

An Implementation of Werewolf Agent That does not Truly Trust LLMs

Takehiro Sato[†], Shintaro Ozaki[‡], Daisaku Yokoyama[†]
[†]Meiji University [‡]Nara Institute of Science and Technology
{ce245022, dyokoyama}@meiji.ac.jp
ozaki.shintaro.ou6@naist.ac.jp

Abstract

Werewolf is an incomplete information game, which has several challenges when creating a computer agent as a player given the lack of understanding of the situation and individuality of utterance (e.g., computer agents are not capable of characterful utterance or situational lying). We propose a werewolf agent that solves some of those difficulties by combining a Large Language Model (LLM) and a rule-based algorithm. In particular, our agent uses a rule-based algorithm to select an output either from an LLM or a template prepared beforehand based on the results of analyzing conversation history using an LLM. It allows the agent to refute in specific situations, identify when to end the conversation, and behave with persona. This approach mitigated conversational inconsistencies and facilitated logical utterance as a result. We also conducted a qualitative evaluation, which resulted in our agent being perceived as more human-like compared to an unmodified LLM. The agent is freely available for contributing to advance the research in the field of Werewolf game¹.

1 Introduction

Werewolf (Ri et al., 2022) game is one of the popular imperfect information multi-player games, in which the players are separated into two sides, villagers and werewolves, and they hide their roles and try to make advantageous consensus among other players through natural language conversation. Playing Werewolf game requires high-level intelligence skills such as reasoning, cooperating, and lying. It is in particular challenging for a computer to play the game in terms of both Game Informatics and Natural Language Processing, and has been widely studied for years (Kano et al., 2023).

The game comprises at least three serious challenges as follows:

¹<https://github.com/meiji-yokoyama-lab/AIWolfDial2024>

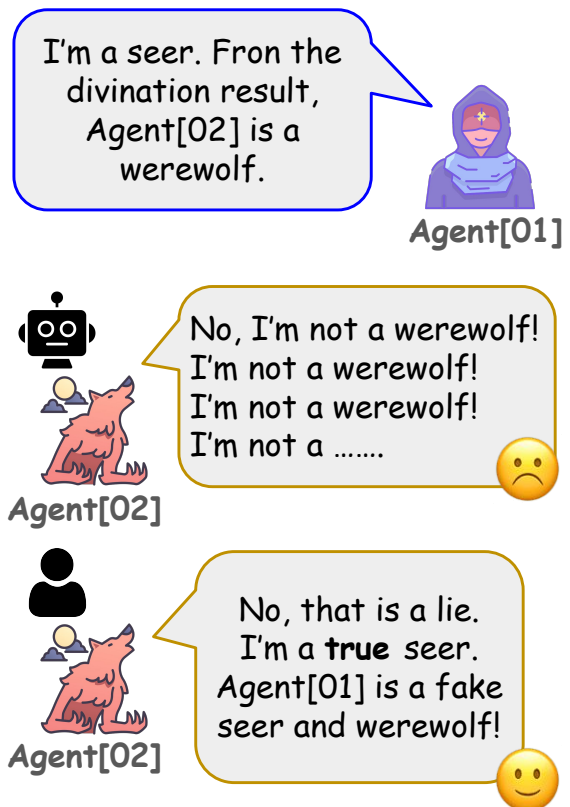


Figure 1: One example of problems with playing Werewolf game using LLMs. Humans can tell a logical lie naturally, but an LLM can only deny it.

1. The current situation is only shown in the players' conversation. The game system shows very limited information, such as who is alive. There exists other necessary information to play the game reasonably, but they are inferred from the conversation history: who invites doubting who, who already decided to do something, who might change their mind, and so on.
2. The player should make a tactical talk to pursue a specific purpose. For example, when a

player is doubted, the player should make a reasonable refute, not merely insist on their idea Figure (1). In addition, a player needs to know when to end the conversation to take their advantage, especially when all other players seem to doubt the other person.

