Initialize agent configurations
      $\mathbb{S} = \{\mathbb{A}_1, \mathbb{A}_2, \dots, \mathbb{A}_N\}$ ,  $\mathbb{A}_i = \{a_i, \mathcal{P}_i, \mathcal{M}_{i,0}\}$  and
     initial attitude  $O_{i,0}$  for  $i = 1 \dots N$ ;
3:   for each timestep  $t = 1$  to  $T$  do
4:     Expose  $\mathcal{S}_t$  for each agent;
5:     for each agent  $\mathbb{A}_i \in \mathbb{S}$  do
6:       Select news  $\mathcal{I}_{i,t}$  via PAIM based on
          $\mathcal{S}_t, \mathcal{W}_t, \mathcal{P}_i$  and  $\mathcal{M}_{i,t-1}$ ;
7:       Integrate data  $\mathcal{D}_{i,t} = \{\mathcal{S}_t, \mathcal{I}_{i,t}\}$ ;
8:       Generate initial attitude  $O_{i,t}$  and behavior
          $\mathcal{B}_{i,t}$  via  $\mathcal{F}_{act}$  based on  $\mathcal{D}_{i,t}, \mathcal{P}_i, \mathcal{R}_{i,t,k}$  and
          $\mathcal{M}_{i,t-1}$ ;
9:     Partition  $\mathbb{S}$  into groups  $\mathbb{G}_t = \{\mathcal{G}_{t,1}, \dots, \mathcal{G}_{t,K}\}$ 
       based on actions  $\mathcal{B}_{i,t}$ ;
10:    for each group  $\mathcal{G}_{t,k} \in \mathbb{G}_t$  do
11:      for each agent  $a_i \in \mathcal{G}_{t,k}$  do
12:        Form reference stream  $\mathcal{R}_{i,t,k}$  using
          visible replies  $\mathcal{C}_{i,t,k}$  and top- $\tau$  filtered
          peers  $\mathcal{V}_{i,t,k}$ ;
13:      Attitude and memory update of the
        Environment ;
14:  return  $\{O_{i,t}, \mathcal{B}_{i,t}\}_{i=1 \dots N}^{t=1 \dots T}$ ;

```

where $\mathcal{P}_i$ is the personal profile, and $\mathcal{F}_{act}(\cdot)$ is the decision engine to generate opinion and news-responsive actions, manifesting as operations including POST, RETWEET, and REPLY. $\mathcal{R}_{i,t,k}$ is the reference information stream for an agent a_i when interacting within a group $\mathcal{G}_{t,k}$. More details are provided in the Appendix.

When the agent conducts the interaction with Eq. (4), a critical problem is how to collect the reference information stream $\mathcal{R}_{i,t,k}$ among the agents from the group $\mathcal{G}_{t,k}$. In the News-Driven Dynamic Interaction, the act of engaging with the same news item establishes a shared information set to make peer actions potentially visible. Moreover, an agent would refer to other agents' responses to the same news to decide its own actions. Therefore, the reference information stream $\mathcal{R}_{i,t,k}$ consists of two terms,

$$\mathcal{R}_{i,t,k} = \mathcal{C}_{i,t,k} \cup \mathcal{V}_{i,t,k}, \quad (5)$$

where $\mathcal{V}_{i,t,k}$ is the dynamic information stream of agent a_i from other agents, and $\mathcal{C}_{i,t,k}$ denotes the set of observable reply contents to news item $\mathcal{I}_{i,t,k}$ at time t that are visible to agent a_i ,

$$\mathcal{C}_{i,t,k} = \{\mathcal{B}_{j,t}^{\text{REPLY}} \mid a_j \in \mathcal{G}_{t,k}, j \neq i\}, \quad (6)$$

where $\mathcal{B}_{j,t}^{\text{REPLY}}$ denotes the agent a_j conducts the same "RE-

PLY” operation as agent a_i .

