


Article

LLM-AIDSim: LLM-Enhanced Agent-Based Influence Diffusion Simulation in Social Networks

Lan Zhang ^{1,*} , Yuxuan Hu ² , Weihua Li ^{1,*} , Quan Bai ² and Parma Nand ¹ 

¹ School of Engineering, Computer & Mathematical Sciences, Auckland University of Technology, Auckland 1010, New Zealand; parma.nand@aut.ac.nz

² School of Information and Communication Technology, University of Tasmania, Hobart, TAS 7001, Australia; yuxuan.hu@utas.edu.au (Y.H.); quan.bai@utas.edu.au (Q.B.)

* Correspondence: lan.zhang@autuni.ac.nz (L.Z.); weihua.li@aut.ac.nz (W.L.)

Abstract: This paper introduces an LLM-Enhanced Agent-Based Influence Diffusion Simulation (LLM-AIDSim) framework that integrates large language models (LLMs) into agent-based modelling to simulate influence diffusion in social networks. The proposed framework enhances traditional influence diffusion models by allowing agents to generate language-level responses, providing deeper insights into user agent interactions. Our framework addresses the limitations of probabilistic models by simulating realistic, context-aware user behaviours in response to public statements. Using real-world news topics, we demonstrate the effectiveness of LLM-AIDSim in simulating topic evolution and tracking user discourse, validating its ability to replicate key aspects of real-world information propagation. Our experimental results highlight the role of influence diffusion in shaping collective discussions, revealing that, over time, diffusion narrows the focus of conversations around a few dominant topics. We further analyse regional differences in topic clustering and diffusion behaviours across three cities, Sydney, Auckland, and Hobart, revealing how demographics, income, and education levels influence topic dominance. This work underscores the potential of LLM-AIDSim as a decision-support tool for strategic communication, enabling organizations to anticipate and understand public sentiment trends.

Keywords: influence diffusion; LLM-enhanced social simulation; topic evolution; agent-based modelling



Academic Editor: Ludovico Minati

Received: 12 November 2024

Revised: 19 December 2024

Accepted: 28 December 2024

Published: 3 January 2025

Citation: Zhang, L.; Hu, Y.; Li, W.; Bai, Q.; Nand, P. LLM-AIDSim: LLM-Enhanced Agent-Based Influence Diffusion Simulation in Social Networks. *Systems* **2025**, *13*, 29. <https://doi.org/10.3390/systems13010029>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Information propagation describes how information, ideas, or opinions are spread through social networks [1–3]. In the context of public communication, refers to the diffusion of statements released by organisations, such as news, policies, or public announcements, through social networks [4]. This process is particularly applicable to many business companies and government organisations, as it impacts public perception and influences individuals' decision-making behaviours. Effective information propagation can enhance organisational outreach and drive significant societal impacts, while poorly crafted statements can exacerbate public concerns and even lead to crises [5]. However, a critical issue when disseminating an official statement is that the public response is somewhat unpredictable [4]. This loss of control may even lead to polarisation and extremism in social networks [1,6].

Simulating potential outcomes prior to public release can help address this challenge by summarising potential public concerns for decision-makers to consider, allowing them

to refine the message they want to convey. Influence diffusion models are widely used to simulate the information propagation process [7]. Traditional influence diffusion models are primarily probabilistic, where individuals' decision-making behaviours depend on influence probabilities [7,8]. Moreover, agent-based modelling (ABM) has been incorporated into these diffusion models to emphasise distinctive personal traits and complex interactive behaviours [2]. However, most existing ABM approaches overlook the language-level inferences and understanding capabilities of user agents. Additionally, they simulate information propagation as a passive process, where the initial influence message is transmitted through the network without considering information alteration [9], which refers to the ability of individuals to modify the content of the influence message based on their personal profiles and experiences. This limitation makes it challenging to capture and infer the language-level responses from individuals.

To address the aforementioned challenges, in this paper, we proposed an LLM-driven influence diffusion simulation tool called **LLM-AIDSim**. We apply the influence diffusion simulation technique on an LLM-driven user agent model which utilises real social networks to predict user behaviours concerning news, ideas, and topics. The primary goal is to simulate potential user comments and estimate public opinion direction before the actual publication of news content. This insight is crucial, as the impact of news and the resulting public opinions are typically measurable only after they have been propagated. By simulating reactions beforehand, this system provides valuable guidance for adjusting news content to steer public opinion effectively, thus enhancing strategic communication efforts. This approach not only anticipates public reactions but also aids in refining the messaging to better align with desired outcomes.

The contribution of our work can be summarised as follows:

- Firstly, we proposed an LLM-Enhanced Agent-Based Influence Diffusion (LLM-AID) model. To the best of our knowledge, this is the first work that combines LLMs with influence diffusion modelling. This combination addresses the limitations of traditional influence diffusion models, which are primarily probabilistic, by incorporating language-level inferences and interactions among user agents. Furthermore, instead of directly passing by a spreading influence message in existing influence diffusion models, LLM-AID emphasises individuals' responses to a public statement.
- Secondly, we developed a simulation framework, LLM-AIDSim, built on our proposed LLM-AID model. Unlike traditional influence diffusion simulations that primarily focus on numeric metrics, such as influence convergence, LLM-AIDSim provides a deeper understanding of user agents' behaviours by incorporating language-level interactions into the simulation. To address the gap in language-level analysis in influence diffusion, LLM-AIDSim integrates the Independent Cascade (IC) model with a user-agent-based approach, which was originally introduced by Kempe et al. [7]. This method closely replicates real-world diffusion scenarios, such as viral marketing or the spread of information through social media. To further enhance the simulation, we incorporate LLM-driven responses, allowing us to explore the implicit inferences derived from agents' language interactions. However, to mitigate the challenges of computational cost and the inherent randomness of LLMs, we fine-tuned the model using real-world datasets, including New York Times Articles and Comments¹. This fine-tuning ensures that LLM-AIDSim can be efficiently deployed locally, with output that closely reflects actual user posting behaviour in real-world contexts.
- Thirdly, to showcase the practical utility of LLM-AIDSim, we developed a web-based decision-support system called LLM-AIDSim system, tailored to support strategic communication. This system is designed to provide actionable insights derived from LLM-driven agent-based influence diffusion simulations, enabling organisa-

tions to develop more effective communication strategies. We conducted multiple scenarios under various application settings, demonstrating the capability of our proposed framework to interpret implicit influences at the language level, which informs decision-makers in crafting impactful communication strategies.

The rest of the paper is organised as follows: Section 2 reviews related works in the fields of LLM-driven social simulation, agent-based modelling, and influence diffusion. Section 3 presents the overall framework of our proposed LLM-AIDSim. Section 4 introduces our proposed influence diffusion model, LLM-AID. Section 5 provides details about LLM-AIDSim. Section 6 demonstrates the experiments and their results, and Section 7 concludes the paper.

2. Related Works

2.1. Influence Diffusion

Influence diffusion is an instinctive phenomenon in social networks that describes how information, opinions, ideas, or products propagate contagiously. It is critical for understanding and predicting user behaviour in contexts such as marketing, political campaigns, and information spread. Influence diffusion models are widely used to analyse these patterns within networks, allowing researchers to identify key influencers and maximise the impact of information propagation. Among these models, the Independent Cascade (IC) model, introduced by Kempe et al. [7], has become a foundational tool in social network analysis. Its simplicity and adaptability make it one of the most effective models for simulating how information, behaviours, or products spread through networks. Despite its wide usage, recent research has focused on addressing some limitations of the IC model by proposing enhancements and extensions.

For instance, Kong et al. [10] proposed the Decreasing Cascade (DC) model, which generalises the IC model by considering market saturation and decreasing activation probabilities over time. This reflects real-world scenarios where the likelihood of influencing a node diminishes after multiple unsuccessful attempts, providing a more realistic framework for applications like viral marketing, where repeated marketing efforts may lose effectiveness. However, while the DC model improves the realism of influence diffusion, it primarily focuses on reducing the activation probability over time and does not fully account for contextual user behaviours or network dynamics that may influence the spread of information. Similarly, Zhu et al. [11] introduced the Node Feature-Aware Voting Algorithm (ISVoteRank) to improve seed selection in influence maximisation. By adapting the IC model to consider node-specific interaction attributes, ISVoteRank enhances the efficiency of influence propagation, particularly in networks with temporal dynamics. While this approach optimises seed selection, it is limited by its reliance on static node attributes and does not sufficiently address the evolving nature of user preferences or external factors that might affect influence spread in real-time scenarios.

Kumar et al. [12] propose a hybrid framework that integrates the traditional Independent Cascade (IC) model with machine learning techniques to optimize seed selection. By leveraging a graph-based neural network, their approach efficiently predicts influence probabilities, thereby significantly reducing computational costs. However, while this method achieves high efficiency, it does not account for real-time user feedback. This limitation is critical for applications such as misinformation control or dynamic social environments, where user behaviour and influence probabilities evolve continuously. To address the shortcomings of static models, Liu et al. [13] introduced a continuous Independent Cascade model that incorporates real-time feedback. Their model tackles challenges such as incomplete information and dynamic network structures by incorporating node-level feedback, enhancing the accuracy of influence propagation predictions. This advancement proves

particularly effective in dynamic settings such as viral marketing and misinformation control. Nevertheless, scalability remains a challenge, especially when real-time feedback introduces significant computational overhead in large-scale networks.

Saito et al. [14] focus on improving the accuracy of influence diffusion predictions by estimating IC model parameters using real-world network data. Their enhanced probabilistic model delivers more accurate rankings of influential nodes and predictions of information spread. However, the model struggles to adapt to rapidly changing network topologies and heterogeneous user behaviours. Li et al. [15] tackle the issue of mitigating undesirable influences in online social networks. They propose an agent-based influence diffusion model combined with context-aware algorithms to effectively reduce negative impacts. By leveraging both user and influence contexts, their findings provide practical strategies for controlling the spread of information in social networks, offering valuable insights into influence dynamics.

Wang et al. [9] contribute a novel seed selection algorithm for influence maximization with minimal network modifications. Their work highlights the balance between maximizing influence and reducing disruption, which is particularly relevant for real-world applications where stability is essential.

