Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)
Demonstrations Track

An LLM-enhanced Agent-based Simulation Tool for Information Propagation

Yuxuan Hu[1], **Gemju Sherpa**[1], **Lan Zhang**[2], **Weihua Li**[2], **Quan Bai**[1], **Yijun Wang**[1] and **Xiaodan Wang**[3]
[1]University of Tasmania, Hobart, Australia
[2]Auckland University of Technology, Auckland, New Zealand
[3]Yanbian University, Jilin, China
{yuxuan.hu, sherpag, quan.bai, ywang134}@utas.edu.au, {lan.zhang, weihua.li}@autuni.ac.nz, 0000007458@ybu.edu.cn

Abstract

Influence diffusion models are used for simulating information propagation in social networks. While most existing influence diffusion models are probabilistic, the emergence of Large Language Model (LLM) sheds light on the language-level inferences and interactions of user agents. This paper presents an LLM-enhanced Agent-based Influence Diffusion model (LAID), and a web-based visualization tool, LAIDSim, for simulating the information propagation in social networks.

1 Introduction

Information propagation describes how pieces of information, such as news, ideas, opinions, and advertisements, spread through social networks [Gomez-Rodriguez *et al.*, 2013; Wu *et al.*, 2019]. Simulating this phenomenon aids in understanding how information diffuses through these networks. Influence diffusion models [Kempe *et al.*, 2003], extensively used to simulate information propagation, follow a 'contagion' pattern based on the network structure [Gomez-Rodriguez *et al.*, 2013], where influence spreads from one node to another through connecting edges. Agent-based modeling (ABM) [Macal, 2016] has been applied to simulate influence diffusion. By leveraging the advantages of ABM [Li *et al.*, 2023; Li *et al.*, 2016; Hu *et al.*, 2019], the diffusion process is driven and captured by the behaviors of individual agents, allowing for a more accurate representation of the complex and dynamic nature of information diffusion.

Most existing influence diffusion models are probability-based, which presents several challenges. Firstly, the content of information plays a critical role in individuals' decision-making behaviors, but is often ignored by most probability-based models. Although some topic-aware influence diffusion models [Chen *et al.*, 2015] consider how tagged content impacts these probabilistic models, they are typically constrained to predefined topics [Chen *et al.*, 2023]. Secondly, probability-based models, including agent-based models, neglect individuals' ability to alter information [Wang *et al.*, 2023]. Specifically, the content of a piece of received information can be altered according to an individual's prior knowledge or interests.

In recent years, the emergence of Large Language Models (LLMs) has shed light on filling some of the above gaps [Gao *et al.*, 2023]. In this paper, we present an LLM-enhanced agent-based influence diffusion model, named **LAID**, along with a software simulation tool, named **LAIDSim**[1][2] for information propagation simulations. LAID is developed based on the conventional Independent Cascade model [Kempe *et al.*, 2003], while leveraging ABM [Li *et al.*, 2023] to capture the personal traits and individual behaviors of social network users. In LAID, users are represented as autonomous agents with distinctive profiles and complex social behaviors. Information content is considered a crucial factor influencing user agents' decision-making behaviors, specifically whether to be influenced by received messages. We particularly emphasize information alteration [Wang *et al.*, 2023] as a key feature in the propagation process, referring to the ability of user agents to modify the content of initially spread information during peer-to-peer propagation. LLMs are utilized to realize this feature, serving as the 'brain' of user agents. The LLM generates the content of an influence message based on the user agent's profile and the content of its received message, facilitating further propagation to its neighbors.

LAIDSim is implemented based on the ABM framework Melodie [Yu and Hou, 2023]. This tool enables end-users to simulate information propagation with self-defined scenarios, visualize the propagation process over time, and generate analytical results at a global level. To demonstrate the capabilities of LAIDSim, we designed and conducted four scenarios within two experiments. These experiments showcase the advantages of LAID in representing language-level behaviors of social network users and highlight the ability of LAIDSim to simulate information propagation across various application settings.