As mentioned above, an agent would refer to other agents’ responses to the same news to decide its own actions. For example, when agent a_i and a_j both conduct the “REPLAY” operation on the same daily news, the response $O_{j,t}$ from the agent a_j can be treated as a reference stream fed to the agent a_i for interaction. Therefore, $\mathcal{V}_{i,t,k}$ is formulated to aggregate the response of agents from group $\mathcal{G}_{t,k}$ having the same operation as a_i with Eq. (7),

$$\mathcal{V}_{i,t,k} = \{\mathcal{B}_{j,t} \mid a_j \in \mathcal{G}_{t,k}, \text{rank}(S_{i,j,t}^{\text{rec}}) \leq \tau\}, \quad (7)$$

where $S_{i,j,t}^{\text{rec}}$ denotes the recommendation score between agent a_j and agent a_i at time t within the news-driven group $\mathcal{G}_{t,k}$. Moreover, agents are ranked in descending order with the top- τ selected. Specifically, $S_{i,j,t}^{\text{rec}}$ is calculated based on the content alignment term and the timeliness factor with Eq. (8),

$$S_{i,j,t}^{\text{rec}} = \underbrace{\Phi(\mathcal{B}_{i,t}, \mathcal{B}_{j,t})}_{\text{content alignment}} \cdot \underbrace{e^{-\beta(t-t_p)}}_{\text{timeliness factor}}, \quad (8)$$

where the first term is the content alignment term that measures the similarity between the response of $\mathcal{B}_{i,t}$ and $\mathcal{B}_{j,t}$ with function $\Phi(\cdot)$. The second term is the timeliness factor, which follows an exponential decay with rate β between the current time t and the response post time t_p . Besides of $\mathcal{R}_{i,t,k}$, the other aspect in Eq. (4) is how to update the memory $\mathcal{M}_{i,t}$ after interaction. Formally, we construct a dynamic memory module to store the agent’s opinion history at simulation step t , which comprises of the self-reflective memory $\mathcal{SF}(\mathcal{M}_{i,t})$ to record the agent’s historical attitude and opinion, and dynamic interaction data $\mathcal{R}_{i,t,k}$. The update of individualized memory is formulated as,

$$\mathcal{M}_{i,t,k} = \{\mathcal{SF}(\mathcal{M}_{i,t}), \mathcal{R}_{i,t,k}\}, \quad (9)$$

where $\mathcal{SF}(\mathcal{M}_{i,t})$ is derived via the function $\mathcal{SF}(\cdot)$ that dynamically queries and synthesizes the agent’s cumulative history to extract historical opinion. More details are provided in the Appendix.

The interaction topology is dynamic. As the simulation progresses to the next $(t+1)$ -th step, the introduction of fresh dual-stream news, $\mathcal{D}_{i,t+1}$, drives agents to re-cluster into new news-driven groups based on their interests. After several iterative simulations, the ES-MAS enables an efficient simulation of U.S. attitude changes toward China, revealing how dynamic social interactions influence the evolution of public attitudes and collective behaviors, is formalized in Algorithm 1.

V. EXPERIMENT

A. Experimental Settings

Configurations: We utilize the GPT-4o-mini model for agent reasoning and the text-embedding-3-large model for generating embeddings, with a memory decay rate of $\beta = -0.1$ and a recommendation list size of $\tau = 8$. The simulation involves $N = 100$ agents across $T = 20$ quarterly steps, spanning from 2021 to 2025. The dataset consists of 258 significant events distributed throughout this timeline. During

TABLE II
RESULT OF MACRO ALIGNMENT EVALUATION. **BOLD** INDICATES THE BEST RESULTS, AND UNDERLINING INDICATES THE SECOND-BEST RESULTS. THE $\downarrow$ SYMBOL IMPLIES LOWER IS BETTER.

Method	U.S. Attitude Changes Toward China			
	$\Delta\text{Bias} \downarrow$	$\Delta\text{Diversity} \downarrow$	DTW $\downarrow$	Frechet $\downarrow$
BC	5.5644	1.9248	22.7863	6.7223
HK	3.9369	1.7856	17.3651	7.5510
LR	2.1805	1.6224	10.9095	3.5467
SJ	3.0801	1.3135	12.8437	4.7100
RA	1.6639	1.2397	8.8880	3.0467
FPS	<u>1.4987</u>	1.1945	<u>7.2168</u>	<u>2.0020</u>
HiSim	1.8936	<u>1.0445</u>	8.9016	5.0950
SOD	4.1018	<u>1.2653</u>	16.0846	5.1000
ES-MAS	0.5767	0.7423	2.4403	1.6150




Fig. 4. Comparison of simulated and real-world U.S. public attitude dynamics toward China over time.

each step, agents process all concurrent significant events and autonomously select $K = 24$ daily news articles from a randomized candidate pool of $\mathcal{N} = 100$. In total, our simulator processes over 70,000 agent-news interactions.

