Modelling Information Spread using Generative Agents

Petra Vidnerová[1,2][0000−0003−3879−3459], Gabriela Kadlecová[1,2,3][0000−0002−4780−0633], Roman Neruda[1,2][0000−0003−2364−5357], and Josef Šlerka[1][0000−0003−3564−1767]

[1] Charles University, Faculty of Arts
[2] The Czech Academy of Sciences, Institute of Computer Science
[3] Charles University, Faculty of Mathematics and Physics

Abstract. In this paper, we introduce a novel multiagent model designed to facilitate communication within small groups of agents. Leveraging the capabilities of large language models (LLMs), our approach enables agents to generate and interpret communications in a natural language, thereby enhancing their collaborative and decision-making processes. The agents are connected in a graph that represents their social network. They are equipped with a memory that is represented by a list of facts in a natural language and an ability to communicate with each other. To verify the feasibility and effectiveness of our model, we conducted several experiments. Since we are generally interested in studying an information transmission dynamics, our experiments are focused on an information transmission. The main goal of the preliminary experiments is to verify if this approach is suitable for analysing the rates of transmission of a particular piece of information throughout the agent community. Our experiments are powered by the publicly available state-of-the-art Mixtral-8x7B-Instruct-v0.1 LLM model.

Keywords: information spread modelling, agent-based models, generative agents, large language models

1 Introduction

The field of artificial intelligence (AI) has witnessed remarkable advancements over the past few decades. Among these advancements, large language models (LLMs) have emerged as pivotal components, driving innovation across various domains. These technologies leverage vast amounts of data and sophisticated algorithms to understand, generate, and interact with human language and behaviour in increasingly complex ways.

Agent-based models (ABM) on the other hand, extend the capabilities of AI by incorporating elements of decision-making, interaction, and autonomous behaviour. These models are designed to perform tasks, make decisions, and adapt to new information in dynamic environments. They can range from simple rule-based systems to complex, adaptive agents capable of learning from their interactions and improving their performance over time. Agent models are employed in various fields, from autonomous vehicles and robotics to virtual assistants and personalized recommendation systems. The overview of two decades of classical ABM in social sciences is well covered in survey [1].

The use of multi-agent systems powered by large language models (LLMs) presents a novel approach to studying communication dynamics in small groups. These systems simulate interactions among multiple autonomous agents, each driven by advanced LLMs capable of generating and understanding human-like text. By emulating real-world communication patterns, these agents can engage in complex dialogues, negotiate, collaborate, and resolve conflicts, providing rich data for analysis. Researchers can manipulate various parameters within these systems, such as communication styles, group roles, and external stressors, to observe their effects on group dynamics and outcomes. This approach allows for a controlled yet flexible experimental environment, offering insights into the mechanisms of effective communication, decision-making processes, and the impact of individual behaviours on group performance. Ultimately, multi-agent systems powered by LLMs not only enhance our understanding of group communication but also pave the way for developing more sophisticated collaborative technologies and strategies in real-world settings.

This paper represents the initial step towards modelling communication and information dissemination in small groups of agents.

The paper is organized as follows: the next section summarizes the related work, the Section 3 introduces the agent model designed by ourselves, the Section 4 presents the results of our experiments, and finally the last section is devoted to conclusion.

2 Related Work

The seminal work of Park et al. [6] presents a successful application of LLM as the control mechanism of agents communicating in natural language. Their environment mimics a Sims-like world with 25 agents simulating limited but believable human behaviour, communicating, planning their activities, and coordinating their behaviour. The work identifies three key components necessary for the simulation to be believable: observation, planning, and reflection of agents.

A similar, albeit simpler, showcase of LLM agents is presented in [4] where small agent groups of 2 to 6 agents participate in tasks of negotiating or playing a simple mystery game. First proof-of-concept tools for the design and deployment of LLM-based agent systems are appearing, such as the Auto-GPT framework for online decision-making tasks [8].

Several examples of LLM application in the context of generative agents have recently appeared. They include a study of a generative agent-based system for the design and execution of scientific experiments [2], or a similar approach proposing LLM agent-based simulations in medical clinic settings [5].

3 Agent Communication Model

Our Agent Communication Model works with a graph that represents a social network of present agents. Each node corresponds to one agent, an edge represents a possibility of contact. By contact, we mean a single chat between the corresponding two agents.

