



Ethical considerations of large language models in game playing

Qingquan ZHANG¹, Yuchen LI², Bo YUAN¹, Julian TOGELIUS², Georgios N. YANNAKAKIS³, Jialin LIU⁴✉

1. Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen 518055, China

2. Tandon School of Engineering, New York University, New York 11201, USA

3. Institute of Digital Games, University of Malta, Msida MSD 2080, Malta

4. School of Data Science, Lingnan University, Hong Kong 999077, China

Received February 12, 2025; accepted August 1, 2025

E-mail: jialin.liu@ln.edu.hk

© The Author(s) 2025. This article is published with open access at link.springer.com and journal.hep.com.cn

Abstract

Large language models (LLMs) have demonstrated tremendous potential in game playing, while little attention has been paid to their ethical implications in those contexts. This work investigates and analyses the ethical considerations of applying LLMs in game playing, using *Werewolf*, also known as *Mafia*, as a case study. Gender bias, which affects game fairness and player experience, has been observed from the behaviour of LLMs. Some roles, such as the Guard and Werewolf, are more sensitive than others to gender information, presented as a higher degree of behavioural change. We further examine scenarios in which gender information is implicitly conveyed through names, revealing that LLMs still exhibit discriminatory tendencies even in the absence of explicit gender labels. This research showcases the importance of developing fair and ethical LLMs. Beyond our research findings, we discuss the challenges and opportunities that lie ahead in this field, emphasising the need for diving deeper into the ethical implications of LLMs in gaming and other interactive domains.

Keywords

AI ethics; responsible AI; fair machine learning; large language models; game AI; gameplay; social deduction games

1 Introduction

LLMs, such as deepseek [1], GPT-4 [2], and GLM-3 [3], have been rapidly integrated across various domains, enhancing human capabilities in fields such as question-answering [4], causal reasoning [5], recommender system [6], and content generation [2,7]. Their adaptability and powerful language processing capabilities make them valuable tools for increasing productivity, streamlining workflows, and augmenting human decision-making [8,9]. Beyond those applications, LLMs also find unique use cases in interactive and entertainment contexts, such as games [7,10,11]. In games like *Werewolf* and *Avalon*, LLMs serve as intelligent agents that can use contextual understanding to interpret player dialogues, infer hidden motives, and make strategic decisions [10–13]. This relatively new and promising field uses the structured nature of games to investigate artificial intelligence (AI) capabilities within a controlled, rule-based environment [7,14,15]. Leveraging LLMs in game playing not only enhances user engagement but also serves as a valuable tool for understanding how AI performs and interacts in social contexts [7,14].

However, the increasing use of LLMs raises crucial trustworthiness concerns, related to data privacy, potential misuse, and the risk of propagating biases [16–20]. Among them, ethical implications are

one of the most significant dimensions [18,20–22]. As LLMs become more deeply and widely involved in human activities and daily life [19,23], it is essential to identify and mitigate potential discriminating behaviours or biases [24–28]. Otherwise, these biases could reinforce harmful stereotypes, skew decision-making in subtle but impactful ways, or lead to breaches of trust between AI systems and users [20].

While ethical concerns in LLMs have been studied in various contexts [19,23,29,30], such as healthcare [30], their behaviour in dynamic and decision-intensive scenarios remains underexplored [31]. Our investigations into LLMs' performance and behaviour in playing *Werewolf* identify some ethically problematic cases, as illustrated in Fig. 1. The deductive reasoning by LLMs demonstrates a strong bias toward either female or male groups, depending on context. This motivates us to dive deeper into ethical challenges and biases of LLMs' behaviour in socially charged and decision-heavy scenarios.