Metrics: To assess the accuracy of ES-MAS in simulating U.S. Attitude Changes Toward China, we evaluate the simulation across two dimensions: 1) *Numerical Alignment*: We use ΔBias and ΔDiv to quantify the mean deviation and variance stability between simulated and real-world attitudes [39]; 2) *Temporal Dynamics*: To assess how well the simulation captures historical evolutionary trends, we employ Dynamic Time Warping (DTW) [40] and Fréchet distance [41];

Baselines: We compare the proposed method with the following baselines: 1) *Agent-Based Models*: BC [6], HK [42], RA [7], SJ [43], and LR [20] which update attitudes via fixed mathematical rules regarding neighbor selection and message acceptance. 2) *LLM-based social network simulators*: FPS [32], which models belief confirmation and consensus-building processes; SOD [35], which focuses on simulating confirmation bias and consensus; and HiSim [11], a hybrid architecture that scales up LLM agents by integrating them with traditional ABM mechanisms.

TABLE III
ABLATION ANALYSIS OF THE COMPONENT IN ES-MAS.

Method	Components			Performance			
	SEIM	PAIM	NDDI	$\Delta\text{Bias} \downarrow$	$\Delta\text{Div.} \downarrow$	$\text{DTW} \downarrow$	$\text{Frech.} \downarrow$
ES-MAS	✓	✓	✓	0.5767	0.7423	2.4403	1.6150
w/o DSDIE	×	×	✓	1.4984	0.7999	6.1764	2.0650
Only PAIM	×	✓	×	1.4529	1.7247	6.7676	3.4300
w/o SEIM	×	✓	✓	1.2386	1.3714	5.2251	2.8650
Only SEIM	✓	×	×	1.1613	1.4073	4.9091	2.6550
w/o PAIM	✓	×	✓	1.0994	0.9710	4.2318	2.5275
w/o NDDI	✓	✓	×	0.9177	1.1642	3.7478	2.5000

B. Comparison with State-of-the-Art Methods

To assess the alignment of U.S. attitudes toward China with reality, we systematically conduct a Macro Alignment Evaluation by comparing the simulated results of ES-MAS with various static baselines. As shown in Table II, ES-MAS consistently demonstrates the best performance across all metrics, highlighting its accuracy, robustness, and generability in alignment. In contrast, existing simulators that rely on fixed interaction rules and pre-existing static data struggle to accurately model public attitudes toward China. This further emphasizes the importance of incorporating real-time news data and the dynamic interactions among agents.

Specifically, regarding Numerical Distribution, ES-MAS achieves a reduction in the mean values of the second-best model by 0.922 ΔBias and 0.3022 $\Delta\text{Diversity}$, thereby substantiating its advantage in marked stability and accurate reflection of public attitude dynamics over long-time simulations. This superior performance indicates that the simulated trajectories generated by ES-MAS are tightly synchronized with the real-world timeline. Regarding trend shape alignment, ES-MAS exhibits notable improvements in both DTW and Frechet metrics, realizing respective average decreases of 4.7765 and 0.387 relative to the second-best method. This affirms its effectiveness in accurately replicating the non-linear fluctuations of historical evolution. For an intuitive comparison of the differences between the simulated and real-world attitude dynamics, we present the visualization results of the U.S. attitude changes toward China in Fig. 4

C. Ablation Study

In this section, we conduct several ablations to analyze the effect of each component.

Dual-Stream Data Integration Engine (DSDIE): DSDIE integrates both significant events and daily news to provide a continuous, real-time data stream for simulation. The setting “ES-MAS w/o DSDIE” without considering DSDIE means that the agents are initialized and internal interactions without any reality-grounding updates. As shown in Table III, it results in the most severe performance degradation, marked by a substantial increase of 0.9217 in ΔBias and 3.7361 in DTW. This shows that continuously injecting dynamic news streams is crucial for accurate simulation, guaranteeing that the simulated public sentiment is consistently influenced by real-world information.




Fig. 5. Sensitivity Analysis on different initial attitudes.

Significant Events Inject Module(SEIM): SEIM directly injects real-time significant events into the simulation environment for establishing a global context accessible to all agents. By excluding SEIM, agents update attitudes relying solely on the cumulative effect of daily news and internal interactions. As shown in Table III, the setting “w/o SEIM” results in a distinct performance decline, *e.g.*, obtaining worse performance of 1.2386, 1.3714, 5.2251, and 2.8650 in terms of ΔBias , $\Delta\text{Div.}$, DTW, and Frech., respectively. This indicates that the injection of significant events plays a crucial role in accurate U.S. Attitude Changes Toward China, as daily news alone fails to capture the critical historical turning points and shifts in public sentiment.