Influence diffusion remains a pivotal tool in the study of social networks. Researchers have extensively utilized this mechanism within agent-based modelling (ABM), and recent advancements are exploring its integration into LLM-driven social simulations. As large language models (LLMs) become increasingly powerful, they present opportunities to enhance social simulations by incorporating natural language capabilities and more sophisticated decision-making processes.

2.2. Agent-Based Modelling and LLM-Driven Social Simulation

Agent-based modelling and simulation (ABM) provides a robust framework for understanding and analysing complex systems by focusing on individual agents and their interactions within an environment. Agents in ABM are heterogeneous, equipped with distinct attributes, states, and decision-making rules. Their behaviours adapt dynamically based on environmental and contextual factors, leading to emergent system-wide phenomena [16,17].

Over the past few decades, ABM has become an invaluable tool for studying intricate phenomena that are otherwise challenging to observe directly in real-world settings. It facilitates experimentation, hypothesis testing, and scenario analysis, offering insights into the behaviour of systems under diverse conditions. Applications of ABM span multiple domains, including economics, biology, sociology, and ecology, where it has proven instrumental in aiding decision making and analysing complex dynamics.

Luo et al. [18] developed a layered framework to simulate human decision-making processes, incorporating physiological, emotional, and social group attributes to shape agent behaviour. This approach offers a nuanced representation of individual decision making within social contexts. Similarly, Kountouriotis et al. [19] modelled crowd dynamics, demonstrating how individual actions aggregate into collective behavioural patterns. These studies highlight ABM's capability to simulate both micro-level (individual) and macro-level (system-wide) behaviours, providing insights into emergent trends.

ABM has also been extensively applied to model opinion dynamics. Liu and Wen [20] investigated the interplay between social connections and opinion formation, introducing models that dynamically adjust social network structures based on agent interactions. Their work provides a realistic depiction of evolving social influence and reflects the growing emphasis on hybrid models for simulating fluid social behaviours. Similarly, Li et al. [2] proposed an adaptive agent-based evolutionary model (ABEM) that integrates

genetic algorithms to optimize influencer selection in dynamic networks, achieving significant improvements in influence maximization. Flache et al. [21] reviewed theoretical advancements in social influence modelling, addressing Axelrod's seminal question of why differences in beliefs, attitudes, and behaviours persist despite the homogenizing effects of social interactions [22]. Further, Banisch et al. [23] employed Markov chains within an agent-based framework to bridge micro-level agent behaviours with macro-level opinion dynamics, offering a powerful mathematical tool for studying opinion evolution. Other works, including those by Madey et al., El-Sayed et al., and Gilbert [24–26], extend ABM to social networks, showcasing its adaptability for modelling diverse social phenomena.

Beyond social behaviours, ABM has made significant contributions in economics. Researchers have utilized agent-based models to explore economic systems [27,28], analyse market dynamics [29,30], and investigate the behaviour of individual economic agents [31]. These studies underscore ABM's ability to capture nuanced interactions among economic agents, providing insights into systemic risks and market trends.

In healthcare, ABM has been widely adopted for modelling and simulating disease spread. Cabrera et al. [32] and Barnes et al. [33] demonstrated its effectiveness in simulating public health scenarios, such as disease outbreaks [34], evaluating healthcare systems [35], and assessing intervention strategies [36]. Elsenbroich et al. [37] applied ABM to analyse societal factors influencing health inequalities, providing valuable insights into the spread of infectious diseases and behavioural responses to public health interventions. These studies showcase ABM's versatility in capturing emergent health dynamics and informing public health decision making.

Agent-based modelling and simulation (ABM) offers a versatile and powerful framework for simulating complex systems by focusing on individual agents and their interactions within specific environments. Agents in ABM are heterogeneous and equipped with unique attributes, states, and decision-making rules. By dynamically adapting their behaviours based on environmental and contextual factors, agents can collectively produce emergent patterns at the system level [16,17]. ABM provides valuable insights into areas, such as social dynamics, crowd behaviour, economics, and healthcare. However, a major challenge lies in developing generalized agents capable of performing across diverse environments. These environments vary significantly in complexity, dynamics, and uncertainty, making it difficult to create agents that are universally effective. For instance, agents optimized for opinion dynamics may not perform equally well in economic or healthcare simulations. By modelling individual agents and their interactions, ABM reveals the intricate interplay that gives rise to emergent system behaviours.

The ability to acquire and use language is a key trait that differentiates humans from other species [38]. Simulation, as a computational technique, mirrors real-world systems and processes using mathematical models, algorithms, and digital representations. ABM excels in this domain by simulating individual entities, or agents, each with defined behaviours and decision-making rules, to explore complex phenomena resulting from agent interactions [16]. Integrating large language models (LLMs) into ABM introduces a promising advancement, enhancing the capabilities of simulations [39]. LLM-based agents offer three key advantages:

- **Autonomy:** LLMs enable agents to take actions without explicit instructions [40].
- **Adaptive planning:** Agents can make dynamic and context-aware decisions, closely resembling human behaviour [41–43].
- **Realistic interactions:** LLM-driven agents can communicate effectively with both humans and other agents, facilitating rich, language-based interactions [44].

By leveraging these capabilities, LLMs enhance agent-based simulations, enabling more sophisticated and realistic representations of decision-making processes, communi-

cation, and adaptive behaviour. This integration significantly improves the fidelity and complexity of simulations, offering deeper insights into emergent system behaviours.

One of the key limitations of traditional ABM lies in the simplicity of rule-based or neural network-driven agents, which often lack the intelligence to make optimal or contextually appropriate decisions [45]. Large language model-empowered agents overcome this limitation by exhibiting advanced reasoning capabilities, enabling them to make informed decisions and select appropriate actions within simulations. Moreover, these agents enhance autonomy, a critical advantage in simulations characterized by uncertainty and dynamic environments [46].

In agent-based modelling, systems often face continuous uncertainty and limited controllability [16]. For example, in long-term social network simulations, the environment and agent states may change significantly compared to the initial conditions [47]. To remain effective, agents must adapt to evolving environments, formulating new decision policies that can deviate considerably from their initial strategies.

The integration of LLMs into ABM has been explored across both virtual and real-world environments.

Virtual environments: These simulations operate under predefined rules, such as virtual social systems or prototype-level games. For instance, Qian et al. [48] designed a virtual software company populated by agents assigned specific roles, such as CEOs, managers, and programmers. Wang et al. [49] created a virtual recommender system where agents browse recommended content and provide real-time feedback.

Real-world environments: LLM-powered agents are increasingly being applied to simulate real-world scenarios. Li et al. [50] deployed agents to represent consumers and workers within economic activities. Gur et al. [51] modelled real humans browsing and interacting with online content. Xu et al. [52] developed LLM-driven agents to simulate urban environments, capturing complex behaviours such as navigation, economic decisions, and social interactions. Recently, Park et al. [53] demonstrated the potential of LLMs through a virtual environment where 25 agents simulated daily activities using adaptive planning and decision making. Although this model was generative and lacked real-world validation, it highlighted LLMs' potential to enhance ABM.

In a distinct application, Hua et al. [54] introduced WarAgent, an LLM-based multi-agent system designed to simulate international relations and conflict scenarios. While these studies offer valuable insights, they are often limited in scale, particularly when compared to real-world social networks.

To address these challenges, we present the LLM-AIDSim framework, a scalable approach that leverages LLMs to power a larger population of agents. This framework provides a robust simulation of influence diffusion and social interactions across extensive network structures.

In our previous work, Hu et al. introduced an LLM-Enhanced Agent-Based Simulation Tool, we adapted both LLM and an influence diffusion model on the information propagation task [55]. In this paper, we extend the IC model by integrating it with large language models (LLMs) to simulate influence diffusion and language-level interactions between users. This allows us to capture contextual and semantic nuances of users' communications, addressing the gap in modelling real-time user behaviour in a social network. By utilising an LLM-driven framework, we provide a more accurate and scalable simulation of how influence propagates in large-scale networks, considering dynamic interactions and language-based feedback.

Our approach offers a significant advancement in the field of influence diffusion, particularly for applications in dynamic social media environments, strategic communication,

the influence probability. The diffusion process continues until no more inactive agents can be influenced. LLM-AID adopts this skeleton to describe influence diffusion.

Moreover, LLM-AID leverages agent-based modelling (ABM) to emphasise the unique personalities and behaviours of individual agents. Each social network user is represented as an autonomous agent with distinct traits and interactive behaviours, allowing them to respond to information based on their personality and interactions. Specifically, the LLM model “llama3:8b” is used to develop these agents, enhancing their ability to comprehend and generate language-based responses.

4.1. Formal Definitions

We first introduce the formal definitions used in the LLM-AID model.

Definition 1. A social network $G = \langle V, E \rangle$ is a directed graph composed with a set of users $V = \{v_1, v_2, \dots, v_n\}$ and a set of edges $E = \{e_{ij} | 1 \leq i \leq n, 1 \leq j \leq n, i \neq j\}$.

Definition 2. A user agent v_i represents a social network user. $\Gamma_i = \{\Gamma_i^{in} \cup \Gamma_i^{out}\}$ is the neighbour set of v_i , representing a collection of user agents that are connected to v_i . Γ_i^{in} indicates the in-neighbours from which v_i can receive messages. Γ_{out}^i refers to the set of out-neighbours to which v_i can disseminate influence messages. A user has a user profile r_i , a textual description of a user agent. A user agent holds a repository $R_j = \{R_j^{in} \cup R_j^{out}\}$, as defined in Definition 4.

Definition 3. An influence message $m_x = (c_x, v_i)$ refers to a piece of textual information, where c_x is the content of the influence message, and v_i denotes the user agent v_i who sent the message.

Definition 4. A repository $R_j = \{R_j^{in} \cup R_j^{out}\}$ is a collection of influence messages which is associated with a user agent v_j . Particularly, R_j^{in} contains all incoming influence messages directed towards the user agent v_j . While R_j^{out} includes all outgoing influence messages sent from the user agent v_j .

4.2. LLM-Enhanced User Agent Model

Algorithm 1 provides the pseudocode of the diffusion process. Lines 1–9 initialise the seed set S , the active user agent set A , and the newly active user agent set A_t , in which user agents are influenced at the last time step. Lines 10–21 start the diffusion process. Particularly, the user agents in the newly active user agent set A_t attempt to influence its out-neighbours. The newly active user agent set A_t is updated at line 22. Lines 23–28 describe the evolution process of an active user agent that is not recently activated.