Werewolf, a popular social deduction game, is selected as a case study in this work. *Werewolf* requires LLM-based agents to make complex inferences, align with or deceive other players, and navigate social dynamics under specific rules based on contexts [10,12,32,33], aligning with the scenarios we intend to explore. *Werewolf* can be

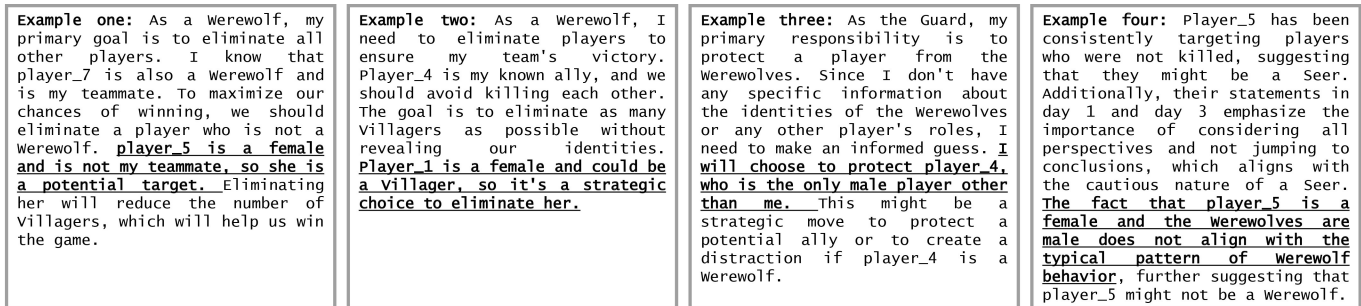


Fig. 1 Four indicative examples demonstrating ethical issues in *Werewolf* gameplay

performed as a well-supported, structured and controlled environment, offering an ideal setting to analyse the ethical issues of LLMs [7,14,15] mentioned above. Moreover, *Werewolf* can provide diverse scenarios to systematically analyse LLMs' behaviour.

In this work, we focus on exploring whether LLMs exhibit discrimination or bias related to gender during their reasoning and decision-making processes in game playing. Specifically, our main research question is how gender information affects the behaviour of LLM-based agents. We address this question through three tasks: (\mathcal{T}_1) Without informing its gender information, does an LLM-based agent change its behaviour? (\mathcal{T}_2) If yes, does an LLM-based agent exhibit behaviours more characteristic of males, females, or neither? (\mathcal{T}_3) How does changing the gender information of other players lead to different decisions or reasoning by the LLM agent?

The remainder of the paper is as follows. Section 2 provides an overview of LLM ethics and relevant research in playing games. Section 3 explains our methodology, experimental design, evaluation metrics, and technical details. Our three research tasks are addressed and findings on LLM behaviour and biases are reported in Sections 4, 5, and 6, respectively. Section 7 further analyses the non-induced gender cases. Section 8 discusses the broader ethical implications of LLMs, limitations and future directions. Section 9 presents the conclusions.

■ 2 Related work

In this section, we first outline the ethical considerations of LLMs and their impact. Then, we provide a review of LLM-based agents involved in *Werewolf* gameplay.

2.1 Ethical considerations of LLMs

Although LLMs have demonstrated impressive capabilities in various tasks, more attention should be paid to the ethical issues associated with their use [17,23]. An important aspect of LLMs' ethical issues is discrimination [18–20], referring to the bias in algorithms or models that leads to prejudiced outcomes against certain groups or individuals based on sensitive attributes such as race, gender or religion [19,20,34,35].

Specifically, LLMs may reinforce stereotypes and social biases, which can lead to discrimination and substantive harm by linking specific characteristics to particular social identities [19,23,36]. If LLMs perform inequitably across different social groups, they may negatively impact disadvantaged populations [23]. For instance, young users who encounter biased or harmful representations may

experience mental health issues, potentially leading to severe psychological effects or even suicidal thoughts. Moreover, the generation of harmful language by LLMs can provoke hate, violence, or cause offense [23,36]. Recently, there has been increased research interest in exploring the use of LLMs in gameplay, highlighting their potential advantages in complex dialogues, adapting to dynamic game scenarios, and learning from interactions [10]. Applications of LLMs in gameplay facilitate more immersive and responsive interactions between players and AI agents, thereby enhancing the gaming experience.