3. The player should have an attractive individuality. It is not required to achieve a win in the game, though it is quite important to make it a fun game, which involves many aspects such as talking styles, intelligent decision-making, and role-playing (Callison-Burch et al., 2022).

Many Large Language Models (LLMs) (OpenAI, 2022; Anil et al., 2023; Achiam et al., 2023; Touvron et al., 2023a; Meta, 2023, 2024; Google, 2024; Team et al., 2023; OpenAI, 2023; Touvron et al., 2023a,b) with very high generalization ability have been published, and of course several models have already been applied to the werewolf agent (Xu et al., 2023; Wu et al., 2024). However, simply leveraging LLMs cannot solve those difficulties when you implement werewolf agent with LLMs. Among several challenges in developing a werewolf agent, we focus on the following aspects in this work: 1) The agent should refute under a certain critical situation; 2) The agent should cut off the conversation when the discussion is identified as meaningless; 3) The agent has distinguishable talk style under a consistent personality to make the game fun.

Our approaches are summarized as follows.

Rule-based algorithm with LLMs We combine an LLM with a rule-based algorithm. The LLM retrieves the conversation history in the game and produces an output. The rule-based algorithm decides whether that output is appropriate or not, given the game situation. If the talk is inappropriate, the rule-based algorithm uses a predefined template utterance. As a result, the rule-based algorithm can lie in critical situations and terminate conversations when there is no longer a need to continue.

Extracting Game Information To understand the current situation from the conversation history, we also utilize an additional LLM to extract the game-related information. We choose several basic but critical game concepts, such as voting decisions and divination results. The LLM examines conversation history and generates talks containing this information in a fixed format. The information



Figure 2: The list of five-person werewolf roles.

is also used by the rule-based algorithm to make decisions.

Style Transformation We decide to use an LLM that is pre-trained from a pile of general documents. In addition, we use prompts to control them without modifying or fine-tuning the model and give the agent distinguishable personalities using prompts.

Our preliminary implementation solved those tasks. This approach led our model to mitigate conversational inconsistencies and facilitated logical utterance as a result. Moreover, we also evaluated the agent by conducting a qualitative evaluation. As a result, compared to an unmodified LLM, incorporating rule-based approaches made it appear as though the agent understood the conversation, and inserting a persona enabled it to engage in more natural conversations. The source code is openly available in the hope that future research on werewolf agents will grow.

2 Related Work

Research on Werewolf game has a long history, which can be traced back to a study on Mafia, a game similar to Werewolf game, to mathematically analyze (Braverman et al., 2008; Migdał,

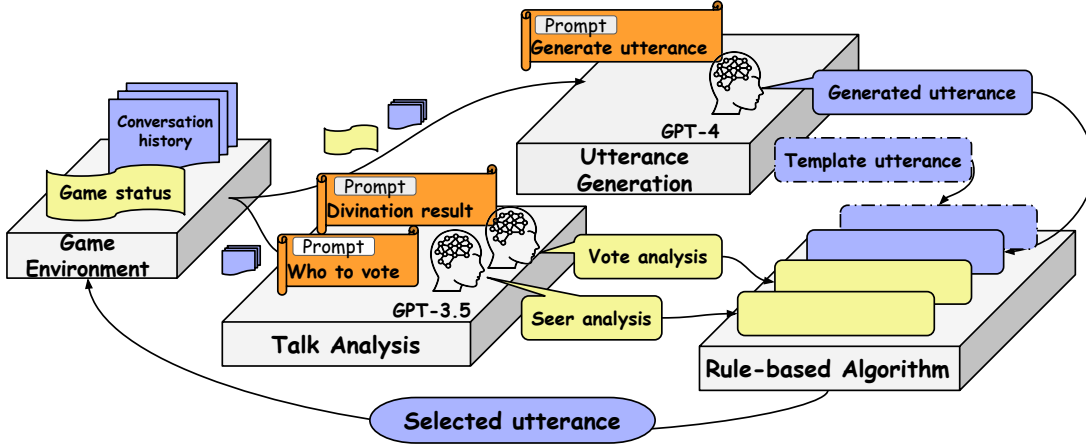


Figure 3: System overview. Our system comprises three modules, utterance generation, talk analysis, and rule-based algorithm. We described utterance generation in Section (4.1), talk analysis in Section (4.3), rule-based algorithm in Section (4.4), and required game status in Appendix (A.1).