A chat is a sequence of natural language statements (a sentence, group of sentences), the statements are generated via LLMs. Each agent has several attributes (name, sex, age) and a description. They are all a part of a setup. A description is a short summary in a natural language.

An example description: *Isabella Rodriguez is a cafe owner of Hobbs Cafe who loves to make people feel welcome. She is always looking for ways to make the cafe a place where people can come to relax and enjoy themselves. Isabella Rodriguez is planning on having a Valentine's Day party at Hobbs Cafe with her customers on February 14th, 2023 at 5 pm. She is gathering party material and is telling everyone to join the party at Hobbs Cafe on February 14th, 2023, from 5 pm to 7 pm.* (this description is used in our experiments and is inspired by the agent from the project [6]).

In addition, agents have a memory. A memory is a list of facts in natural language. In a dialogue, after each step, an LLM is asked to summarize it into fact and this fact is saved into the agent's memory. In addition, agents have a list of information about agents that they already met.

Memory is a key part of an agent, it enables it to store gained knowledge and continue the communication on a particular topic the next time it meets the partner or use the gained knowledge in a conversation with other agents.

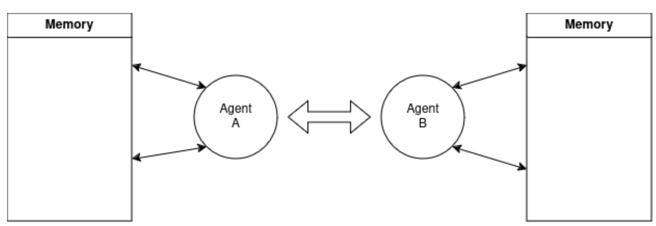

Fig. 1. The chat between two agents. The memory is used the generate the chat and the new facts learnt during the chat are stored back in the memory.

A contact (a chat) between two agents is generated as follows. A task-dependent scheduler is responsible for generating pairs of agents for a chat. Once a pair of agents is chosen, the LLM gets information (name, sex, age, description, memory) about the starting agents and whether it has already met the second agent (the communication starts differently if the agents do not know each other, i.e. it is their first meeting, or if the agents already met). The first agent is asked to generate the first part of a dialogue. Then the second agent is primed with corresponding information, and the agents alternate in replies until one of them "says goodbye" or the maximum number of replies is met. The replies are generated in the same way, the LLM gets information about the agent, its memory and the dialogue so far and it is asked to generate the next piece of dialogue. To decide whether the agent wants to end the dialogue ("says goodbye") the average cosine similarity with several good-bye-phrases is calculated.

A simulation of the whole system runs in iterations. Typically, during one iteration, all edges are activated once, i.e. all couples of agents connected by an edge have one chat with each other, the order of edges depends on the task scheduler.

4 Experiments

The goal of our preliminary experiments is to verify the feasibility of the proposed model. We study if information spreads among the agents, how it depends on the properties of agents, and if the results are believable.

As an LLM we used the `Mixtral-8x7B-Instruct-v0.1` [3]. The code of our experiments is publicly available on GitHub [7].

4.1 Experiment I

The setup of the first experiment is depicted in Fig. 2. There are three agents, Isabella, Maria and Klaus, whose characteristics are taken from the project [6].

Isabella is 34 years old, she is an owner of Hobbs Cafe and is planning a Valentine's Day party. Maria, 21 years old, is a physics student and game streamer. Klaus, 20 years old, is a sociology student, and he is currently writing a research paper on gentrification.

In each iteration of the simulation, first, the edge between Isabella and Maria is activated, and then the edge between Maria and Klaus. We are interested if the information about the Valentine's Day party reaches Klaus (since he never speaks directly with Isabella). Therefore, after each iteration, we query each agent, i.e. the LLM gets information about the agent and listing of its memory and question whether the agent knows about the party.

The simulation is repeated ten times and answers "yes" are interpreted as 1, while answers "no" as zero. Answers are averaged over the 10 runs. The results are displayed in the Fig. 3.

We can see that Isabella speaks about her party in all 10 simulation runs in her first meeting with Maria since Maria knows about the party in all cases from the first iteration. On the other hand, Maria mentions it immediately only in a portion of runs, the probability that the information reaches Klaus grows with time.

4.2 Experiment II

In the second experiment, we added a new agent called Julia to the graph (see Fig. 4),. It represents a person, who is a sociable florist and loves partying. We

Fig. 2. Experiment I: a graph setup.