2 The LLM-enhanced Agent-based Influence Diffusion Model

2.1 Preliminaries

In LAID, a social network is modelled as is a directed graph $G = (V, E, M)$ with a set of user agents $V =$

[1]https://youtu.be/AaxDqK96RZE
[2]https://github.com/shaunahu/LAIDSim

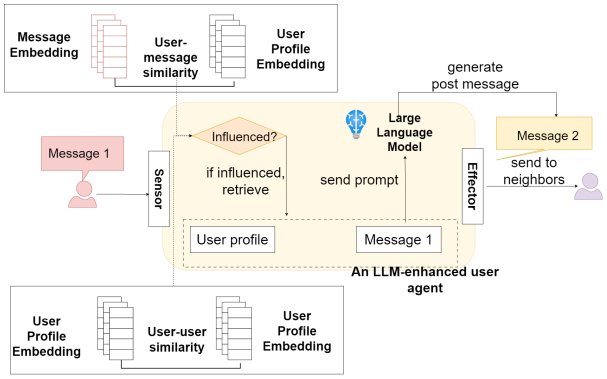

Figure 1: LLM-enhanced Agent-based Influence Diffusion

$\{v_1, v_2, ..., v_n\}$, a set of edges $E = \{e_{ij|1\leq i\leq n, 1\leq j\leq n, i\neq j}\}$, and a set of influence messages $M = \{m_1, m_2, ..., m_x\}$.

A use agent $v_i = (r_i, \Gamma_i)$ represents a social network user with a profile r_i. $\Gamma_i = \Gamma_i^{in} \cup \Gamma_i^{out}$ is the neighbor set of v_i which refers to a collection of user agents who are connected with v_i. Notably, $\Gamma_i^{in} = \{v_j | e_{ji} \in E\}$ is the in-neighbors of v_i, so that v_i receives influence messages from these neighbors. While $\Gamma_i^{out} = \{v_j | e_{ij} \in E\}$ is the out-neighbors of v_i, so that v_i sends influence messages to these users. An influence message $m_x = (c_x, v_i)$ is a piece of textual information spread from a user agent v_i to its neighbors, where c_x is the content of the message. Since language representation is not the focus of our study, we used BERT [Devlin *et al.*, 2018], a pre-trained model, to learn user profile embedding $\mathbf{r_i}$ of a user agent v_i based on its textual profile r_i, and content embedding $\mathbf{c_k}$ of an influence message m_k based on its content c_k.

2.2 Overall Framework

Figure 1 shows the overall framework of LAID with an example. The yellow square represents a LLM-enhanced user agent. In the social network, each user agent has its unique user profile, a repository to store received influence messages, and access to an LLM. The diffusion process starts from a user agent with an initial influence message, spreading to its neighbors through network topology. Once an agent receives an influence message from its neighbor, it decides whether to be influenced by the message or not based on the similarities between itself and the sender of the message, as well as the content of the message. If it was influenced, it retrieves its personal profile and the received message, then sends a request to an LLM, asking it to generate a new message for dissemination. The diffusion process continuous until there is no more user agents can be influenced.

2.3 Initialisation

We adopted GPT-3.5-turbo to generate user profiles with the following rules: (1) a user profile is a summary of a social network user within 50 words; (2) the ages of these users follow a normalised Gaussian distribution; (3) the genders of these users are evenly distributed, half male and half female. At the beginning of the influence diffusion, k user agents are randomly selected to spread an influence message m_s to their neighbors.

2.4 Influence Diffusion

Influence diffusion is developed based on the conventional IC model. In the IC model, each user agent has two possible states: active and inactive. An inactive user agent has a single chance to be activated by its active neighbors with an influence probability. The diffusion process stops if there were no more user agents could be influenced.

The influence probability p_{ij}^k is the probability for a user agent v_i to accept an influence message m_k from its neighbor v_j. This is calculated by the consideration of the content of the message m_k and the relationship between two user agents:

$$p_{ij}^k = \alpha * \text{sim}(m_k, v_i) + (1 - \alpha) * \text{sim}(v_i, v_j). \quad (1)$$

In Equation 1, $\text{sim}(m_k, v_i)$ is the cosine similarity between the content embedding $\mathbf{c}_k$ of m_k and the embedding $\mathbf{r}_i$ of v_i. It suggests how the content c_k of an influence message m_k matches a user agent v_i:

$$\text{sim}(m_k, v_i) = \frac{\mathbf{c}_k \cdot \mathbf{r}_i}{\|\mathbf{c}_k\| \, \|\mathbf{r}_i\|}. \quad (2)$$

Similarly, the similarity between two user agents v_i and v_j is measured by the cosine similarity between the user profile embeddings $\mathbf{r}_i$ and $\mathbf{r}_j$, where r_i and r_j are the profiles of v_i and v_j:

$$\text{sim}(v_i, v_j) = \frac{\mathbf{r}_i \cdot \mathbf{r}_j}{\|\mathbf{r}_i\| \, \|\mathbf{r}_j\|}. \quad (3)$$

In the influence diffusion process, we identify content influence as a critical factor in individuals accepting an influential message, denoted as $\text{sim}(m_k, v_i)$. The influence parameter α in the proposed diffusion model is a trade-off between these two factors. A higher α indicates that the user agent is more likely to be influenced by social influence rather than content, while a lower α suggests that the user agent is more sensitive to the content of the information but pays little attention to its neighbors' decisions.