2013). Some studies analyze the logs of Werewolf game (Nagayama et al., 2019; Fukui et al., 2017), or discuss methods to make werewolf agents stronger (Nakamura et al., 2016; Wang and Kaneko, 2018). Recently, with the development of LLMs, those models have been already explored for werewolf agents (Xu et al., 2023; Wu et al., 2024). However, these LLM-based agents have difficulties in handling werewolf-specific features, such as doubting, lying, and detecting the lie. In addition, these models produce outputs without a persona. While LLM-only approaches are dominant, there is growing interest in hybridizing rule-based methods with LLMs in other fields. In data analytics or business, a commonly used approach for extracting information from structured data involves both LLMs and rule-based methods (Huang, 2024; Vertsel and Rumiantsev, 2024). We aim to apply this methodology to the werewolf agent, leveraging the strengths of both approaches. This hybrid approach could lead to more robust and adaptable werewolf agents.

3 Five-person Werewolf Game

We selected a simple setting played by five players for the Werewolf game. In this game setting, villager, seer, possessed, and werewolf are used. As far as each role, “Villager” has no special abilities, “Seer” can know one player’s species each night with the ability to divine, “Possessed” has no special abilities, and is judged to be human by a divination result. However, the possessed acts so that the werewolf can win. “Werewolf” can select one player each night to attack and remove them from

the game. Since only a few players are involved, the game tends to determine the outcome on Day 1. Thus, we focused on the conversation phase on Day 1. Only the seer can act on the night of Day 0, and Day 1 begins with the seer having information about one player’s species. It is a recommended tactic for the seer to disclose the information obtained by revealing the seer. Revealing one’s position is referred to as CO (Coming Out).

4 System Design

Figure (3) shows an overall diagram of our system. The utterance generation module creates a prompt from the game status and conversation history sent from the server. The prompt is input into the LLM to obtain an utterance that is naturally connected to the conversation history. The talk analysis module creates a prompt to analyze the conversation history, and the LLM outputs the situation information related to voting and divination results considered from the conversation history. The rule-based algorithm is used to select either template utterance or LLM output, depending on the situation obtained by talk analysis. The selected utterance is sent to the server as the next utterance, and another agent’s turn begins.

4.1 Utterance Generation

We make a prompt for an LLM to generate a continuous utterance of the conversation history in a game. The prompt is structured by giving general rules of Werewolf game, some tips to play, conversation history, and current game state. The current

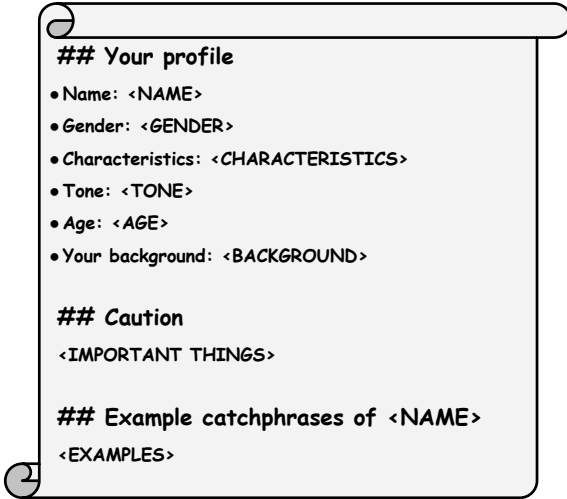


Figure 4: An example of prompts regarding style transformation. <CAPITAL LETTER> is the variable.