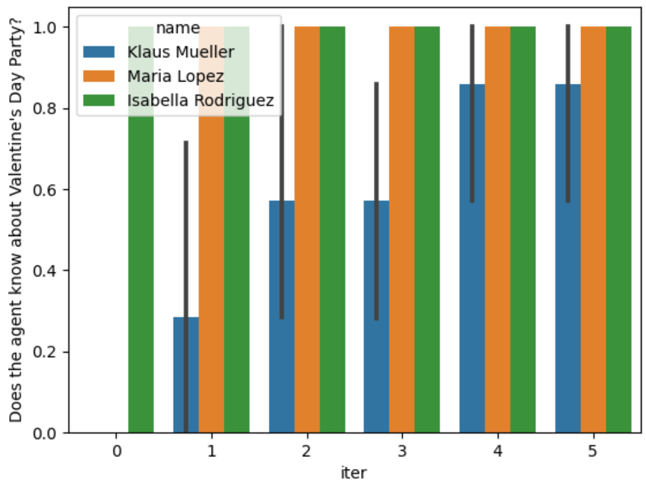

Fig. 3. Experiment I: the barplot shows the relative portion of runs in which the agent knows about Isabella's party (mean from 10 runs with 95% CI). Each iteration contains all (two) chats from the communication graph.

expected that she would tend to chat about parties more likely than Klaus and therefore she is more likely to get the information about Isabella's party.

In this case, all three edges in the graph are activated once in each iteration, the edge between Isabella and Maria is always the first; the other two edges follow in a random order.

The simulation was again repeated 10 times and the results are displayed in Fig. 5. We can see that the information about the party reached Julia earlier than Klaus.

4.3 Experiment III

The last experiment is a modification of the second one, we added one edge to the graph, connecting Julia and Klaus (Fig. 6).

Again, the edge between Maria and Isabella is activated as the first one, and then the rest of the edges is activated in a random order.

The results are displayed in Fig. 7 and are similar to those in experiment II, but the information spreads slightly faster.

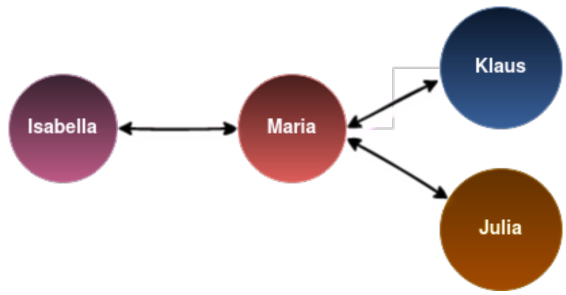

Fig. 4. Experiment II: a graph setup.

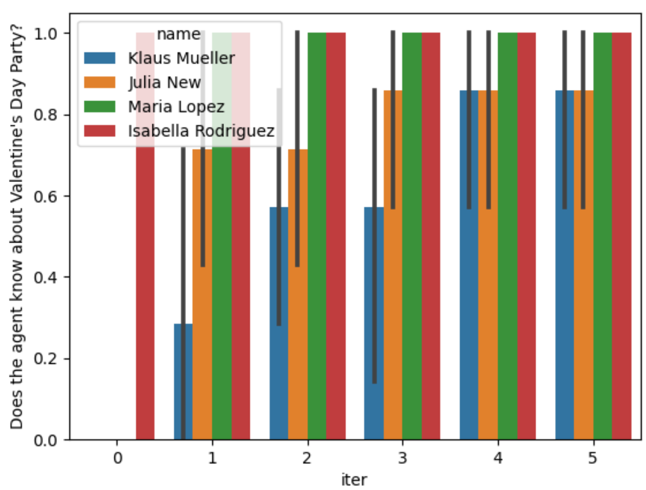

Fig. 5. Experiment II: the barplot shows the portion of runs in which the agent knows about Isabella's party (mean from 10 runs with 95% CI). Each iteration contains all (three) chats from the communication graph..