Personalized Autonomous Ingestion Module (PAIM): Operating within the context of significant event injection, PAIM constructs a daily news pool that enables agents to autonomously select news matching their preferences. By discarding PAIM, agents are restricted to reacting solely to significant events without considering their personalized information intake. Consequently, they are driven entirely by macro trends, rendering them devoid of independent opinion evolution and prone to artificial consensus. As shown in Table III, the removal of PAIM results in a performance drop, causing an increase of 0.5227 in ΔBias and 1.7915 in DTW. This indicates that the autonomous selection from the daily news pool is essential for cognitive diversity in the evolution of China-U.S. relations.

News-Driven Dynamic Interaction Module (NDDI): NDDI aims to collect agents into local groups for facilitating dynamic interactions among agents belonging to the same group. Once excluding this NDDI from the ES-MAS, agents’ interaction becomes trapped in information cocoons because of the lack of information exchange, rendering them unable to adapt to the evolution of the China-U.S. relations. As shown in Table III, the setting “w/o NDDI” results in a distinct decline, marked by an increase of 0.3410 in ΔBias and 0.4219 in $\Delta\text{Diversity}$. This indicates that local interactions within dynamic clusters are the driving force behind the emergence of macro-level trends.

Effect of Significant Events: Significant events represent the impact of macro-level dynamic trends. In the Only




Fig. 6. The attitude trajectories of two representative agents (Mike and Ella) toward China from 2021 to 2025, illustrates how individual opinions dynamically adapt to macro-level geopolitical shocks (e.g., the Chip Export Ban and Woodside Summit) and micro-level influences (e.g., personalized news selection and dynamic peer interactions).

SEIM setting, the simulation is driven exclusively by significant events, where agents evolve solely based on macro-level shocks, devoid of personalized daily news and dynamic interactions. As shown in Table III, this setting yields a relatively robust performance, achieving a ΔBias of 1.1613. This result outperforms the w/o SEIM setting (1.2386) in terms of bias alignment, suggesting that significant events are the primary determinant of the overall attitude baseline. This demonstrates that SEIM successfully captures the fundamental turning points and the structural trend of China-U.S. relations, ensuring the macro-direction remains correct.

Effect of Daily News: In the Only PAIM setting, agents are exposed solely to daily news, meaning their opinion evolution is driven exclusively by the autonomous selection of information, lacking both macro-level guidance and dynamic interactions among agents. As shown in Table III, this setting suffers from significant misalignment, yielding a ΔBias of 1.4529 and a DTW of 6.7676. Notably, this performance is inferior to the Only SEIM setting (ΔBias 1.1613), this indicate that agents tend to get lost in diverse daily information. Therefore, the daily news provides the necessary details for preventing the opinion trajectory from reality.

Effect of Initial Attitudes: The initial attitude is the initial opinion of each agent about the U.S. Attitudes Toward China, while a different attitude would cause a different simulation trend. Therefore, a good simulator should be insensitive to different initializations. To evaluate the robustness of ES-MAS, we conducted a sensitivity analysis with a variety of initial attitudes. Formally, we initialized the agents with five distinct mean attitude values, i.e., $-7.6, -3, 0, 3, 7$.

As shown in Figure 5, ES-MAS exhibits a consistent simulation trend for the distinct initial attitudes, thereby successfully reproducing the macro-level evolutionary trend.

Effect of Micro-level: At the individual cognitive level, Figure 6 tracks the attitude trajectories of two representative agents, Mike and Ella, to illustrate how our framework's mechanisms drive opinion shifts. When encountering macro-level shocks and micro-level influences, the agents generate distinct attitude fluctuations supported by personalized rationales. This visualization clearly illustrates how individual opinions real-

TABLE IV
SCALABILITY ANALYSIS REGARDING FPS (BASELINE) AND ES-MAS.