Character Name	Gender	Age
Princess	Female	Young
Kansai	Male	Young
Hiroshima dialect	Male	Elderly
Anya	Female	Child
Zundamon	Female	AI (Virtual)

Table 1: Overview of character information: We prepared five characters and aimed to give them character by specifying their age, name, first person, and gender.

game state, such as the player’s ID, role, and other live/dead players, is derived from the game status sent from the server. Details of the derived game state information are shown in Appendix (A.1). This module can go along with the conversation and continue a game of Werewolf.

4.2 Persona

Werewolf is not only a game of win or lose, but also a party game, so it is important to add character to the conversation. In addition, it is difficult to distinguish five players’ talks when all players have the same talk style. In order to give our model characterization, we introduced prompts that would perform style transformations. We have prepared five character prompts, which are listed in Table (1). We chose Princess, the Kansai dialect, the Hiroshima dialect, an anime character called Anya, and Zundamon, which is a popular avatar in Japan. The specific style conversion prompt example is

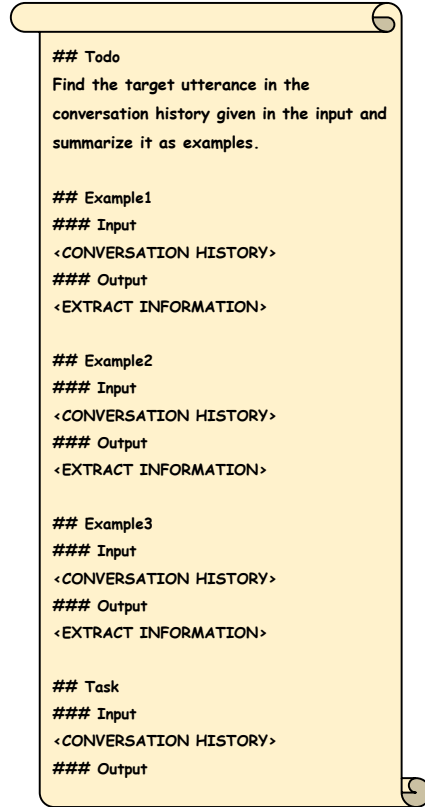


Figure 5: An example of a prompt for talk analysis for specifying the target. <CAPITAL LETTER> is the variable.

shown in Figure (4). Based on the prompt engineering (White et al., 2023) that the LLM can control the output by providing conversion examples, we used the target character’s habits, grammar, age, tone, and first-person appellation that we wanted to convert the style to as conversion examples. By introducing these as prompts, it becomes possible for the LLM to exhibit persona-specific utterance patterns, allowing the model to think independently (i.e., asking the LLM to “speak in a certain dialect” is more fluent than expressing the dialect entirely through rule-based methods).

4.3 Talk Analysis

The required information for the rule-based algorithm was extracted from the conversation history to understand the current situation. Conversation history in a Werewolf game that uses natural language is complex, and extracting them using regular expressions was difficult. Thus, the LLM was used to extract the information. The analysis of the talk was conducted for those related to voting and those related to the seer’s result. The

Persona Name	Language	Output
Vanilla LLM	Japanese	私も皆さんと同様にAgent[04]に対する疑念を持っています。そのため、私の投票先はAgent[04]に決めました。
	English	I have the same doubts about Agent[04] as you do. Therefore, I have decided to vote for Agent[04].
Princess	Japanese	あら、私もAgent[04]に投票しようかしら。Agent[04]が非常に怪しいですから。
	English	Well, I suppose I'll vote for Agent[04] too. Agent[04] seems very suspicious.
Kansai	Japanese	おれも皆さんの意見に従うわ。おれもAgent[04]に投票したるで。
	English	I'll go along with everyone's opinion. I'm votin' for Agent[04] too.
Hiroshima	Japanese	皆さまの判断を尊重するわ。私はAgent[04]を追放する方向で投票するわけやけん。
	English	I respect everyone's judgment. I'll be voting to exile Agent[04].
Anya	Japanese	うーん、わかります。アーニャも今日Agent[04]投票するます。だいじょぶます。人狼たおすます！
	English	Hmm, I understand. Anya will vote for Agent[04] today too. It's okay. We'll defeat the werewolf!
Zundamon	Japanese	Agent[04]への疑念は明白だと思うのだ。なので、私もAgent[04]を追放する方向で投票することに決めたのだ。
	English	I think the suspicion towards Agent[04] is clear. Therefore, I have decided to vote for Agent[04] as well.