Once a user agent is influenced, it accepts the spreading influence message m_k and further propagates to its neighbors. Specifically, LLM model is adopted for information alteration. That is, a user agent can alter the content of an existing influence message m_k that it received from any of its in-neighbors $v_j \in \Gamma_i^{in}$, adding or changing words based on its profiles, and further diffuses the altered influence message m_k' to its out-neighbors Γ_i^{out}.

2.5 Information Alteration

GPT-3.5 Turbo is used in this stage for user agents to generate altered influence messages based on their profiles and the influence messages they have accepted. To evaluate the changes between two influence messages, we adopted alteration degree [Wang *et al.*, 2023] to measure the extent of information alteration by a user agent.

ID	Content	Application Setting
s_1	Meet your local election candidates: [Candidate 1] focuses on education reform and community safety. [Candidate 2] advocates for sustainable development and environmental initiatives. [Candidate 3] prioritizes small business support and infrastructure improvements. Get informed, attend candidate forums, and choose who aligns with your vision for our community's future. #LocalElections	Political election
s_2	Introducing our latest product: The SmartHome Hub! Seamlessly control all your devices, from lights to thermostats, in one central hub. Enjoy energy savings, convenience, and security. Compatible with popular smart home platforms. Upgrade your home today and experience the future of intelligent living. Available now at authorized retailers.	Viral marketing
s_3	Introducing a novel concept: Shared Skill Spaces. A communal area where diverse talents intersect, fostering collaboration and innovation. Unleash creativity, maximize collective expertise, and elevate the potential of our community. Explore Shared Skill Spaces today!	Spread new idea
s_4	How are you today?	Daily conversation

Table 1: Scenarios with different message contents and application settings

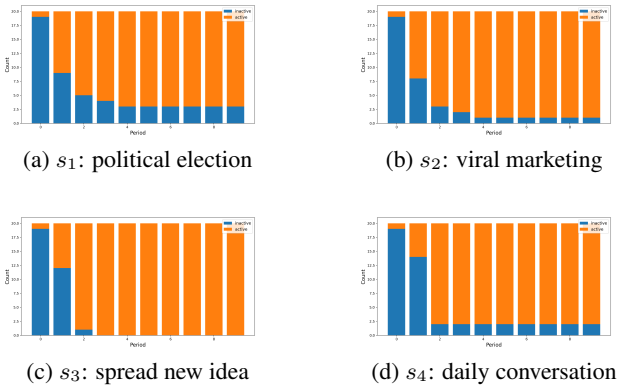

(a) s_1: political election (b) s_2: viral marketing

(c) s_3: spread new idea (d) s_4: daily conversation

Figure 2: Influence coverage of different scenarios

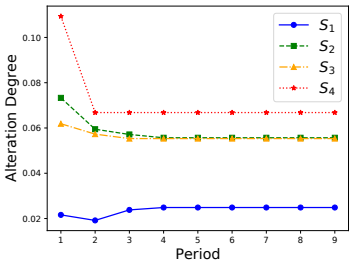

Figure 3: Information alteration of different scenarios

3 Experiments

3.1 Experimental Settings

Parameters: We consider seed size k and influence parameter α as two parameters of our proposed model that can be adjusted from the front-end. The seed size k represents the number of initially active user agents. A higher k suggests a quicker diffusion but also entails higher costs. As explained in Subsection 2.4, the influence parameter α potentially impacts the influence probability by controlling the trade-off between social influence and content influence.

Scenarios: We designed four scenarios with varying initial influence message content and application settings for the proposed model to simulate. The initial influence message is also generated by the GPT-3.5-Turbo model. Table 1 lists these scenarios.

Initialization Setup: At the beginning of each scenario simulation, a synthetic network based on Erdos-Renyi model [Erdős *et al.*, 2012] is randomly constructed with a size of 20 nodes, and a probability of 0.1 for these nodes to connect to each other.