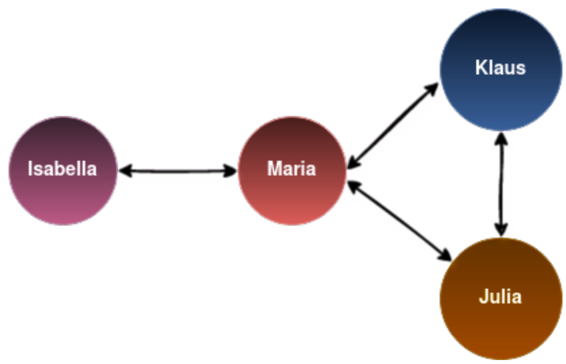

Fig. 6. Experiment III: a graph setup

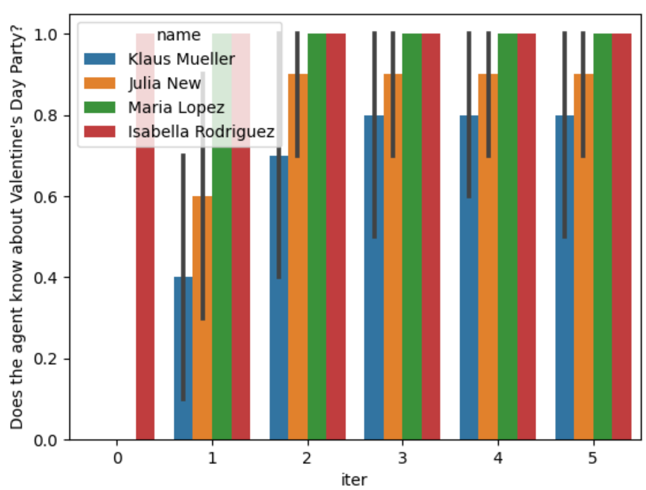

Fig. 7. Experiment III: the barplot shows the portion of runs in which the agent knows about Isabella's party (mean from 10 runs with 95% CI). Each iteration contains all (four) chats from the communication graph..

4.4 Examples of dialogues

Fig. 8 presents three examples of dialogues that were generated during our experiments. We see the difference in the start of a chat when the agents meet for the first time and when they already know each other. Also example of a situation where information about Isabella's party is passed to Klaus is presented.

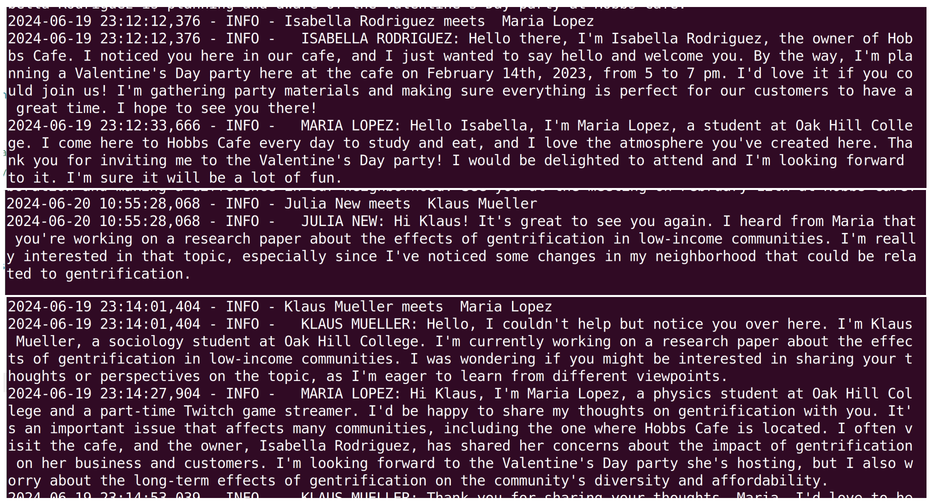

```
2024-06-19 23:12:12,376 - INFO - Isabella Rodriguez meets  Maria Lopez
2024-06-19 23:12:12,376 - INFO -   ISABELLA RODRIGUEZ: Hello there, I'm Isabella Rodriguez, the owner of Hob
bs Cafe. I noticed you here in our cafe, and I just wanted to say hello and welcome you. By the way, I'm pla
nning a Valentine's Day party here at the cafe on February 14th, 2023, from 5 to 7 pm. I'd love it if you co
uld join us! I'm gathering party materials and making sure everything is perfect for our customers to have a
 great time. I hope to see you there!
2024-06-19 23:12:33,666 - INFO -   MARIA LOPEZ: Hello Isabella, I'm Maria Lopez, a student at Oak Hill Colle
ge. I come here to Hobbs Cafe every day to study and eat, and I love the atmosphere you've created here. Tha
nk you for inviting me to the Valentine's Day party! I would be delighted to attend and I'm looking forward
to it. I'm sure it will be a lot of fun.
```

```
2024-06-20 10:55:28,068 - INFO - Julia New meets  Klaus Mueller
2024-06-20 10:55:28,068 - INFO -   JULIA NEW: Hi Klaus! It's great to see you again. I heard from Maria that
 you're working on a research paper about the effects of gentrification in low-income communities. I'm reall
y interested in that topic, especially since I've noticed some changes in my neighborhood that could be rela
ted to gentrification.
```

```
2024-06-19 23:14:01,404 - INFO - Klaus Mueller meets  Maria Lopez
2024-06-19 23:14:01,404 - INFO -   KLAUS MUELLER: Hello, I couldn't help but notice you over here. I'm Klaus
 Mueller, a sociology student at Oak Hill College. I'm currently working on a research paper about the effec
ts of gentrification in low-income communities. I was wondering if you might be interested in sharing your t
houghts or perspectives on the topic, as I'm eager to learn from different viewpoints.
2024-06-19 23:14:27,904 - INFO -   MARIA LOPEZ: Hi Klaus, I'm Maria Lopez, a physics student at Oak Hill Col
lege and a part-time Twitch game streamer. I'd be happy to share my thoughts on gentrification with you. It'
s an important issue that affects many communities, including the one where Hobbs Cafe is located. I often v
isit the cafe, and the owner, Isabella Rodriguez, has shared her concerns about the impact of gentrification
 on her business and customers. I'm looking forward to the Valentine's Day party she's hosting, but I also w
orry about the long-term effects of gentrification on the community's diversity and affordability.
```