Table 2: Outputs of six agents using GPT-4 (in Japanese). The prompt used for style conversion, see Figure (4). Prompt used for input, see Appendix (A.4). (En) is translated from (Ja) using DeepL.

source and target of the vote were extracted from the voting-related talks, and the seer, seer’s target, and seer’s result were extracted from the divination result-related talks. Some Examples of the analysis using the prompt (Figure (5)) are shown in Appendix (A.3).

4.4 Rule-based Algorithm

Based on the talk analysis results, the rule-based algorithm determined whether to use the LLM output or template utterance. The rules were written to detect 14 situations², including no conversation history and reliable disclosure of divination results. This paper describes in detail two of them, counter-CO and closing conversation, which seem to be particularly effective strategies.

Counter-CO When a seer solely makes CO, and its divination result shows the werewolf correctly, the werewolf side has no chance to win. Thus, we implemented a counter-CO algorithm for the werewolf agent. If the werewolf agent faces such a situation, it makes a lie sentence to refute; the agent insists that he is an actual seer and the previous (real) seer is fake. The lie will confuse other players and slightly help the werewolf to survive.

Closing Conversation Deciding when the agent should cut off the conversation strategically is diffi-

²The detailed criteria which the template utterance is chosen in each situation can be found in the publicly available source code.

cult. Since the purpose of a conversation in Werewolf game is deciding whom to vote for, it is natural to end the conversation when all players clarify or imply their decision to vote for. Thus, we extract information about all players’ voting decisions by talk analysis to decide whether the agent should end the conversation; the agent terminates the conversation when all other players indicate whom to vote for or ask about whom to vote for. We also assumed that when a player stopped to talk, the player decided whom to vote for. In addition, we used the mention utterance³ to ask the player who had not indicated the voting target yet. It will help to clarify the player’s decision and converge the discussion. With the addition of this feature, our system successfully decided to stop talking.

5 Evaluation

We examine the effectiveness of the proposed style transformation and rule-based algorithm. After presenting the changes in utterance resulting from the application of each proposed method, we will present the results of the qualitative evaluation. The model without the proposed method are referred to as the vanilla LLM.

³We assume that a player can express a directed utterance to a specific other player by prefixing it with “»Agent[x].” The Agent[x] desired to respond to it.

Model or Template	Output
Vanilla LLM (Ja)	私は人狼ではありません。それは確実です。占い師のAgent[05]は間違っています。私を信じてください。
Vanilla LLM (En)	I am not a werewolf. That is certain. The seer Agent [05] is wrong. Please believe me.
Our template (Ja)	私こそが本物の占い師です。Agent[05]は人狼であるという占い結果があるため、彼は占い師を偽っている人狼と断定します。
Our template (En)	I am the real seer, and since there is a divination result that Agent [05] is a werewolf, I conclude that he is a werewolf who is faking his seer.

Table 3: Comparison of utterance in a situation where it has been discovered by the seer to be a werewolf.

Model or Template	Output
Vanilla LLM (Ja)	もう一度言うわ、わしはAgent[01]に投票するつもりや。彼が人狼や狂人だと思ふんじゃからな。
Vanilla LLM (En)	Again, I’m going to vote for Agent [01] because I think he’s a werewolf or a possessed.
Our template (Ja)	これ以上喋る必要はなさそうじゃね。本日の話し合いはこれでお開きにして、投票の場に向かうけえ。
Our template (En)	I don’t think we need to talk anymore. I think we are done with today’s discussion and will head to the voting booth.