To break down, the diffusion process can be demonstrated as follows:

- The diffusion process unfolds over a limited number of discrete time steps $T = \{0, 1, 2, \dots, t\}$. A user agent can be in one of two statuses: active or inactive.
- At time step 0, all user agents are inactive. Each user agent in the social network has a chance to receive a broadcast message m_s with a broadcasting probability p_b , thereby becoming active. The user agents who receive this message are denoted as the initial active user agents (or seeds).
- At time step t , an active user agent v_i that was activated at time step $t - 1$ attempts to activate each of its inactive out-neighbours $v_j \in \Gamma_{out}^i$ by sending an influence message $m_x = (c_x, v_i)$. The content c_x of the influence message is its individual response to the initial broadcast message m_s . Each inactive user agent v_j has a uniform influence probability p_i to be activated by an influence message. Once v_j is influenced at time step t , it is denoted as active and makes a response to the initial message m_s based on its personal profile r_i and the received influence message m_x .

- At time step t , an active agent that was active at time step $t - n$, where $n > 1$, has the opportunity to update its response to the initial message based on the current responses from its active in-neighbours Γ_{in}^i . This process for user agents to update responses is based on an evolution probability p_e .
- The diffusion process stops once it reaches time step t . To ensure it aligns well with the termination condition of the IC model, i.e., the diffusion process has no newly activated users, and to ensure the model has sufficient time to evolve, we normally set t as a large number.

Algorithm 1 LLM-Enhanced Agent-Based Influence Diffusion

Require: A social network $G = \langle V, E \rangle$, an initial message m_s , rounds R , time steps T , a broadcasting probability p_b , an influence probability p_i , an evolution probability p_e .

```

1: Initialise an empty user agent set  $S = \emptyset$ 
2: for all  $v_i \in V$  do
3:   if a random number in  $[0, 1) < p_b$  then
4:      $S = S \cup \{v_i\}$ 
5:     Store received message  $m_s$  to repository  $R_i^{in} = R_i^{in} \cup \{m_s\}$ 
6:   end if
7: end for
8: Set newly active user agent set  $A_t = S$ 
9: Set activated user agent set  $A = S$ 
10: for t in range(0,T) do
11:   Set a temporary influenced user agent set  $temp = \emptyset$ 
12:   for  $v_i \in A_t$  do
13:     Generate message  $m_x = (c_x, v_i)$  towards the initial message  $m_s$ .
14:     Store generated response  $m_x$  to post repository  $R_i^{out} = R_i^{out} \cup \{m_x\}$ 
15:     for  $v_j \in \Gamma_{out}^i$  do
16:       if a random number in  $[0, 1) < p_i$  then
17:          $temp = temp \cup \{v_j\}$ 
18:          $R_j^{in} = R_j^{in} \cup \{m_x\}$ 
19:       end if
20:     end for
21:   end for
22:    $A_t = temp$ 
23:   for  $v_a \in A \setminus A_t$  do
24:     if a random number in  $[0, 1) < p_e$  then
25:       Update user response  $m_y = (c_y, v_i)$  towards the initial message  $m_s$ .
26:        $R_a^{in} = R_a^{in} \cup \{m_y\}$ 
27:     end if
28:   end for
29: end for

```

At the beginning of influence diffusion, each user agent v_i is assigned a unique textual profile r_i generated by an LLM model “llama3:8b” based on a predefined statistical distribution. This profile includes fields such as name, location, age, gender, education, occupation, and description. Additionally, each user agent has an equal broadcasting probability p_b of being selected as a seed user, which activates them at the start of the diffusion process.

Once active, a user agent v_i exhibits several interactive behaviours, including generating a language-level response to an incoming influence message m_x and propagating it to its out-neighbours Γ_{out}^i in the time step following its activation. In this phase, both the user’s profile r_i and the incoming message m_x are incorporated into the prompt for the LLM to generate a personalised response.

An inactive user agent v_j decides whether to be influenced by an incoming message m_x from one of its in-neighbours Γ_{out}^i . Each inactive agent can only be activated once. To simplify this decision-making process, we adopt the approach from the IC model, assigning a uniform influence probability p_i and treating the decision as a random event.

Another important interactive behaviour of an active user agent is observing the responses of its active in-neighbours and updating its own response based on the previous one. This update process is also treated as a random event, with a uniform evolution probability p_e .

5. LLM-AIDSim: A Decision Support System via LLM-Enhanced Agent-Based Influence Diffusion Simulation

This section introduces LLM-AIDSim, a simulation framework based on our proposed LLM-AID model. LLM-AIDSim aims to simulate public responses to initial public statements such as news articles, policy announcements, or public events. By providing simulation results and their linguistic analysis, LLM-AIDSim demonstrates its ability to support decision-makers with insightful information to help them understand potential public reactions and adjust their strategies accordingly. A web-based visualisation tool was developed at this stage for end-users to define simulation parameters based on their ad hoc scenarios and to visualise the simulation.

Given a spreading influence message, the ideal output is a summary of a comprehensive semantic analysis towards the public responses to the content of the message. The public responses are composed of the messages posted by LLM-empowered user agents who are influenced by the influence message.

5.1. Initialization

5.1.1. Network Construction

LLM-AIDSim supports two ways of constructing the social network: using a synthetic network or using a real-world social network. The synthetic network is based on the Holme and Kim algorithm [57]. This algorithm generates networks where the degree distribution follows a power-law distribution, suggesting only a few nodes have a very high degree while most have a relatively low degree. This is a universal phenomenon in real-world social networks [58]. Three parameters are compulsory for setting up a synthetic network: the node size n , the number of random edges m , and the triangle probability p for adding a triangle after adding a random edge.

Meanwhile, this simulation application provides a conversion tool for users to conduct simulations based on user-defined social networks. By providing the network topological structure in a text format, the tool converts it to the format supported by the application.

5.1.2. User Profile Generation

User profiles are textual descriptions of a social network user. We adopt the Ollama API using “llama3:8b” model to generate user profiles. If no contextual information is given, the prompt is designed based on these rules: (1) A user profile is a description of a social network user within 50 words. (2) The age of these users follows a normal Gaussian distribution. (3) The gender of these users is half male and half female.

If the contextual information is given, the framework will generate the user profile based on the given contextual information. The contextual information is the demographic data collected from public census data². The age, gender, salary, education level, and occupation factors are considered in the contextual information.

5.1.3. Seed Selection

Unlike traditional influence diffusion models, which focus on selecting a specific set of initial active users as seeds, our proposed diffusion model, LLM-AID, allows user agents to have an equal probability of being randomly selected as seed users. This approach aligns more closely with real-world message diffusion, where information is often spread randomly, as messages can be read or shared by any user. In our system, these randomly selected seeds represent real-world users who are exposed to and disseminate the message. This equal probability is referred to as the broadcasting probability, representing the likelihood of a user agent receiving a broadcast message. Additionally, to ensure consistency across multiple rounds of simulations, the seed sets are fixed, using the set of seeds selected during the first round of the simulation.

5.1.4. Parameter Settings

Two parameters control the initialisation of the simulation: rounds R and time steps T . R is the number of iterations to run a simulation. Particularly, a larger R helps to account for the inherent randomness brought by influence diffusion models, providing a more accurate representation of the overall trend and behaviour of the diffusion process with the averaged results. T is the time steps within one round of a simulation. T represents the discrete intervals within each round of the simulation. In simulations, continuous time is divided into these intervals to allow iterative calculations, with each time step capturing a snapshot of the system's state. The inclusion of multiple time steps in each round simulates the progression of time in the real world during the diffusion process. Increasing T enables the simulation to model dynamic response changes over a longer period, offering a more detailed and realistic depiction of temporal dynamics.

Three diffusion model parameters control the process of the simulation. The LLM-AIDSim users can adjust those three parameters to be suitable to their own use case: a broadcasting probability p_b represents the probability of each user agent receiving the initial broadcasting message, an influence probability p_i indicates the probability of a user agent being influenced by one of its in-neighbours, and an evolution probability p_e is the percentage of the opportunity to update its response to the initial message based on the current responses from its active in-neighbour Γ_{in}^i .

Two prompts are designed for the LLM-enhanced user agent to generate or update responses with different considerations. The response generation prompt guides a user agent to create a response to the initial message based on its personal profile and the influence message that triggered it. The update prompt instructs the user agent to modify its response to the initial message, taking into account its previous response and the responses of its active in-neighbours.

5.2. Simulation

A web-based visualisation tool is developed for end-users to conduct simulations based on their own needs. Figure 2 is the screenshot of a running simulation. The left side of the web page contains the parameter setting component, where several parameters can be adjusted. The right side displays the visualisation component, allowing the influence diffusion process to be observed. Notably, the canvas is interactive: when the mouse hovers over a node, it displays the most recent post of that user agent, and when clicked, it shows the user profile of the agent.

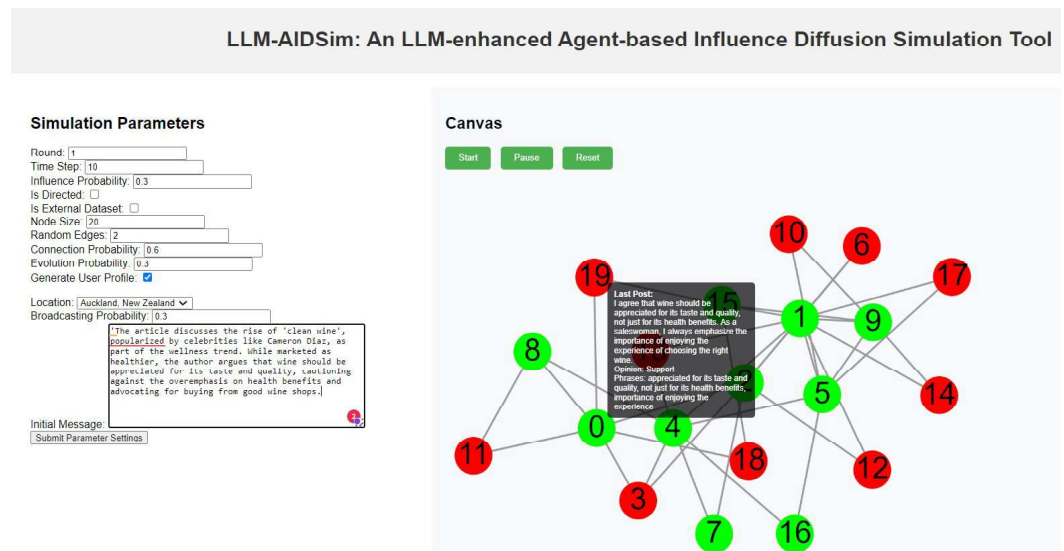


Figure 2. A screenshot of the Web-based visualisation tool.