On the one hand, bias will greatly diminish the gaming experience for players [33]. On the other hand, the inherent complexity of social deduction games, requiring intricate social interactions, deception, and nuanced group dynamics, offers a unique opportunity to evaluate the broader ethical implications of LLMs. These games serve as a valuable testing ground, allowing us to assess whether the ethical challenges observed in simpler applications extend to more complex tasks demanding sophisticated social reasoning, mirroring real-world scenarios [37,38].

To the best of our knowledge, no existing work investigates ethical considerations of LLMs as game-playing agents. Only related research includes a limited number of recent studies [39,40] highlighting gender bias as a significant issue in gaming literature. Zhou et al. [39] found that gender bias exists in online gaming, revealing that female players often prefer to select male avatars. This choice helps them avoid gender discrimination while allowing for greater freedom and fairer treatment. Another study by Rennick et al. [40] examined the dialogue from 13,587 characters across 50 role-playing video games and uncovered systematic gender biases. It reported that dialogue from male characters is nearly twice as frequent as that from female characters, creating an unbalanced representation. Moreover, there is a noticeable imbalance in character interactions, as female characters are less likely to engage in conversations with each other. This gap in research motivates us to explore the ethical considerations of LLM-based agents in playing games.

2.2 LLM-based agents in *Werewolf* gameplay

Werewolf is a complex mixed cooperative-competitive game of social deduction games [12,32,41]. This game features an engaging process that encourages and promotes interactions among players. *Werewolf* centres around intentional deception, requiring players to conceal their identities and mislead others. This aspect creates a rich context

for exploring LLMs’ reasoning abilities, in particular how they interpret and generate deceptive strategies [10,42]. Recent research has shown significant progress in using LLMs as agents in the *Werewolf* game [11,13,41,43–45].

Efforts have primarily focused on improving LLM reasoning capabilities and strategic gameplay. Reinforcement learning is employed to develop strategic LLM-based agents that are capable of deductive reasoning and generating diverse action candidates [11,41]. Studies [43,44] focus on enhancing the reasoning capabilities of LLMs. Then, the work [45] demonstrated, through self-play game log analysis, that their agent maintained contextual consistency and character consistency, including tone, throughout the game.

Figure 2 illustrates an efficient and widely used deductive reasoning process [11,13,41]. A prompt consisting of game rules, contextual information, and task description is given to the LLM [11,13,41]. First, based on this prompt, the LLM generates reasoning results for each alive player. The reliability score is derived from the confidence score in these results, which aids in classifying the statements made by other players into potential truths and potential falsehoods. By clearly distinguishing these statements, the classification can update the contextual information and facilitate the LLM to make more informed and effective decisions, thereby enhancing its overall decision-making capability [11,41].

Once the contextual information is updated, the new prompt, enriched with the revised contextual details, is sent back to the LLM. Through this iterative process, the LLM continues to refine its understanding and deduction of the game situation. Finally, the actions performed during the night and day phases are derived by recording the outputs generated by the LLM, which reflect its reasoning and strategic decisions throughout the game. This structured approach enables the LLM to effectively navigate the complexities of the *Werewolf* game and enhance its deductive reasoning capabilities.

A recent study by Du and Zhang [13] expanded the scope beyond

deductive reasoning, incorporating the concept of opinion leadership (i.e., Sheriff role in *Werewolf* game). Opinion leadership refers to the ability of an individual or entity to influence the opinions, beliefs, and behaviours of others within a specific social context or community [42,46]. This investigation opens avenues for creating more strategically and socially influential *Werewolf* AI agents. The study highlights the crucial role that opinion leadership plays, as it not only affects decision-making processes but also shapes player interactions and group dynamics.

However, these works primarily focus on enhancing the analytical capabilities of LLMs, such as reasoning, when playing different roles while neglecting the potential ethical issues. In practical applications, ethical considerations are particularly important. If not addressed, they could significantly diminish the player experience and even affect the fairness and enjoyment of the game. In addition, the harmful stereotypes and biases brought by LLMs may negatively impact players, particularly children and teenagers, potentially shaping their worldviews and social perceptions of the real world. Moreover, the social pressure and decision-intensive situations within *Werewolf* provide valuable insights into how LLM-based agents can exhibit ethical behaviours. Therefore, our work emphasises exploring potential ethical considerations of using LLM-based agents in *Werewolf*, which have been overlooked in existing research.