3.2 Experiment 1: Influence Coverage Analysis

This experiment analyses the cumulative number of user agents in two states: active and inactive over a particular time period, i.e., 10 time steps. Figure 2 shows the simulation results. The slight differences of the numbers of active and inactive agents across the four scenarios suggest that user agents possess the capability to respond to information with different contents. In all four scenarios, the diffusion process adheres to a contagion pattern, indicating LAID's capacity to effectively simulate information propagation.

3.3 Experiment 2: Information Alteration Analysis

Figure 3 shows the simulation results of four scenarios monitored from time step 1. It can be seen that all four scenarios have been slightly altered at the first time step compared to the initial influence message. S_4 shows the most significant alteration. This is because the initial influence message in S_4 is an 'open' question. The alteration degree scenarios s_2 to s_4 decreases, implying that user agents generate information with a similar pattern to their received influence messages. The degree of alteration converges after a few time steps, indicating that the diffusion process stops, and no more user agents can be influenced to generate new influence messages.

4 Conclusion

In this paper, we presented an LLM-enhanced agent-based simulation tool, LAIDSim, for simulating information propagation in social networks, based on our proposed diffusion model LAID. We showcased the interpretability of the proposed model across four scenarios by demonstrating language-level interaction and inference of user agents. In the future, we plan to extend our work by improving the user agent model with more comprehensive user behaviors.

Contribution Statement

The first two authors contributed equally to this work and should be considered as joint first authors.

References

[Chen *et al.*, 2015] Shuo Chen, Ju Fan, Guoliang Li, Jianhua Feng, Kian-lee Tan, and Jinhui Tang. Online topic-aware influence maximization. *Proceedings of the VLDB Endowment*, 8(6):666–677, 2015.

[Chen *et al.*, 2023] Hsi-Wen Chen, De-Nian Yang, Wang-Chien Lee, Philip S Yu, and Ming-Syan Chen. Cminet: a graph learning framework for content-aware multi-channel influence diffusion. In *Proceedings of the ACM Web Conference 2023*, pages 545–555, 2023.

[Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[Erdős *et al.*, 2012] László Erdős, Antti Knowles, Horng-Tzer Yau, and Jun Yin. Spectral statistics of erdős-rényi graphs ii: Eigenvalue spacing and the extreme eigenvalues. *Communications in Mathematical Physics*, 314(3):587–640, 2012.

[Gao *et al.*, 2023] Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. S ^{3}: Social-network simulation system with large language model-empowered agents. *arXiv preprint arXiv:2307.14984*, 2023.

[Gomez-Rodriguez *et al.*, 2013] Manuel Gomez-Rodriguez, Jure Leskovec, and Bernhard Schölkopf. Modeling information propagation with survival theory. In *International conference on machine learning*, pages 666–674. PMLR, 2013.

[Hu *et al.*, 2019] Yuxuan Hu, Quan Bai, and Weihua Li. Context-aware influence diffusion in online social networks. In *Knowledge Management and Acquisition for Intelligent Systems: 16th Pacific Rim Knowledge Acquisition Workshop, PKAW 2019, Cuvu, Fiji, August 26–27, 2019, Proceedings 16*, pages 153–162. Springer, 2019.

[Kempe *et al.*, 2003] David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003.

[Li *et al.*, 2016] Weihua Li, Quan Bai, and Minjie Zhang. Agent-based influence propagation in social networks. In *2016 IEEE International Conference on Agents (ICA)*, pages 51–56, 2016.

[Li *et al.*, 2023] Weihua Li, Yuxuan Hu, Chenting Jiang, Shiqing Wu, Quan Bai, and Edmund Lai. Abem: An adaptive agent-based evolutionary approach for influence maximization in dynamic social networks. *Applied Soft Computing*, 136:110062, 2023.

[Macal, 2016] Charles M Macal. Everything you need to know about agent-based modelling and simulation. *Journal of Simulation*, 10:144–156, 2016.

[Wang *et al.*, 2023] Guan Wang, Weihua Li, Quan Bai, and Edmund MK Lai. Maximizing social influence with minimum information alteration. *IEEE Transactions on Emerging Topics in Computing*, 2023.

[Wu *et al.*, 2019] Le Wu, Peijie Sun, Yanjie Fu, Richang Hong, Xiting Wang, and Meng Wang. A neural influence diffusion model for social recommendation. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 235–244, 2019.

[Yu and Hou, 2023] Songmin Yu and Zhanyi Hou. Melodie: Agent-based modeling in python. *Journal of Open Source Software*, 8(83):5100, 2023.