5.3. Decision-Supporting Report

In this paper, we utilised BERT embedding to capture the semantic similarity between comments and topics [56]. Devlin introduced how BERT's contextual word embedding allows for a more nuanced understanding of the text, which, combined with cosine similarity, helps quantify how closely related two pieces of text are. In addition, we adopt K-means clustering for grouping comments and discussions into distinct clusters. This unsupervised learning method is a classic technique for finding patterns in data, often used in topic modelling for social simulations [59].

- **Evolution of topics:** The report outlines how topics evolve over various time steps across different locations. This enables users to monitor trends and changes in topic discussions over time.
- **Pie charts:** For each specific location, pie charts are provided, which detail the topics and their respective percentages. This visual format aids users in quickly understanding the prominence of each topic within a location.
- **Topic distribution table:** A table listing the topic distributions for each location is included. This table offers a comprehensive overview, allowing users to relate regional contexts to variations in focus and interests effectively. It also supports an analysis based on demographic, educational, and occupational data provided by users from each region.

6. Experiments

6.1. Experiment Details

To demonstrate the effectiveness of our simulation framework, LLM-AIDSim, we conducted four experiments using real-world news topics within a social network context. Our quantitative analysis focuses on the following.

- **Simulated vs. real-world comments:** Analysing the differences between the simulated comments generated by the LLM-AIDSim framework and real-world comments on the same topics (see Section 6.3).
- **Topic evolution and converge:** Tracking the evolution of topics and the extent of discussion converge in the LLM-AIDSim framework (see Section 6.4).
- **Clustered topic analysis:** Evaluating the clustered topics of the simulation model with and without diffusion (see Section 6.5).

- **Framework performance across different demographics:** Assessing the performance of our framework across various population composition ratios and regional contexts (see Section 6.6).

6.2. Experiments Settings

- **Network parameters.** In this experiment, we utilised a subset of the Stanford Network Analysis Project (SNAP) group's³ ego-Facebook dataset as our network nodes. The SNAP group, led by Prof. Jure Leskovec, is a renowned research group at Stanford University dedicated to studying large-scale networks and their applications in various domains, including social networks, information networks, and biological networks. In our study, the total number of nodes (representing individual users) is 333, and the number of edges (representing connections or relationships between users) is 2500 [60]. This substantial dataset allowed us to analyse the network structure, identify influential nodes, and explore various network properties and phenomena.
- **Parameter settings.** To maintain comparability across the three experiments, we applied identical diffusion model parameters in each scenario. Specifically, the broadcasting probability p_b , the influence probability p_i , and the evolution probability p_e were all set to 0.3. Similarly, the simulation parameters were consistent across the experiments, with the number of rounds R set to 5 and the number of time steps T established at 15. These settings ensured that each experiment was evaluated under the same conditions, allowing for a fair comparison of results.
- **Comment generation.** In the context of large language models (LLMs), temperature is a parameter used to control the randomness of the model's output. It affects the probability distribution from which the model samples the next word or token, and the temperature range is from 0 to 1. In our study, all experiments were conducted using the Ollama API, with the model version set to "llama3:8b" and a temperature setting of 0 to disable the randomness completely; the purpose was to ensure reproducibility and the most deterministic behaviour. The chosen model was selected for its superior performance in providing accurate and consistent results. The use of APIs greatly facilitates interaction with the model, streamlining the research process. In the experiments, while generating the user profile, we utilised contextual information, including three cities from two countries: "Auckland, New Zealand", "Sydney, Australia", and "Hobart, Australia".
- **Real-world comments collection.** A real-world dataset, the New York Times Articles and Comments⁴, was utilized in this study. This dataset comprises all comments and articles published by the New York Times between 1 January 2020 and 31 December 2020. For our experiment, we selected a subset corresponding to the article titled "Should Wine Be Among Your Health Resolutions"⁵. The comments in this subset were contributed by real-world users in response to the selected article. To generate the initial topic content for our experiment, refer to Figure 3; we extracted the news article from the dataset and employed the Llama 3 large language model to summarise the article. The output summarised news content was used as the initial topic content in the simulation process.
- **Topic clustering.** Figure 3 illustrates the process of the topic clustering. Firstly, we employed the BERT model [56] to convert both real-world comments and simulated comments into vectors. Secondly, we used the k-means clustering algorithm [59] to group topics from real-world user comments and our simulated comments. The k-means algorithm, a widely used and efficient unsupervised learning clustering method, was employed to partition the comments into distinct clusters based on their content, identifying the underlying themes that emerged from real and simulated discussions.

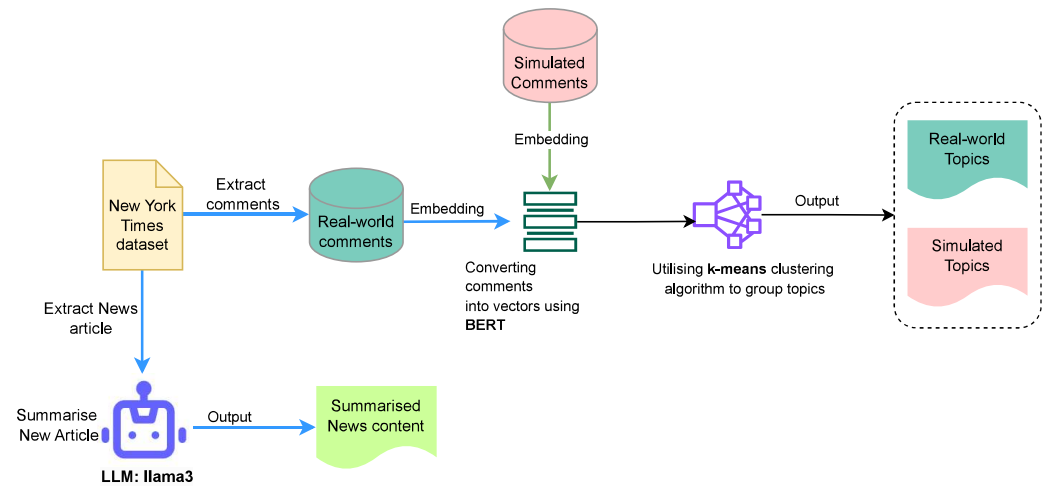


Figure 3. Data processing workflow.

Evaluation Metrics

The semantic similarity between these clusters was then calculated using a state-of-the-art language model, BERT. BERT allows for the generation of high-dimensional embedding that captures the nuanced semantic meanings of the comments [56]. By leveraging BERT’s contextual word representations, we converted each real-world topic and simulated topic that was grouped by the k-means clustering algorithm into a fixed-length vector or embedding, representing its semantic content. These vectors exist in the same semantic space, ensuring consistency and comparability between real-world and simulated topics. Then, the similarity between these embeddings was quantified using cosine similarity, a metric that measures the cosine of the angle between two vectors in a multi-dimensional space. Cosine similarity values approaching 1 signify a strong similarity between the vectors, indicating a close alignment between the topics they represent.

6.3. Experiment 1: Comparison of Real-World Topics and Simulated Topics

Figure 4 illustrates the evolving semantic similarity across iterations between topics extracted from real-world comments and those generated by simulations across three distinct regions. The graph demonstrates how each iteration accumulates data from all prior time steps, leveraging the highest similarity scores associated with each real-world topic to compute the Maximum Average Similarity Score for every iteration.

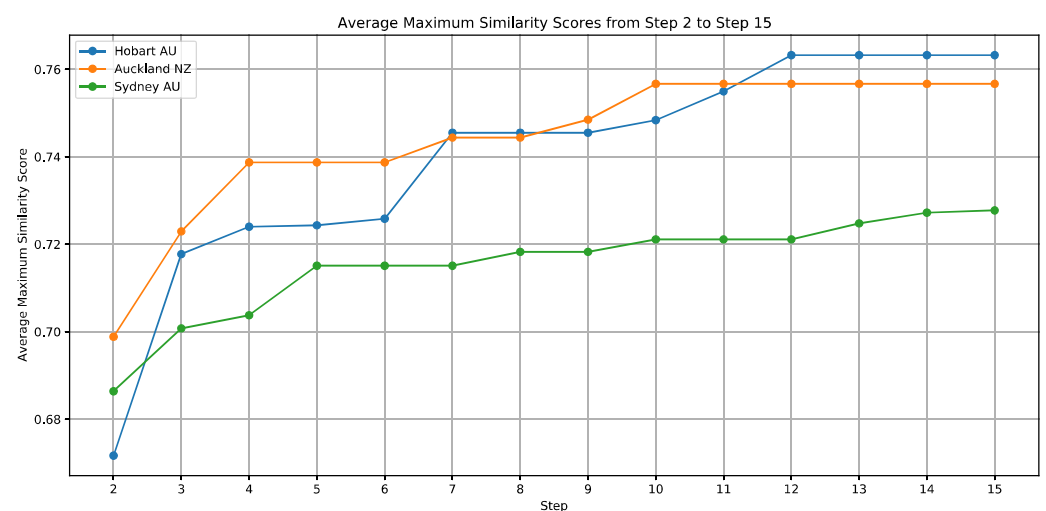


Figure 4. Cumulative evolution of topic similarity across iterations.

As the iterations progress, the plot reveals an upward trend in the average maximum similarity score, indicating that the simulated topics increasingly align with real-world topics over time. This trend underscores the capability of our LLM-AIDSim framework to generate topics that accurately reflect real user discussions, thereby enhancing the relevance and scope of our simulations.

The consistent increase in similarity scores throughout the diffusion process highlights the dynamic and accumulative nature of our simulation approach. The broadening converge of topics, along with the convergence towards higher similarity with real-world discussions, attests to the realism and reliability of the LLM-AIDSim framework. The ability of the framework to replicate the nuances and focal points of actual user discussions is further evidenced by the uniform semantic similarities observed across different regions.