Method	Size (N)	Scalability Analysis (N agents)			
		$\Delta\text{Bias} \downarrow$	$\Delta\text{Diversity} \downarrow$	DTW $\downarrow$	Frechet $\downarrow$
ES-MAS	25	0.8388	1.0134	3.5158	1.6900
	50	0.9783	1.1771	4.2899	2.3300
	100	0.5767	0.7423	2.4403	1.6150
	200	1.5027	1.7857	7.1429	3.0050
	500	1.3275	1.6052	6.2304	3.5056
FPS	25	1.3650	1.4309	6.7985	2.1800
	50	1.2461	1.3043	6.3538	1.9900
	100	1.4987	1.1945	7.2168	2.0020
	200	6.0051	1.8115	25.3306	9.1530
	500	6.3244	1.6678	28.5097	8.8335

istically adapt to a variety of complex exogenous events and endogenous social interactions over time. Consequently, this micro-level analysis validates the internal mechanisms within ES-MAS, demonstrating its effectiveness in driving macro-level trends through bottom-up individual behaviors.

D. Scalability Analysis

To evaluate the scalability of the proposed ES-MAS, we conduct further evaluation based on varying the agent numbers $N \in \{25, 50, 100, 200, 500\}$. Moreover, we compare ES-MAS against the second-best FPS across four alignment metrics, and summarize the related results in Table IV. It can be seen that increasing the number of agents can improve the accuracy of the simulation. However, when the number of agents exceeds 100, the accuracy of the simulation decreases. The reason is that an excessive number of agents introduces complex interaction dynamics, which facilitates the emergence of more polarized opinions and impedes the agents' to accurately judge the current real-time situation.

Furthermore, compared to the existing FPS simulator, the proposed ES-MAS demonstrates strong performance across various scales, with minimal differences in performance regardless of the number of agents. For instance, in terms of

numerical stability, ES-MAS consistently maintains low error rates across all four metrics, even as the number of agents, N , increases to 500. When N reaches 200, FPS experiences a significant spike in ΔBias , measuring 6.0051, whereas ES-MAS retains a much lower ΔBias of 1.5027. At $N = 500$, FPS fails to converge, showing a ΔBias greater than 6.3, while ES-MAS remains stable with a ΔBias of 1.3275. This demonstrates that our simulator effectively reduces the collective polarization typically observed in scaling Multi-Agent Systems through dynamic interactions among agents and real-time news data.

Regarding trend shape alignment, our simulator maintains the temporal accuracy of the simulation, while FPS experiences a severe decline. At $N = 500$, the Dynamic Time Warping (DTW) for FPS skyrockets to 28.5097, indicating a complete disconnection from the historical timeline. In contrast, ES-MAS achieves a low DTW of 6.2304. This suggests that ES-MAS ensures that macro-level evolutionary trends are kept closely aligned with reality, even as the agent population increases.

VI. CONCLUSION

In this work, we propose ES-MAS, an Event-Steered Multi-Agent Simulator designed to understanding the evolution of U.S. public attitudes toward China. By integrating dual-stream data and news-driven dynamic interactions among agents, ES-MAS ensuring a degree of alignment between simulated societal trends and historical reality. Unlike traditional approaches that rely on static rules and data, our framework effectively captures the continuous, bottom-up emergence of public opinion. Extensive experiments demonstrate that ES-MAS offers a approach for analyzing complex geopolitical dynamics, providing insights into the mechanisms driving international relations.

ETHICAL STATEMENT

The development and application of the Event-Steered Multi-Agent Simulator (ES-MAS) require careful ethical consideration, particularly given its focus on sensitive geopolitical topics. To prevent misuse for manipulating public opinion, transparency and strict supervision are essential, necessitating standardized usage scopes, clear permissions, and defined ethical boundaries. Regarding data usage, the CURE dataset is constructed entirely from publicly available news sources and historical event records (e.g., The Guardian API, policy briefings). No personally identifiable information or private user data was involved in the simulation process. We are committed to ensuring data reliability while strictly respecting privacy boundaries. To mitigate this, robust verification and control mechanisms will be implemented to prevent the spread of harmful information. By adhering to stringent standards and safeguards, ES-MAS will be used solely for legitimate purposes, avoiding any form of information manipulation or illegal activities.

REFERENCES

- [1] G. T. Allison, “Destined for war?” *The National Interest*, no. 149, pp. 9–21, 2017.
- [2] X. Li and Y. Y. Ang, “Media consumption patterns and public hostility toward china in the united states,” *Communication and the Public*, p. 20570473251401708, 2025.
- [3] K. M. Campbell and E. Ratner, “The china reckoning: how beijing defied american expectations,” *Foreign Aff.*, vol. 97, p. 60, 2018.