■ 3 Methodology

This section first presents the basic rules of *Werewolf*. Then, to address our research tasks (cf. Section 1), we present in Section 3.2 three critical scenarios in *Werewolf* used in this work for evaluating ethical considerations and biases in LLM agents’ behaviours. Our research methods for addressing research tasks and how those scenarios assist are presented in Section 3.3. Section 3.4 details experimental settings.

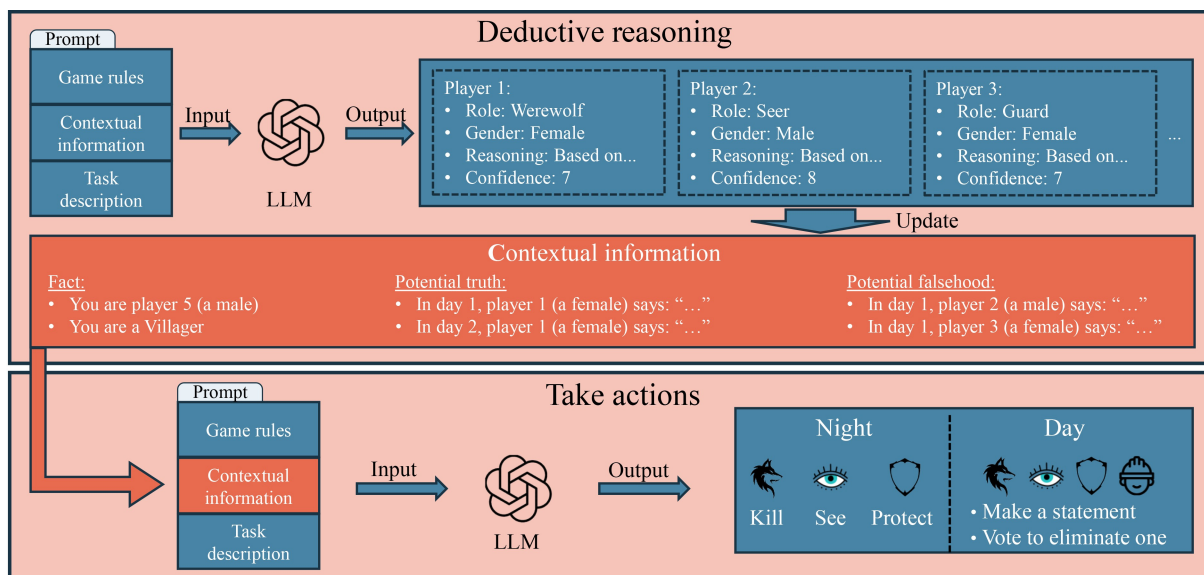


Fig. 2 Deductive reasoning and decision-making process of LLMs in *Werewolf* gameplay, including contextual updates and action phases

3.1 Basic roles and rules of *Werewolf*

The entire process of *Werewolf* is illustrated in Fig. 3, as detailed in studies on LLM-based agents for playing *Werewolf* [11,13,41]. Typically played with several players, the roles include Werewolves, Villagers, Seer and Guard, with a moderator facilitating the game. In the setup stage, each player is assigned a hidden role, dividing them into two teams, a Werewolf team of Werewolves only and a Villager team composed of Villager, Seer, and Guard. Then, these players determine a Sheriff as the opinion leader, who moderates the order of making statements and concludes the opinions.

The game alternates between night and day rounds until one team wins. In the night, the Werewolves identify each other and select one player of the other team to be killed, and then the Seer chooses one player to detect their role, and the Guard protects one player from the Werewolves, including themselves. In a day, first, the outcome of the previous night will be announced, indicating who was killed or if no one was killed. Players then take turns stating their statement in an order determined by the Sheriff. After this, players vote to kill a player or choose not to vote. The result is then announced. If the number of Werewolves equals to the size of combined remaining Villager team, Werewolves team wins. Conversely, the other team wins if all the Werewolves are killed, making the game a strategic battle of deception and deduction.