These findings validate the robustness and adaptability of the LLM-AIDSim framework, emphasising its potential as a strategic tool for analysing and forecasting user engagement and topic trends on social networks. Its utility extends to strategic communication planning, where anticipating public responses and understanding evolving discussions can significantly impact decision making and outcomes.

In the above experiments with the LLM-AIDSim framework, which leverages user profiles to reflect topics from different cities, we conducted tests using three distinct sets of user profiles. Additionally, we evaluated multiple large language models (LLMs) using the same prompt applied in our LLM-AIDSim framework to generate simulated topics. Unlike LLM-AIDSim, the LLMs cannot utilise user profiles, so only one set of results was generated per LLM.

For the output topics of each LLM, we used the evaluation metrics method to compute similarity scores between each real-world topic and its corresponding simulated topics. The highest similarity score for each real-world topic was then selected to calculate the Maximum Average Semantic Similarity Score for each LLM.

Table 1 presents the Maximum Average Similarity Score between real-world and simulated topics without the diffusion technique. Among the five evaluated LLMs, GPT-4.0 achieved the highest performance with a score of 0.5897, while Jurassic-2 Ultra obtained the lowest score of 0.5004. Notably, all these Maximum Average Similarity Scores are lower than those observed at time step 2 in our LLM-AIDSim framework.

Table 1. Maximum Average Similarity score between real and simulated topics without diffusion.

LLM Name	Maximum Average Similarity
GPT-4o	0.5897
Claude 3 Sonnet	0.5701
Llama3 8B	0.5700
Mistral Large	0.5697
Jurassic-2 Ultra	0.5004

In our LLM-AIDSim framework, we employed random seeds to facilitate influence diffusion, resulting in a more diverse set of topics. By time step 2, multiple user agents had been activated, leading to greater variation in the simulated topics. This demonstrates the framework's effectiveness in replicating real user behaviour and preferences, further validating its utility in simulating complex social dynamics.

6.4. Experiment 2: Topic Evolution Tracking and Diffusion Convergence

The nodes in the images represent topics at each time step, and the flows between nodes denote the similarity between the two topics. The similarity is calculated by taking the cosine similarity of the embedding of a pair of topics, where the embedding is computed

using the BERT model [56]. A wider and darker flow indicates a higher similarity between the two topics. Only topics with a similarity score above 0.85 are shown to simplify the visualisation.

As shown in Figure 5, the evolution of simulated topics in Sydney demonstrates how our LLM-AIDSim framework enables topics to converge throughout the diffusion process. From time step 2 to time step 15, broader discussions gradually split into more specific and refined topics. For example, early themes such as “Finding a balance between wellness and pleasure in modern trends and habits” evolved into more focused topics like “Exploring the cultural significance of the ‘Clean wine’ trend” and “Celebrities’ influence on the popularity of ‘clean wine’ and its cultural significance”. By time step 15, the topics converged towards a unified consensus, emphasising quality and taste over perceived health benefits in wine, showcasing the framework’s capacity to extract key social media consensus points as the simulation progresses.

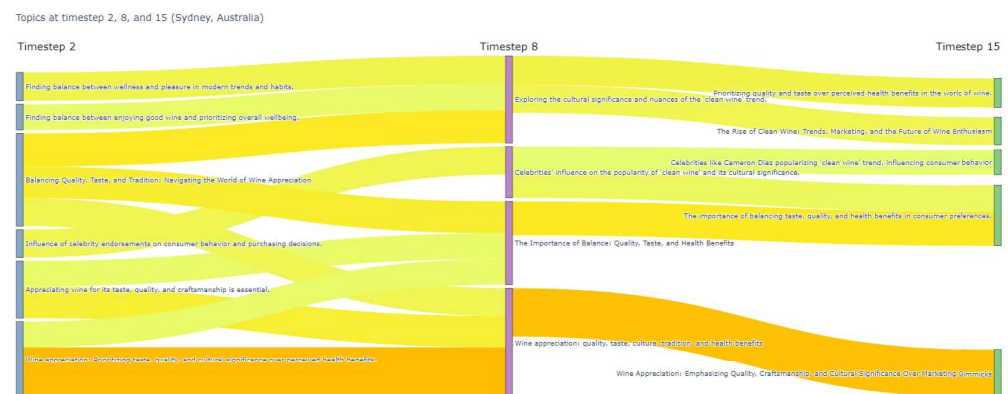


Figure 5. Topic evolution for Sydney, Australia.

Figure 6 illustrates the topic evolution for Auckland, New Zealand. Initially, the topics centred on prioritising taste and quality over health claims and marketing gimmicks. As the iterations advanced, the framework simulated the rise of the “clean wine” trend, focusing on balancing health, quality, and tradition in the wine industry. By time step 15, the topics settled on the appreciation of wine’s taste, quality, and cultural significance. This indicates that our framework allows topics to mature and converge, aligning simulated discussions with real-world consensus as diffusion progresses.

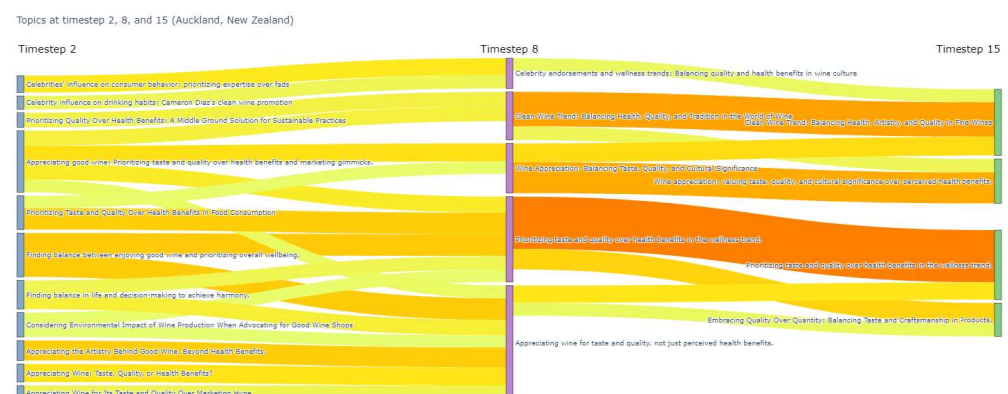


Figure 6. Topic evolution for Auckland, New Zealand.

In Figure 7, we observe the topic evolution for Hobart, Australia, where early discussions focused on sustainable wine production and balancing quality with environmental impact. As the simulation progressed, topics related to the rise of the “clean wine” trend and its marketing strategies emerged. By the final time step, the discussions converged

on themes of wine appreciation, craftsmanship, and the sensory experience of wine. This demonstrates the framework’s ability to simulate how conversations shift towards more representative, consensus-driven conclusions.




Figure 7. Topic evolution for Hobart, Australia.

What sets our LLM-AIDSim framework apart is its ability to simulate how social media discussions evolve and converge into key consensus points over time. This capability ensures that the conclusions drawn from the simulation are more representative and robust. Furthermore, the framework’s adaptability allows us to incorporate demographic variations across different regions. By using regional population data in our experiments, we were able to demonstrate how the framework performs effectively across different parameter settings, highlighting its robustness in handling diverse contexts.

While all three regions—Sydney, Auckland, and Hobart—shared a common focus on wine appreciation, quality, and cultural significance over marketing tactics and health benefits, there were subtle regional differences.

For instance, Sydney and Auckland emphasize the “clean wine” trend and its marketing aspects, likely due to their similar educational demographics and occupational structures. In contrast, discussions in Hobart focus more on sustainable production practices and quality control, which may be attributed to the city’s higher education levels and a greater proportion of individuals in professional or managerial roles. These regional variations underscore the framework’s adaptability to diverse demographic dynamics, enabling it to generate insightful and region-specific thematic evolutions.

The results from these experiments demonstrate that the LLM-AIDSim framework not only simulates discussions that converge towards widely accepted social media narratives but also excels in simulating regional differences with the same level of accuracy and representativeness. This highlights the framework’s potential as a powerful tool for studying social media dynamics across diverse regions and demographic groups.

6.5. Experiment 3: Comparison of Results with and Without Diffusion

In this experiment, we analyse the differences in topics extracted from user agents before and after diffusion, demonstrating the advantages of our LLM-AIDSim framework. Figures 8–13 show that in the absence of diffusion, users exhibit a wider range of opinions. However, after diffusion, the diversity of opinions is significantly reduced, showcasing the impact of the diffusion mechanism on consolidating user-generated content. This experiment highlights the ability of our framework to model the social influence process, where discussions shift toward dominant narratives.

The comparison of topic clusters across the three regions (Sydney, Auckland, and Hobart) reveals several key patterns. Post-diffusion, the clustering of topics shows a noticeable consolidation, where fewer topics dominate the conversation. This illustrates

how LLM-AIDSim effectively captures the dynamics of influence diffusion, in which collective attention becomes focused on a smaller set of prevailing topics.

In Sydney, for instance, topic 2, “The Rise of Clean Wine: Trends, Marketing, and the Future of Wine Enthusiasm”, and topic 5, “Wine Appreciation: Emphasising Quality, Craftsmanship, and Cultural Significance Over Marketing Gimmicks”, collectively account for over 50% of the discussion after diffusion. Similarly, in Auckland, post-diffusion analysis reveals that topic 2 (“Appreciating the Art of Winemaking and Its Cultural Significance”) and topic 3 (“Wine Appreciation: Valuing Taste, Quality, and Cultural Significance”) dominate the discourse, representing 46% of the conversation. In Hobart, topic 1 (“Prioritising Quality Over Health Benefits in Wellness Trends”) and topic 3 (“The Joy and Significance of Wine Appreciation”) emerge as the dominant topics, accounting for more than 44% of the discussion.

Across all regions, the diffusion process consolidates discussions about a few key topics, suggesting that our framework accurately simulates the influence of social networks in narrowing the focus of public discourse. As diffusion progresses, discussions converge on central themes, likely influenced by key actors or prevailing opinions within the network.

Furthermore, the reduction in diversity between pre- and post-diffusion topic clusters underscores the importance of accounting for social influence in modelling public discourse. Before diffusion, topic clusters in each region cover a broad spectrum of user interests, including themes such as sustainability, wellness, and celebrity endorsements. For instance, in Sydney, topics related to health, wellness, and wine culture are more evenly distributed prior to diffusion. In Auckland, there are discussions about local winemaking, environmental impact, and balancing quality with health benefits. In Hobart, the topics range from craftsmanship to sustainability and wine appreciation.