- [4] L. Silver, C. Huang, L. Clancy, and M. Fagan, *Americans Are Critical of China’s Global Role-as Well as Its Relationship With Russia*. JSTOR, 2023, vol. 12.
- [5] M. D. Wang, “Artificial intelligence techniques for political risk management: An nlp analysis of the 2019 us-china trade war,” in *Machine Learning and Data Mining Annual Volume 2023*. IntechOpen, 2023.
- [6] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, “Mixing beliefs among interacting agents,” *Advances in Complex Systems*, vol. 3, no. 01n04, pp. 87–98, 2000.
- [7] G. Deffuant, F. Amblard, G. Weisbuch, and T. Faure, “How can extremism prevail? a study based on the relative agreement interaction model,” *Journal of artificial societies and social simulation*, vol. 5, no. 4, 2002.
- [8] D. Lazer, A. Pentland, L. Adamic, S. Aral, A.-L. Barabási, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann *et al.*, “Computational social science,” *Science*, vol. 323, no. 5915, pp. 721–723, 2009.
- [9] W. Hua, L. Fan, L. Li, K. Mei, J. Ji, Y. Ge, L. Hemphill, and Y. Zhang, “War and peace (waragent): Large language model-based multi-agent simulation of world wars,” *arXiv preprint arXiv:2311.17227*, 2023.
- [10] Z. Guan, X. Kong, F. Zhong, and Y. Wang, “Richelieu: Self-evolving llm-based agents for ai diplomacy,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 123 471–123 497, 2024.
- [11] X. Mou, Z. Wei, and X.-J. Huang, “Unveiling the truth and facilitating change: Towards agent-based large-scale social movement simulation,” in *Findings of the Association for Computational Linguistics: ACL 2024*, 2024, pp. 4789–4809.
- [12] C. Gao, X. Lan, Z. Lu, J. Mao, J. Piao, H. Wang, D. Jin, and Y. Li, “S3: Social-network simulation system with large language model-empowered agents,” *arXiv preprint arXiv:2307.14984*, 2023.
- [13] N. Sukiennik, Y. Xu, Y. Kan, J. Piao, Y. Yan, C. Gao, and Y. Li, “Debiasing international attitudes: Llm agents for simulating us-china perception changes,” *arXiv preprint arXiv:2508.08837*, 2025.
- [14] T. C. Schelling, *Micromotives and macrobehavior*. WW Norton & Company, 2006.
- [15] M. H. DeGroot, “Reaching a consensus,” *Journal of the American Statistical association*, vol. 69, no. 345, pp. 118–121, 1974.
- [16] M. McPherson, L. Smith-Lovin, and J. M. Cook, “Birds of a feather: Homophily in social networks,” *Annual review of sociology*, vol. 27, no. 1, pp. 415–444, 2001.
- [17] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [18] D. Frey, “Recent research on selective exposure to information,” *Advances in experimental social psychology*, vol. 19, pp. 41–80, 1986.
- [19] H. Tajfel and J. C. Turner, “The social identity theory of intergroup behavior,” in *Political psychology*. Psychology Press, 2004, pp. 276–293.
- [20] J. Lorenz, M. Neumann, and T. Schröder, “Individual attitude change and societal dynamics: Computational experiments with psychological theories,” *Psychological Review*, vol. 128, no. 4, p. 623, 2021.
- [21] B. J. Berry, L. D. Kiel, and E. Elliott, “Adaptive agents, intelligence, and emergent human organization: Capturing complexity through agent-based modeling,” *Proceedings of the National Academy of Sciences*, vol. 99, no. suppl_3, pp. 7187–7188, 2002.
- [22] P. Törnberg, D. Valeeva, J. Uitermark, and C. Bail, “Simulating social media using large language models to evaluate alternative news feed algorithms,” *arXiv preprint arXiv:2310.05984*, 2023.
- [23] L. P. Argyle, E. C. Busby, N. Fulda, J. R. Gubler, C. Rytting, and D. Wingate, “Out of one, many: Using language models to simulate human samples,” *Political Analysis*, vol. 31, no. 3, pp. 337–351, 2023.
- [24] G. V. Aher, R. I. Arriaga, and A. T. Kalai, “Using large language models to simulate multiple humans and replicate human subject studies,” in *International conference on machine learning*. PMLR, 2023, pp. 337–371.