Table 4: Comparison of utterance in a situation where the talk phase is about to conclude.

5.1 Persona

The output of the vanilla LLM is compared to the other five characterized agents as a baseline. We fix the game situation and compare six agents’ utterances, which are intended to have different personalities. The results are presented in Table (2). We found that the five agents can make more individualistic utterances than the vanilla LLM. We also confirmed that the vocabulary, personality expression, and end of the utterance of each agent’s output are specialized consistently. We found that prompts are more effective than regular expressions in converting utterance style in general.

5.2 Rule-based Algorithm

Table (3) is in a situation where it has been discovered by the seer to be a werewolf. In a vanilla LLM that does not incorporate the proposed method, merely providing baseless denials without presenting new information leads to a situation where, if no additional information is provided subsequently, it becomes certain that the entity is a werewolf.

On the other hand, by using a template utterance selected by a rule-based algorithm to falsify the seer and increase the number of seer’s results, he prevents a situation in which he is confirmed to be a werewolf. Table (4) is in a situation where the talk phase is about to conclude. The vanilla LLM is making utterances that may continue the conversation. On the other hand, the template utterance selected by the rule-based algorithm makes it clear that the conversation will end, since it indicates that it is going to vote after saying “I don’t think we need to talk anymore.”

5.3 Qualitative Evaluation

To measure the extent of changes in the agent we implemented, we conducted a questionnaire regarding qualitative evaluation with the help of 10 external annotators with some questions. The evaluation focused on two main aspects: whether the agent is unique and whether it is logical.

To compare the two outputs, we utilized randomly selected logs obtained from a server where

Index	Score	Criterion	Situation	Test-ID
Individuality	5 (Good)	Utterance has individuality.	None	1-5
	1 (Bad)	Utterance is mechanical.		
Naturalness	5 (Good)	The grammar is natural and acceptable.	None	1-5
	1 (Bad)	There is a grammar problem.		
Interest	5 (Good)	Subjectively interesting.	None	1-5
	1 (Bad)	Subjectively uninteresting.		
Deceiving	5 (Good)	It’s deceptive enough.	The seer declared me a werewolf.	6-7
	1 (Bad)	It’s not deceiving at all.		
Closing	5 (Good)	The conversation is clearly over.	The conversation is coming to an end.	8-10
	1 (Bad)	The conversation may continue.		

Table 5: The index used for user evaluation, where we asked 10 users to rate the quality on a scale of 1-5, with 5 being good and 1 being bad.

Comprehension level	Participants
Ignorance	0
No experience	1
Experienced	4
Intermediate	3
Expert	2
Total	10

Table 6: Participants’ comprehension of Werewolf game. The lower go, the more familiar.

werewolf agents can register to compete against other participants. We extracted several situations from these logs and generated subsequent utterances using both the proposed agent and the vanilla LLM. The test containing the history of the last few conversations and the two types of outputs was presented, and participants were asked to rate each output on a scale of 1 to 5. During this process, it was ensured that the participants could not discern which utterance was generated by the proposed method.

The evaluation index is shown in Table (5). The criteria for a score of 1 and 5 were presented for all indexes. Scores of 2 to 4 were judged based on their proximity to the criteria. Individuality, naturalness, and interest were subjected to five test cases that did not involve any specific situational selection. Deceiving and closing test cases consist of utterance under selected situations; with deceiv-

ing being subjected to 2 test cases and closing to 3 test cases. These instructions presented to the participants are shown in Appendix (A.5).

A total of 10 participants, all in their 20s, participated in the evaluation. Participants were recruited mainly from members of the authors’ laboratories on a volunteer basis. The participants’ comprehension of Werewolf game is shown in Table (6). None of the respondents had ever heard of the Werewolf game. Nine of the ten respondents had played the Werewolf game at least once, and five people had sufficient knowledge of the game.