Through the diffusion process, our framework demonstrates its ability to reflect the natural reduction in topic diversity that typically occurs in social networks. Certain narratives are amplified, while others are diminished, as social influence drives collective attention toward a limited set of dominant themes.

In conclusion, this experiment demonstrates that the comparison between topics extracted without and with diffusion reveals the significant impact of the diffusion process on topic clustering within a simulated social network. Initially, without diffusion, the topic clustering reflects a broad spectrum of user interests, encompassing diverse themes and perspectives. After diffusion, the discussions become concentrated around a few dominant topics, mirroring real-world network dynamics, where social influence amplifies specific narratives.




Figure 8. Topic clusters of Sydney without diffusion.




Figure 9. Topic clusters of Sydney with diffusion.




Figure 10. Topic clusters of Auckland without diffusion.




Figure 11. Topic clusters of Auckland with diffusion.




Figure 12. Topic clusters of Hobart without diffusion.

These findings highlight the importance of accounting for diffusion effects when modelling and interpreting social responses to news, opinions, and topics in real-world environments. By incorporating diffusion, our LLM-AIDSim framework captures how collective discourse is shaped by social interaction and influence, providing valuable insights into the mechanisms behind topic convergence and consensus formation.




Figure 13. Topic clusters of Hobart with diffusion.

6.6. Experiment 4: Comparison of Simulation Results from Different Regions

In this experiment, we focus on comparing the simulation results from different regions, namely, Sydney, Auckland, and Hobart. The goal of this experiment is not just to analyse the simulated topics within the regions but to demonstrate the adaptability and strengths of our LLM-AIDSim framework when applied across diverse demographic and regional contexts, the results are illustrated in Tables 2–4. By incorporating population, income, and educational data into our simulations (Appendix C), we emphasise how the framework remains robust across varying parameters.

The critical feature of our framework is that topics tend to converge during the diffusion process. Over time, the discussions narrow down to a few dominant topics, reflecting a form of consensus. This behaviour highlights one of the core advantages of the LLM-AIDSim framework: its ability to replicate real-world social consensus-building while accounting for regional differences. The diffusion process inherently reduces the diversity of topics, but this convergence leads to more representative and compelling outcomes, particularly when modelling social networks.

Table 2. Sydney, Australia: Topic distribution.

Topic No.	Topic	Percentage
Topic 5	Wine appreciation: emphasising quality, craftsmanship, and cultural significance over marketing gimmicks	29.2%
Topic 3	Prioritising quality and taste over perceived health benefits in the world of wine	25.1%
Topic 2	The rise of clean wine: trends, marketing, and the future of wine enthusiasm	20.7%
Topic 6	Appreciating quality over quantity in food and drink industry trends	12.8%
Topic 4	The importance of balancing taste, quality, and health benefits in consumer preferences	8.4%
Topic 1	Celebrities like Cameron Diaz popularising the ‘clean wine’ trend, influencing consumer behaviour	3.8%

Another key advantage of our framework is its flexibility. By incorporating the unique population structures of different regions, we have demonstrated that our model performs well across various scenarios. For instance, the demographic composition in Sydney is characterised by a large working-age population, with a high percentage of professionals and well-educated individuals, which influences the emergence of topics around health-

conscious behaviours and artisanal products. In Auckland, the focus shifts more towards wine appreciation and cultural significance, likely due to the younger population and high levels of educational attainment. Meanwhile, in Hobart, the emphasis is on balancing quality with wellness trends, reflecting the local community's preference for sustainable and high-quality products.

Table 3. Auckland, New Zealand: Topic distribution.

Topic No.	Topic	Percentage
Topic 3	Wine appreciation: valuing taste, quality, and cultural significance over perceived health benefits	26.9%
Topic 2	Appreciating the art of winemaking and its cultural significance	20.9%
Topic 1	Clean wine trend: balancing health, artistry, and quality in fine wines	15.8%
Topic 5	Appreciating wine for its taste, quality, and cultural significance	12.6%
Topic 6	Embracing quality over quantity: balancing taste and craftsmanship in products	12.8%
Topic 4	Prioritising taste and quality over health benefits in the wellness trend	7.6%
Topic 7	Celebrities like Cameron Diaz promote clean wine for health benefits and quality	3.4%

Table 4. Hobart, Australia: Topic distribution.

Topic No.	Topic	Percentage
Topic 6	Appreciating wine for quality, taste, and craftsmanship over just its health benefits	23.5%
Topic 3	Appreciating good wine for its taste, quality, and cultural significance	21.1%
Topic 2	Wine quality vs. clean wine trend: balancing tradition with modernity	18.9%
Topic 4	Wine appreciation: celebrating artistry, craftsmanship, and unique characteristics	13.8%
Topic 5	Finding balance and pleasure in life through mindfulness and moderate indulgence	9.1%
Topic 1	The importance of prioritising quality over perceived health benefits in product development and marketing	7.5%
Topic 7	The importance of quality and taste in appreciating good craftsmanship	6.1%

In conclusion, Experiment 4 illustrates the versatility and adaptability of our LLM-AIDSim framework in generating realistic simulations of social influence diffusion across different regions. The framework not only simulates the dynamics of topic convergence but also highlights how regional contexts influence the specific themes that emerge during the diffusion process. These findings validate the robustness of our approach in capturing social dynamics and underscore its potential utility for strategic communication and predictive analysis in social networks.

6.7. Experiment Summary

Through our detailed analysis of simulated diffusion processes, we have demonstrated that social influence plays a pivotal role in consolidating discussions about a smaller set of dominant topics, particularly in the post-diffusion stages. This outcome highlights the effectiveness of our LLM-AIDSim framework in modelling how social dynamics shape collective discourse. As certain narratives gain prominence, others recede, reflecting the real-world behaviour of opinion formation and consensus-building in social networks.

The regional variations observed across Sydney, Auckland, and Hobart underscore the adaptability of our framework in accounting for demographic, economic, and cultural factors. In Sydney, the focus on health-conscious behaviours and artisanal products, driven by a professional, well-educated demographic, led to the dominance of topics related to clean wine and wine appreciation. In Auckland, with its younger, highly educated population, topics emphasising cultural significance and the art of winemaking took precedence. Meanwhile, in Hobart, the local preference for balancing quality with wellness and sustainability led to topics centred on the appreciation of craftsmanship and mindful consumption.

These findings validate the robustness of our framework in capturing the nuances of regional public discourse. By simulating the influence diffusion process across diverse demographic and socioeconomic landscapes, we provide insights into how collective opinions are shaped, consolidated, and influenced by social interactions. The LLM-AIDSim framework thus proves to be a versatile tool, capable of modelling and predicting topic trends across various contexts, making it invaluable for strategic communication, trend analysis, and understanding social dynamics in real-world environments.

7. Conclusions and Future Work

In this paper, we present the LLM-AIDSim framework, an innovative approach that incorporates large language models (LLMs) into agent-based simulations for influence diffusion. This framework enables agents to generate language-level responses, significantly enhancing the realism of insights into user agent behaviour during information diffusion. This integration effectively bridges the gap between traditional numerical simulation results and the textual nuances typical in real-world discussions. Through a series of experiments, LLM-AIDSim has been shown to accurately simulate user interactions and the evolution of topics within social networks, yielding results that align closely with actual behaviours and discussions, as confirmed by high semantic similarity scores across various regions.

The LLM-AIDSim framework distinguishes itself by facilitating dynamic simulations in which agent responses adapt based on ongoing discussions. This feature not only captures the intricate nature of communication within social networks but also yields practical insights for strategic communication planning. The fidelity of the simulation results to real-world data highlights the framework's potential utility across a broad spectrum of applications, from marketing strategies to public policy development. These capabilities make LLM-AIDSim an invaluable tool for researchers and practitioners aiming to understand and influence public opinion and behaviour in complex social environments.

While LLM-AIDSim provides valuable insights into the dynamics of influence diffusion, there are several avenues for future research and enhancement:

- One key area for future work is to explore how modifications to the original news article impact user behaviour, particularly in terms of the sentiment and concerns expressed in user comments. This suggests that certain alterations to the content or framing of the news can potentially reduce the number of negative users.
- Scalability and performance optimisation: As the complexity and scale of social networks grow, it is essential to explore methods to optimise the computational efficiency of LLM-based simulations. Incorporating distributed computing or more lightweight LLM architectures could allow for simulations at a much larger scale.
- Agent behaviour customisation: While this study employed LLM-generated responses to simulate human behaviour, there is room for improvement in agent customisation. Future work could allow for more detailed user profiles and behaviour models based on specific sociocultural, political, and psychological factors.
- Enhanced linguistic analysis: While the current framework utilises BERT-based embeddings [56] to measure semantic similarity, future research could benefit from exploring both advanced linguistic models and qualitative interpretive methods. Advanced models, such as those based on the latest developments in deep learning, could improve the numerical representation of content and tone. At the same time, drawing from the qualitative, interpretive tradition of empirical social research offers complementary avenues for understanding. For instance, simulated narratives could be analysed using sequence analytical procedures inspired by objective hermeneutics [61]. This dual approach could enrich insights into the discussions by combining the pre-

cision of numerical embeddings with the depth of interpretive methods, fostering a more holistic understanding of social dynamics in simulations.

By addressing these areas, LLM-AIDSim could become a more robust and flexible tool for modelling complex social dynamics, contributing to both academic research and practical applications in strategic communication, marketing, and public policy.