- [25] J. Zaller, *The nature and origins of mass opinion*. Cambridge university press, 1992.

- [26] Z. Xi, W. Chen, X. Guo, W. He, Y. Ding, B. Hong, M. Zhang, J. Wang, S. Jin, E. Zhou *et al.*, “The rise and potential of large language model based agents: A survey,” *Science China Information Sciences*, vol. 68, no. 2, p. 121101, 2025.
- [27] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin *et al.*, “A survey on large language model based autonomous agents,” *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.
- [28] Y. Sun, W. Liu, H. Ming, X. Wang, H. Yao, and Y. Zhang, “Cyberjurors: A multi-agent simulation task for e-commerce disputes verdict,” *arXiv preprint arXiv:2605.28369*, 2026.
- [29] J. S. Park, J. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, “Generative agents: Interactive simulacra of human behavior,” in *Proceedings of the 36th annual acm symposium on user interface software and technology*, 2023, pp. 1–22.
- [30] Z. Yang, Z. Zhang, Z. Zheng, Y. Jiang, Z. Gan, Z. Wang, Z. Ling, J. Chen, M. Ma, B. Dong *et al.*, “Oasis: Open agent social interaction simulations with one million agents,” *arXiv preprint arXiv:2411.11581*, 2024.
- [31] J. Piao, Y. Yan, J. Zhang, N. Li, J. Yan, X. Lan, Z. Lu, Z. Zheng, J. Y. Wang, D. Zhou *et al.*, “Agentsociety: Large-scale simulation of llm-driven generative agents advances understanding of human behaviors and society,” *arXiv preprint arXiv:2502.08691*, 2025.
- [32] Y. Liu, X. Chen, X. Zhang, X. Gao, J. Zhang, and R. Yan, “From skepticism to acceptance: Simulating the attitude dynamics toward fake news,” *arXiv preprint arXiv:2403.09498*, 2024.
- [33] X. Zhou, Y. Sun, H. Yao, A. He, Y. Zhang, and W. Liu, “Gasim: A graph-accelerated hybrid framework for social simulation,” *arXiv preprint arXiv:2605.07692*, 2026.
- [34] Y. Sun, W. Liu, W. Wang, H. Yao, J. Luo, and Y. Zhang, “Dynamix: Large-scale dynamic social network simulator,” *arXiv preprint arXiv:2507.19929*, 2025.
- [35] Y.-S. Chuang, A. Goyal, N. Harlalka, S. Suresh, R. Hawkins, S. Yang, D. Shah, J. Hu, and T. Rogers, “Simulating opinion dynamics with networks of llm-based agents,” in *Findings of the association for computational linguistics: NAACL 2024*, 2024, pp. 3326–3346.
- [36] J. S. Park, L. Popowski, C. Cai, M. R. Morris, P. Liang, and M. 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*, 2022, pp. 1–18.
- [37] X. Zhang, J. Lin, L. Sun, W. Qi, Y. Yang, Y. Chen, H. Lyu, X. Mou, S. Chen, J. Luo, X. Huang, S. Tang, and Z. Wei, “Electionsim: Massive population election simulation powered by large language model driven agents,” *CoRR*, vol. abs/2410.20746, 2024. [Online]. Available: <https://doi.org/10.48550/arXiv.2410.20746>
- [38] J. Tang, H. Gao, X. Pan, L. Wang, H. Tan, D. Gao, Y. Chen, X. Chen, Y. Lin, Y. Li *et al.*, “Gensim: A general social simulation platform with large language model based agents,” in *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (System Demonstrations)*, 2025, pp. 143–150.
- [39] M. Müller, *Information retrieval for music and motion*. Springer, 2007.
- [40] T. Eiter, H. Mannila *et al.*, “Computing discrete fréchet distance,” 1994.
- [41] F. P. Santos, Y. Lelkes, and S. A. Levin, “Link recommendation algorithms and dynamics of polarization in online social networks,” *Proceedings of the National Academy of Sciences*, vol. 118, no. 50, p. e2102141118, 2021.
- [42] R. Hegselmann, “Opinion dynamics and bounded confidence: models, analysis and simulation,” *The Journal of Artificial Societies and Social Simulation*, 2015.
- [43] W. Jager and F. Amblard, “Uniformity, bipolarization and pluriformity captured as generic stylized behavior with an agent-based simulation model of attitude change,” *Computational & Mathematical Organization Theory*, vol. 10, no. 4, pp. 295–303, 2005.