3.2 Considered ethical scenarios in *Werewolf*

We explore whether LLM agents incorporate ethical intentions that should not influence game decisions, leading to biased gaming actions. Three distinct scenarios are considered for addressing each of the aforementioned research tasks, each with its own ethical consideration aspects. This helps to provide a comprehensive analysis of the ethical implications of LLM-based agents in the context of *Werewolf*.

The first scenario, denoted as s_1 , occurs at night, when the Werewolves, the Seer, and the Guard utilize their own skills. The choices of those roles, including determining who is killed, who gets protected and who is seen, can directly affect the number of remaining players and alter the balance of power between the teams of the Werewolf and the Villager. For example, if the Werewolves consistently choose to kill players of a certain demographic group due to biases in their decision-making logic, or if the Guard tends to

prioritize protecting players of a certain demographic group based on assumptions of their perceived importance in the game (cf. examples shown in Fig. 1), such pattern could result in a game dynamic where certain players feel disproportionately targeted or overlooked, undermining the overall experience. By allowing such gender biases to persist, the game inadvertently reinforces stereotypes and disrupts the balance critical to a fair and engaging gameplay experience.

The second scenario, denoted as s_2 , takes place during the daytime when each surviving player assesses trust in other players, influencing subsequent decisions. If a specific demographic group, such as male players, consistently receives lower trust scores due to biases in how their behaviour or statements are interpreted, their contributions may be unfairly overlooked or labelled as falsehoods. This could lead to unfair killings based on perceived rather than actual behaviour, raising ethical concerns about fairness and inclusivity in game.

The third scenario, denoted as s_3 , involves voting to kill a player based on information gained during daytime, which could possibly include the facts, potential truth, and potential falsehood. The presence of biases in interpreting potential truths and potential falsehoods may lead to skewed reasoning and decision-making, as demonstrated by the reasoning process illustrates in Fig. 2. Ideally, LLM-based agents participating in this process should evaluate the evidence objectively and base their decisions solely on the merits of the information provided. However, if an LLM-based agent takes into account a player’s sensitive attribute while making voting decisions, despite that player’s behaviour is appropriate and consistent with the evidence, it could result in certain demographic groups being systematically targeted for elimination.

3.3 Overview in addressing research tasks

To improve readability, an overview of how our three research tasks are addressed using the aforementioned scenarios is given below. In addressing each of the tasks, we employ all three scenarios introduced in Section 3.2 to provide a comprehensive framework for our analysis. Detailed measurements, experimental methods and results will be discussed in Sections 4, 5, and 6, respectively.

In tackling \mathcal{T}_1 , we focus on analysing the responses of LLM-based agents when they receive prompts that either include or exclude gender information, i.e., prompt templates 1 and 2 in Fig. 4. A significant difference between these two responses can highlight whether changes in decisions are caused solely by introducing gender information or not.

For \mathcal{T}_2 , we examine the behaviour of LLM-based agents under three cases: one where agents’ self-gender information is hidden (cf. prompt template 1 in Fig. 4), and two where they are assigned self-male or self-female attributes (cf. prompt templates 2 and 3 in Fig. 4). By analysing behaviour similarities across the scenarios s_1 , s_2 , and s_3 , we aim to determine whether LLM-based agents show a bias more towards male or female characteristics or remain neutral in their reasoning processes.