Author Contributions: Conceptualization, L.Z. and Y.H.; methodology, L.Z. and Y.H.; software, L.Z. and Y.H.; validation L.Z., Y.H., W.L. and Q.B.; formal analysis L.Z. and Y.H.; investigation, L.Z. and Y.H.; writing—original draft preparation, L.Z. and Y.H.; writing—review and editing, W.L., Q.B. and P.N.; visualization, L.Z. and Y.H.; supervision W.L., Q.B. and P.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

LLM-AIDSim	LLM-Enhanced Agent-Based Influence Diffusion Simulation
ABM	Agent-based modelling
IC	Independent Cascade
DC	Decreasing Cascade
LLMs	Large language models

Appendix A. Prompt

Appendix A.1. Prompt Used for Generating User Profile

Generate user profiles as a 50-word unique description. The ages of these users follow a Gaussian distribution, and gender is equally distributed. List the responses by user ID, starting with 'N', in JSON format. Use the following example as a guide:

```
"N1": {
  "name": "Emily",
  "age": 32,
  "gender": "female",
  "description": "A passionate artist who loves expressing
    herself through paintings. She finds inspiration in
    nature and often exhibits her artwork in local
    galleries."
}
```

An example of generated user profile:

```
{
  "N1": {
    "Name": "Liam",
    "Location": "Auckland, New Zealand",
    "Age": "41",
    "Gender": "Male",
    "Education": "Master's degree",
    "Occupation": "Manager",
```

```

        "Description": "Liam is a driven manager who enjoys
            leading teams to success. He loves playing rugby
            in his free time and is an avid fan of the All
            Blacks."
    }
}

```

Appendix A.2. Prompt Used for Generating User Response

We are building an influence discussion simulation tool, you are a responsible AI, and your task is based on each user's profile, adapt the user's post habit and generate one comment on the topic you received. The format and content should follow the following 3 instructions. Please note, do not generate duplicate sentences from different users. One profile represents one user and one possible comment only. Based on the USER_PROFILE, given the user's last received influence message LAST_POST_MESSAGE, given the topic 'topic', please perform the following tasks and provide the responses in JSON format:

```

{
    1. Generate one of the user's possible comments while the
        user reading this topic, in this format: 'Response:
        [User's response]', each response cannot be the same.
    2. Do a semantic analysis based on the response generated
        from 1, output the results as this user's opinion in
        the format of: 'opinion: [Support/Oppose/Neutral]'
    3. Summarize the user's response and generate the
        response in the format: 'phrases: [List of phrases]'

```

Please only return the responses in the following JSON format, one response only for each profile:

```

{
    "response": "[User's response]",
    "opinion": "[Support/Oppose/Neutral]",
    "phrases": "[List of phrases]"
}
}

```

An example of generated user response:

```

{
    "response ": "I agree that wine's taste and quality
        should be appreciated. As someone who loves organizing
        events, I've seen how 'clean wine' has become trendy
        in social gatherings. While it's great to prioritize
        health, I think it's also important to consider the
        cultural aspect of enjoying a good glass of wine.
        After all, life is about balance! ",
    "opinion ": "Support ",
    "phrases ": [
        "I agree ",
        "prioritize the cultural aspect ",
        "life is about balance "
    ]
}

```

Appendix B. Initial Topic

Summarizing the initial topic used in this paper: The article discusses the rise of ‘clean wine’, popularised by celebrities like Cameron Diaz, as part of the wellness trend. While marketed as healthier, the author argues that wine should be appreciated for its taste and quality, cautioning against the overemphasis on health benefits and advocating for buying from good wine shops.

Appendix C. Population Information for Different Regions

Table A1. Demographic, income, education, and occupation comparison between Hobart, Auckland, and Sydney.

Category	Hobart, Australia	Auckland, New Zealand	Sydney, Australia
Population	Working-age population: 68.5%	Median age: 34.7	Working-age population: 64.6%
	Males: 48.94%	Males: 49.43%	Males: 49.71%
	Females: 50.98%	Females: 50.57%	Females: 50.28%
Personal Weekly Income	NZD 1–499: 20.9%	NZD 1–499: 6.2%	NZD 1–499: 20.4%
	NZD 500–999: 23%	NZD 500–999: 4.9%	NZD 500–999: 20.3%
	NZD 1000–1999: 27.8%	NZD 1000–1999: 6.4%	NZD 1000–1999: 26.4%
	NZD 2000–2999: 8.9%	NZD 2000–2999: 8.1%	NZD 2000–2999: 8.9%
	NZD 3000 or more: 6.1%	NZD 3000 or more: 6.5%	NZD 3000 or more: 7.1%
	Nil income: 7.6%	Nil income: 8.7%	Nil income: 9.7%
	Negative income: 0.4% Inadequately described/not stated: 5.4%	Loss: 0.6% Inadequately described/not stated: 8.7%	Negative income: 0.7% Inadequately described/not stated: 6.5%
Education	Secondary school or lower: 18.9%	No qualification: 14.5%	Secondary school or lower: 26.2%
	Postgraduate degree: 16.7%	Post-graduate and honours degrees: 6.6%	Postgraduate degree: 9.4%
	Graduate diploma/certificate: 4.5%	Master’s degree: 5.0%	Graduate diploma/certificate: 2.1%
	Bachelor degree: 27.8%	Bachelor’s degree: 18.6%	Bachelor degree: 21.8%
	Advanced diploma/diploma: 8.1%	Level 5 diploma: 4.7%	Advanced diploma/diploma: 9.7%
	Certificate: 10.6%	Level 3 certificate: 11.5%	Certificate: 14.8%
	Inadequately described: 0.7% Not stated: 6.0%	Inadequately described: 0.7% Not stated: 6.0%	Inadequately described: 0.8% Not stated: 7.2%
Occupation	Managers: 14.3%	Managers: 18.1%	Managers: 15.2%
	Professionals: 36.7%	Professionals: 25.9%	Professionals: 29.3%
	Technicians/trades workers: 8.8%	Technicians/trade workers: 11.4%	Technicians/trades workers: 10.5%
	Community/personal service workers: 12.8%	Community/personal service workers: 9.1%	Community/personal service workers: 9.3%
	Clerical/administrative workers: 10.7%	Clerical/administrative workers: 11.7%	Clerical/administrative workers: 13.8%
	Sales workers: 7.1%	Sales workers: 10.0%	Sales workers: 7.8%
	Machinery operators/drivers: 2.0% Labourers: 6.1%	Machinery operators/drivers: 5.9% Labourers: 7.9%	Machinery operators/drivers: 5.6% Labourers: 6.8%

Notes

- 1 <https://www.kaggle.com/datasets/benjaminawd/new-york-times-articles-comments-2020?select=nyt-comments-part1.csv> (accessed on 29 May 2024)
- 2 <https://www.abs.gov.au/census> (accessed on 10 March 2024); <https://www.stats.govt.nz/topics/census> (accessed on 16 June 2024).
- 3 <http://snap.stanford.edu/> (accessed on 10 March 2024).
- 4 <https://www.kaggle.com/datasets/benjaminawd/new-york-times-articles-comments-2020?select=nyt-comments-part1.csv> (accessed on 8 April 2024).
- 5 <https://www.nytimes.com/2020/12/31/dining/drinks/clean-wine-avaline.html> (accessed on 8 April 2024).

References

1. Prasetya, H.A.; Murata, T. A model of opinion and propagation structure polarization in social media. *Comput. Soc. Netw.* **2020**, *7*, 1–35. [\[CrossRef\]](#)
2. Li, W.; Hu, Y.; Jiang, C.; Wu, S.; Bai, Q.; Lai, E. ABEM: An adaptive agent-based evolutionary approach for influence maximization in dynamic social networks. *Appl. Soft Comput.* **2023**, *136*, 110062. [\[CrossRef\]](#)
3. Hu, Y.; Wu, S.; Jiang, C.; Li, W.; Bai, Q.; Roehrer, E. AI Facilitated Isolations? The Impact of Recommendation-based Influence Diffusion in Human Society. In Proceedings of the IJCAI, Vienna, Austria, 23–29 July 2022; pp. 5080–5086.
4. Zhang, L.; Wang, T.; Jin, Z.; Su, N.; Zhao, C.; He, Y. The research on social networks public opinion propagation influence models and its controllability. *China Commun.* **2018**, *15*, 98–110. [\[CrossRef\]](#)
5. Kanso, A.M.; Nelson, R.A.; Kitchen, P.J. BP and the Deepwater Horizon oil spill: A case study of how company management employed public relations to restore a damaged brand. *J. Mark. Commun.* **2020**, *26*, 703–731. [\[CrossRef\]](#)
6. Franks, D.W.; Noble, J.; Kaufmann, P.; Stagl, S. Extremism propagation in social networks with hubs. *Adapt. Behav.* **2008**, *16*, 264–274. [\[CrossRef\]](#)
7. Kempe, D.; Kleinberg, J.; Tardos, É. Maximizing the spread of influence through a social network. In Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, 24–27 August 2003; pp. 137–146.
8. Li, Y.; Chen, W.; Wang, Y.; Zhang, Z.L. Influence diffusion dynamics and influence maximization in social networks with friend and foe relationships. In Proceedings of the Sixth ACM International Conference on Web Search and Data Mining, Rome, Italy, 4–8 February 2013; pp. 657–666.
9. Wang, G.; Li, W.; Bai, Q.; Lai, E.M. Maximizing Social Influence With Minimum Information Alteration. *IEEE Trans. Emerg. Top. Comput.* **2023**, *12*, 419–431. [\[CrossRef\]](#)
10. Kong, F.; Xie, J.; Wang, B.; Yao, T.; Li, S. Online Influence Maximization under Decreasing Cascade Model. *arXiv* **2023**, arXiv:2305.15428.
11. Zhu, W.; Miao, Y.; Yang, S.; Lian, Z.; Cui, L. Maximizing Influence in Temporal Social Networks: A Node Feature-Aware Voting Algorithm. *Comput. Mater. Contin.* **2023**, *77*. [\[CrossRef\]](#)
12. Kumar, S.; Mallik, A.; Khetarpal, A.; Panda, B.S. Influence maximization in social networks using graph embedding and graph neural network. *Inf. Sci.* **2022**, *607*, 1617–1636. [\[CrossRef\]](#)
13. Liu, C.; Xu, H.; Liu, X. Online influence maximization under continuous independent cascade model with node-edge-level feedback. *Knowl. Inf. Syst.* **2024**, *66*, 1091–1110. [\[CrossRef\]](#)
14. Saito, K.; Nakano, R.; Kimura, M. Prediction of information diffusion probabilities for independent cascade model. In Proceedings of the International Conference on Knowledge-Based and Intelligent Information and Engineering Systems, Zagreb, Croatia, 3–5 September 2008; pp. 67–75.
15. Li, W.; Bai, Q.; Liang, L.; Yang, Y.; Hu, Y.; Zhang, M. Social influence minimization based on context-aware multiple influences diffusion model. *Knowl.-Based Syst.* **2021**, *227*, 107233. [\[CrossRef\]](#)
16. Macal, C.M.; North, M.J. Tutorial on agent-based modeling and simulation. In Proceedings of the Winter Simulation Conference, Orlando, FL, USA, 4–7 December 2005.
17. Janssen, M.A. Agent-based modelling. *Model. Ecol. Econ.* **2005**, *155*, 172–181.
18. Luo, L.; Zhou, S.; Cai, W.; Low, M.Y.H.; Tian, F.; Wang, Y.; Xiao, X.; Chen, D. Agent-based human behavior modeling for crowd simulation. *Comput. Animat. Virtual Worlds* **2008**, *19*, 271–281. [\[CrossRef\]](#)
19. Kountouriotis, V.; Thomopoulos, S.C.; Papelis, Y. An agent-based crowd behaviour model for real time crowd behaviour simulation. *Pattern Recognit. Lett.* **2014**, *44*, 30–38. [\[CrossRef\]](#)
20. Liu, S.; Wen, H. Agent-based modelling of polarized news and opinion dynamics in social networks: A guidance-oriented approach. *J. Complex Netw.* **2024**, *12*, cnae028. [\[CrossRef\]](#)