Fig. 8. Examples of generated dialogues between agents. (top) Isabella meets Maria who she has not met before. Therefore she introduces herself. (middle) Julia meets Klaus, she already has a reference to Klaus in her memory, so she starts the conversation with "It's great to see you again." (bottom) Example of passing the information about Isabella's party.

5 Conclusion

Our research introduced a novel multi-agent network model, driven entirely by large language models (LLMs), to simulate agent communication. Through our preliminary experimental validation, we have demonstrated the model's potential. The results are believable in terms of the parameters of agents and the relative speed of information spread.

Looking ahead, this model will serve as a foundational tool in our continued exploration of communication patterns and information transmission. In our future work, we will create larger agent communities with up to a dozen of agents, and study communication patterns supported by real data from experiments with human respondents. Further improvements of the model include more diverse communication graphs, different chat schedulers, enriching the agents about the planning component, and possibly a better reflection of agents.

Future experiments will also include comparisons of different LLMs.

Acknowledgments. This work was supported by the European Regional Development Fund project "Beyond Security: Role of Conflict in Resilience-Building" (reg. no.: CZ.02.01.01/00/22_008/0004595). This research was partially supported by SVV project number 260 698.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bianchi, F., Squazzoni, F.: Agent-based models in sociology. WIREs Comput. Stat. **7**(4), 284–306 (jul 2015). https://doi.org/10.1002/wics.1356, https://doi.org/10.1002/wics.1356
2. Boiko, D.A., MacKnight, R., Gomes, G.: Emergent autonomous scientific research capabilities of large language models (2023), https://arxiv.org/abs/2304.05332
3. Jiang, A.Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D.S., de las Casas, D., Hanna, E.B., Bressand, F., Lengyel, G., Bour, G., Lample, G., Lavaud, L.R., Saulnier, L., Lachaux, M.A., Stock, P., Subramanian, S., Yang, S., Antoniak, S., Scao, T.L., Gervet, T., Lavril, T., Wang, T., Lacroix, T., Sayed, W.E.: Mixtral of experts (2024), https://arxiv.org/abs/2401.04088
4. Junprung, E.: Exploring the intersection of large language models and agent-based modeling via prompt engineering (2023), https://arxiv.org/abs/2308.07411
5. Mehandru, N., Miao, B.Y., Almaraz, E.R., Sushil, M., Butte, A.J., Alaa, A.: Evaluating large language models as agents in the clinic. npj Digital Medicine **7**(1) (Apr 2024). https://doi.org/10.1038/s41746-024-01083-y, http://dx.doi.org/10.1038/s41746-024-01083-y
6. Park, J.S., O'Brien, J.C., Cai, C.J., Morris, M.R., Liang, P., Bernstein, M.S.: Generative agents: Interactive simulacra of human behavior (2023), https://arxiv.org/abs/2304.03442
7. Vidnerová, P.: Chat simulation software, https://github.com/PetraVidnerova/simchat
8. Yang, H., Yue, S., He, Y.: Auto-gpt for online decision making: Benchmarks and additional opinions (2023), https://arxiv.org/abs/2306.02224