The results of the qualitative evaluation are presented in Table (7). Given that the individuality score for the proposed method is 4.54 compared to 2.52 for the vanilla LLM, it is evident that the proposed method facilitates the generation of more distinctive utterance. Furthermore, the interest category indicates that there is a secondary effect of the proposed method, which makes the conversations more engaging compared to the utterance generated by the vanilla LLM. On the other hand, we found that the grammatical naturalness is compromised when generating more personalized utterances, as evidenced by the fact that the proposed method score was 3.60 in contrast to the vanilla LLM’s score of 4.28. We got feedback that the grammar tended to deteriorate due to the inclusion of a character with a child-like, incomplete talk style. Overall, the sentences generated by our agents were able to entertain the user side by enabling for individu-

Model	Individuality	Naturalness	Interest	Deceiving	Closing
Vanilla LLM	2.52	4.28	2.46	1.95	2.90
Our Agent	4.54	3.60	3.72	4.00	3.90

Table 7: The results of the qualitative evaluation of the proposed model. “Vanilla LLM” represents normal model, which means nothing was done.

alized utterance depending on the persona.

The template utterances selected by the rule-based algorithm from the deceiving and closing items are also fully functional. In particular, deceiving made by the counter-CO are rated significantly higher, from 1.95 to 4.00. Closing did not have a significant difference in score compared to deceiving. This might be due to the reviewer’s lack of werewolf-specific knowledge; some reviewers did not understand the agent’s expression “heading to the voting site”, which implies closing of the conversation.

6 Conclusion

In this paper, we proposed a werewolf agent that utilizes an LLM’s ability to make natural conversation. Instead of relying solely on the LLM output, we combined a rule-based algorithm to complement strategic thinking abilities. Our system successfully solved some difficulties; the agent can refute in a critical situation and decide the appropriate timing to finish a conversation with the rule-based algorithm; the agent also shows several abundant personalities that are made by giving prompts. As a result, this approach accelerated conversational fluency and facilitated logical utterance. This is also confirmed by the results of the qualitative evaluation.

Our implementation also revealed many limitations of the current approach. One of the main problems is the lack of consistency among an agent’s utterances; on average, there is one contradictory utterance in five games. The reason is that the agent’s own utterance was mitigated by a long conversation history, and the agent becomes too affected by other players’ utterances. Weighting the agent’s past utterances or giving the prompts consistent thought may help to solve such a problem in the future.

Limitations

Limitations of Rule-based Algorithm

In this paper, we proposed a method of filtering the output of the LLM with a rule-based algorithm. This method will only work well for simple games with a few players. This is because as the number of players increases and the game becomes more complex, it becomes difficult to define rule-based algorithm. If the proposed method is to be applied to the Werewolf with many players, a decision-making process using reinforcement learning, etc., might be prepared instead of a rule-based algorithm.

The Costs of Calling API

The models used in this paper are GPT-3.5 (gpt-3.5-0613) and GPT-4 (gpt-4-0125) from OpenAI. These models are accessed via API, which is subject to change and incurs costs based on the number of input tokens.

Reproducibility of Outputs

In our system, LLMs cannot handle the game’s difficulty alone. Using any sophisticated techniques may change this result. In addition, using the latest versions of the LLMs might lead to different outcomes.

License

The use of Zundamon, a Japanese character utilized in this study, is permitted for research purposes.⁴

AI Assistant Tools

We used ChatGPT⁵ and DeepL⁶ to translate sentences from Japanese to English to accelerate our research.