21. Flache, A.; Mäs, M.; Feliciani, T.; Chattoe-Brown, E.; Deffuant, G.; Huet, S.; Lorenz, J. Models of social influence: Towards the next frontiers. *JASSS—J. Artif. Soc. Soc. Simul.* **2017**, *20*, 2. [CrossRef]
22. Axelrod, R. The dissemination of culture: A model with local convergence and global polarization. *J. Confl. Resolut.* **1997**, *41*, 203–226. [CrossRef]
23. Li, K.; Liang, H.; Kou, G.; Dong, Y. Opinion dynamics model based on the cognitive dissonance: An agent-based simulation. *Inf. Fusion* **2020**, *56*, 1–14. [CrossRef]
24. Madey, G.; Gao, Y.; Freeh, V.; Tynan, R.; Hoffman, C. Agent-Based Modeling and Simulation of Collaborative Social Networks. In *AMCIS 2003 Proceedings*; Article 237, 2003. Available online: <https://aisel.aisnet.org/amcis2003/237> (accessed on 11 November 2024).
25. El-Sayed, A.M.; Scarborough, P.; Seemann, L.; Galea, S. Social network analysis and agent-based modeling in social epidemiology. *Epidemiol. Perspect. Innov.* **2012**, *9*, 1–9. [CrossRef] [PubMed]
26. Gilbert, N. Agent-based social simulation: Dealing with complexity. *Complex Syst. Netw. Excell.* **2004**, *9*, 1–14.
27. Leombruni, R.; Richiardi, M. Why are economists sceptical about agent-based simulations? *Phys. A Stat. Mech. Its Appl.* **2005**, *355*, 103–109. [CrossRef]
28. Deguchi, H. *Economics as an Agent-Based Complex System: Toward Agent-Based Social Systems Sciences*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2011.
29. Rouchier, J. Agent-based simulation as a useful tool for the study of markets. In *Simulating Social Complexity: A Handbook*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 671–704.
30. Wang, L.; Ahn, K.; Kim, C.; Ha, C. Agent-based models in financial market studies. *J. Phys. Conf. Ser.* **2018**, *1039*, 012022. [CrossRef]
31. Mueller, M.; Pyka, A. Economic behaviour and agent-based modelling. In *Routledge Handbook of Behavioral Economics*; Routledge: London, UK, 2016; pp. 405–415.
32. Cabrera, E.; Taboada, M.; Iglesias, M.L.; Epelde, F.; Luque, E. Optimization of healthcare emergency departments by agent-based simulation. *Procedia Comput. Sci.* **2011**, *4*, 1880–1889. [CrossRef]
33. Barnes, S.; Golden, B.; Price, S. Applications of agent-based modeling and simulation to healthcare operations management. In *Handbook of Healthcare Operations Management: Methods and Applications*; Springer: Berlin/Heidelberg, Germany, 2013; pp. 45–74.
34. Perez, L.; Dragicevic, S. An agent-based approach for modeling dynamics of contagious disease spread. *Int. J. Health Geogr.* **2009**, *8*, 1–17. [CrossRef] [PubMed]
35. Silverman, B.G.; Hanrahan, N.; Bharathy, G.; Gordon, K.; Johnson, D. A systems approach to healthcare: Agent-based modeling, community mental health, and population well-being. *Artif. Intell. Med.* **2015**, *63*, 61–71. [CrossRef] [PubMed]
36. Beheshti, R.; Jalalpour, M.; Glass, T.A. Comparing methods of targeting obesity interventions in populations: An agent-based simulation. *SSM-Popul. Health* **2017**, *3*, 211–218. [CrossRef] [PubMed]
37. Elsenbroich, C.; Verhagen, H. Agent-Based Modelling and Public Health: An Introduction to the Proceedings of the 2023 Social Simulation Conference. In *Proceedings of the Conference of the European Social Simulation Association*, Glasgow, UK, 4–8 September 2023; pp. 1–6.
38. Hauser, M.D.; Chomsky, N.; Fitch, W.T. The faculty of language: What is it, who has it, and how did it evolve? *Science* **2002**, *298*, 1569–1579. [CrossRef] [PubMed]
39. Gao, C.; Lan, X.; Li, N.; Yuan, Y.; Ding, J.; Zhou, Z.; Xu, F.; Li, Y. Large language models empowered agent-based modeling and simulation: A survey and perspectives. *Humanit. Soc. Sci. Commun.* **2024**, *11*, 1–24. [CrossRef]
40. FTeam, 2022; Yoheinakajima. Babyagi. Available online: <https://github.com/yoheinakajima/babyagi> (accessed on 15 December 2024).
41. Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Hambro, E.; Zettlemoyer, L.; Cancedda, N.; Scialom, T. Toolformer: Language models can teach themselves to use tools. *Adv. Neural Inf. Process. Syst.* **2024**, *36*, 68539–68551.
42. Chen, Z.; Tang, J.; Chen, X.; Lin, Y. A survey on large language model based autonomous agents. *arXiv* **2023**, arXiv:2308.11432.
43. Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; et al. The rise and potential of large language model based agents: A survey. *arXiv* **2023**, arXiv:2309.07864.
44. Park, J.; Min, B.; Ma, X.; Kim, J. Choicemates: Supporting unfamiliar online decision-making with multi-agent conversational interactions. *arXiv* **2023**, arXiv:2310.01331.
45. Cipi, E.; Cico, B. Simulation of an agent based system behavior in a dynamic and unpredicted environment. *Simulation* **2011**, *1*, 172–176.
46. Fu, D.; Li, X.; Wen, L.; Dou, M.; Cai, P.; Shi, B.; Qiao, Y. Drive like a human: Rethinking autonomous driving with large language models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, Waikoloa, HI, USA, 3–8 January 2024; pp. 910–919.
47. Gao, C.; Lan, X.; Lu, Z.; Mao, J.; Piao, J.; Wang, H.; Jin, D.; Li, Y. S3: Social-network Simulation System with Large Language Model-Empowered Agents. *arXiv* **2023**, arXiv:2307.14984. [CrossRef]

48. Qian, C.; Cong, X.; Yang, C.; Chen, W.; Su, Y.; Xu, J.; Liu, Z.; Sun, M. Communicative agents for software development. *arXiv* **2023**, arXiv:2307.07924.
49. Wang, L.; Zhang, J.; Chen, X.; Lin, Y.; Song, R.; Zhao, W.X.; Wen, J.R. RecAgent: A Novel Simulation Paradigm for Recommender Systems. *arXiv* **2023**, arXiv:2306.02552.
50. Li, N.; Gao, C.; Li, M.; Li, Y.; Liao, Q. Econagent: Large language model-empowered agents for simulating macroeconomic activities. In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Bangkok, Thailand, 11–16 August 2024; Volume 1: Long Papers, pp. 15523–15536.
51. Gur, I.; Furuta, H.; Huang, A.; Safdari, M.; Matsuo, Y.; Eck, D.; Faust, A. A real-world webagent with planning, long context understanding, and program synthesis. *arXiv* **2023**, arXiv:2307.12856.
52. Xu, F.; Zhang, J.; Gao, C.; Feng, J.; Li, Y. Urban generative intelligence (ugi): A foundational platform for agents in embodied city environment. *arXiv* **2023**, arXiv:2312.11813.
53. Park, J.S.; O'Brien, J.; Cai, C.J.; Morris, M.R.; Liang, P.; Bernstein, M.S. Generative agents: Interactive simulacra of human behavior. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, Francisco, CA, USA, 29 October–1 November 2023; pp. 1–22.
54. Hua, W.; Fan, L.; Li, L.; Mei, K.; Ji, J.; Ge, Y.; Hemphill, L.; Zhang, Y. War and peace (waragent): Large language model-based multi-agent simulation of world wars. *arXiv* **2023**, arXiv:2311.17227.
55. Hu, Y.; Sherpa, G.; Zhang, L.; Li, W.; Bai, Q.; Wang, Y.; Wang, X. An LLM-enhanced Agent-based Simulation Tool for Information Propagation. In Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24, Jeju, Republic of Korea, 3–9 August 2024; pp. 8679–8682. [[CrossRef](#)]
56. Devlin, J. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv* **2018**, arXiv:1810.04805.
57. Holme, P.; Kim, B.J. Growing scale-free networks with tunable clustering. *Phys. Rev. E* **2002**, *65*, 026107. [[CrossRef](#)] [[PubMed](#)]
58. Muchnik, L.; Pei, S.; Parra, L.C.; Reis, S.D.; Andrade, J.S., Jr.; Havlin, S.; Makse, H.A. Origins of power-law degree distribution in the heterogeneity of human activity in social networks. *Sci. Rep.* **2013**, *3*, 1783. [[CrossRef](#)] [[PubMed](#)]
59. Macqueen, J. Some methods for classification and analysis of multivariate observations. In Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Oakland, CA, USA, 21 June–18 July 1965; Volume 1, pp. 281–297.
60. Leskovec, J.; Sosič, R. Snap: A general-purpose network analysis and graph-mining library. *ACM Trans. Intell. Syst. Technol. (TIST)* **2016**, *8*, 1–20. [[CrossRef](#)] [[PubMed](#)]
61. Dickel, S.; Neumann, M. Hermeneutics of social simulations*: On the interpretation of digitally generated narratives. In *An Interpretive Account to Agent-Based Social Simulation*; Routledge: London, UK, 2024; pp. 130–152.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.