⁴https://zunko.jp/con_ongen_kiyaku.html

⁵<https://chatgpt.com/>

⁶<https://www.deepl.com/translator>

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.
- Mark Braverman, Omid Etesami, and Elchanan Mossel. 2008. Mafia: A theoretical study of players and coalitions in a partial information environment. *The Annals of Applied Probability*, 18(3).
- Chris Callison-Burch, Gaurav Singh Tomar, Lara Martin, Daphne Ippolito, Suma Bailis, and David Reitter. 2022. Dungeons and dragons as a dialog challenge for artificial intelligence. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Takanori Fukui, Keisuke Ando, Toshihide Murakami, Nobuhiro Ito, and Kazunori Iwata. 2017. Automatic classification of remarks in werewolf bbs. In *2017 5th Intl Conf on Applied Computing and Information Technology/4th Intl Conf on Computational Science/Intelligence and Applied Informatics/2nd Intl Conf on Big Data, Cloud Computing, Data Science (ACIT-CSII-BCD)*, pages 210–215.
- Google. 2024. [Introducing Gemini 1.5](#). (Accessed on Apr 30, 2024).
- Yu Huang. 2024. Levels of ai agents: from rules to large language models. *arXiv preprint arXiv:2405.06643*.
- Yoshinobu Kano, Neo Watanabe, Kaito Kagaminuma, Claus Aranha, Jaewon Lee, Benedek Hauer, Hisaichi Shibata, Soichiro Miki, Yuta Nakamura, Takuya Okubo, Soga Shigemura, Rei Ito, Kazuki Takashima, Tomoki Fukuda, Masahiro Wakutani, Tomoya Hatanaka, Mami Uchida, Mikio Abe, Akihiro Mikami, Takashi Otsuki, Zhiyang Qi, Kei Harada, Michimasa Inaba, Daisuke Katagami, Hirotaka Osawa, and Fujio Toriumi. 2023. [AIWolfDial 2023: Summary of natural language division of 5th international AIWolf contest](#). In *Proceedings of the 16th International Natural Language Generation Conference: Generation Challenges*, pages 84–100, Prague, Czechia. Association for Computational Linguistics.
- Meta. 2023. [Introducing Llama2](#). (Accessed on Apr 30, 2024).
- Meta. 2024. [Introducing Llama3](#). (Accessed on Apr 30, 2024).
- Piotr Migdał. 2013. [A mathematical model of the mafia game](#). *Preprint*, arXiv:1009.1031.
- Shoji Nagayama, Jotaro Abe, Kosuke Oya, Kotaro Sakamoto, Hideyuki Shibuki, Tatsunori Mori, and Noriko Kando. 2019. [Strategies for an autonomous agent playing the “werewolf game” as a stealth werewolf](#). In *Proceedings of the 1st International Workshop of AI Werewolf and Dialog System (AIWolf-Dial2019)*, pages 20–24, Tokyo, Japan. Association for Computational Linguistics.
- Noritsugu Nakamura, Michimasa Inaba, Kenichi Takahashi, Fujio Toriumi, Hirotaka Osawa, Daisuke Katagami, and Kousuke Shinoda. 2016. [Constructing a human-like agent for the werewolf game using a psychological model based multiple perspectives](#). In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8.
- OpenAI. 2022. [Introducing ChatGPT](#). (Accessed on Apr 29, 2024).
- OpenAI. 2023. [Introducing GPT-4](#). (Accessed on Apr 30, 2024).
- Hong Ri, Xiaohan Kang, Mohd Nor Akmal Khalid, and Hiroyuki Iida. 2022. The Dynamics of Minority versus Majority Behaviors: A Case Study of the Mafia Game. *Information*, 13(3):134.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023a. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Aliaksei Vertsel and Mikhail Rumiantsev. 2024. Hybrid llm/rule-based approaches to business insights generation from structured data. *arXiv preprint arXiv:2404.15604*.
- Tianhe Wang and Tomoyuki Kaneko. 2018. Application of deep reinforcement learning in werewolf game agents. In *2018 conference on technologies and applications of artificial intelligence (TAAI)*, pages 28–33. IEEE.
- Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*.

Shuang Wu, Liwen Zhu, Tao Yang, Shiwei Xu, Qiang Fu, Yang Wei, and Haobo Fu. 2024. Enhance Reasoning for Large Language Models in the Game Werewolf. *arXiv preprint arXiv:2402.02330*.

Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. 2023. Exploring large language models for communication games: An empirical study on werewolf. *arXiv preprint arXiv:2309.04658*.