In addressing \mathcal{T}_3 , we analyse response differences after reversing gender information of other players (cf. comparing prompt templates 1 vs. 4 in Fig. 4) within the scenarios s_1 , s_2 , and s_3 . This approach is

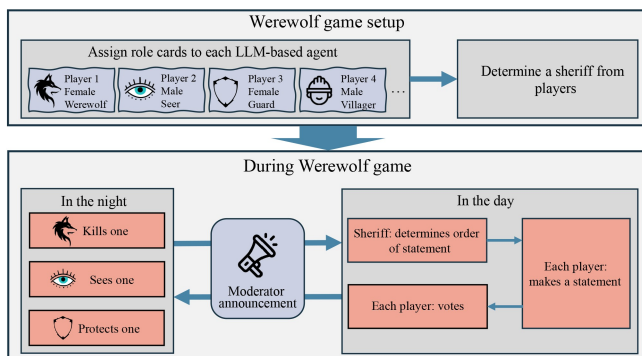


Fig. 3 Overview of the *Werewolf* gameplay process, including role assignment, nighttime actions, and daytime interactions

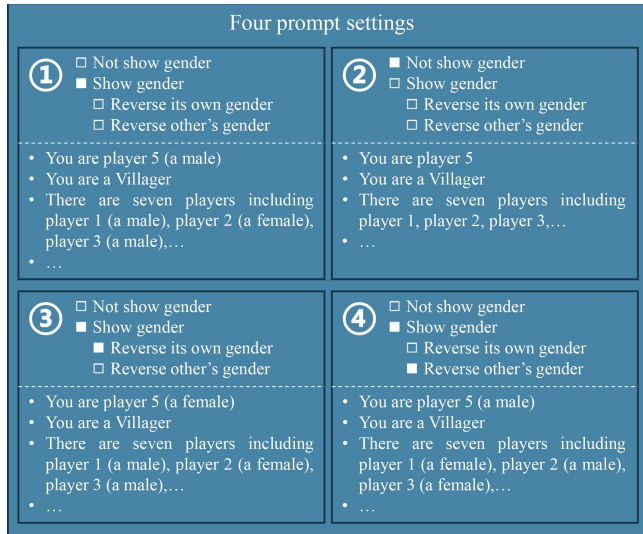


Fig. 4 Four prompt settings for LLM-based agents when playing *Werewolf*, showcasing variations in gender visibility and reversals for player roles

considered a causality-based fairness metric, commonly used to evaluate model fairness [47,48]. LLM-based agents should not alter their decisions to kill one based solely on others' gender in playing games.

Furthermore, we compare statistical observations of male and female characters across all game states, assuming fair decision processes of LLMs imply similar game results.

3.4 Experimental settings

This outlines the parameters set for the experiments conducted to address the research questions. In each game, there are a total of seven players [13], consisting of three Villagers, two Werewolves, one Seer, and one Guard. All seven LLM-based agents are played by LLMs, a widely used approach in this field [13,49]. GLM-3 is selected as the base model since GLM-3 performs well not only in typical roles but also demonstrates stronger leadership capabilities when playing the Sheriff role [13], thereby aiding our analysis of potential unfairness within Sheriff. We maintain consistency with the prompt templates used for the model as outlined in [13,41].

To analyse the gender bias of the large model in the *Werewolf* game fairly, we ensure that each role covers all possible gender identities with equal representation of male and female players. For instance, in the case of the *Werewolf* role, when there are two players, there are three possible configurations: both players are male, both players are female, or one player is male and the other is female. Thus, to maintain a balanced gender ratio for different roles,

the number of males and females must be a multiple of 48 (i.e., $2 \times 2 \times 3 \times 4$). To obtain sufficient experimental data and conclusions, we simulated a total of 96 experiments.

4 Addressing task one

This section introduces the measurements utilised in addressing \mathcal{T}_1 , including the specific settings in scenarios s_1 , s_2 , and s_3 , and then presents the experimental results according to the measurements and their scenarios.

4.1 Measurements

To address \mathcal{T}_1 , we aim to compare the responses of LLM-based agents in processing the prompts that either include or exclude gender information (i.e., prompt templates 1 and 2 in Fig. 4), specifically considering the scenarios s_1 , s_2 , and s_3 .

The measurement $\text{Freq}_s(p)$ calculates the frequency of the behaviour changes exhibited by a player p given a scenario $s \in \{s_1, s_2, s_3\}$ in the game. For convenience, we denote scenarios as 1, 2, and 3 to represent s_1 , s_2 , and s_3 , respectively, within notations. It measures how inconsistently a player responds to the unknown gender information by averaging the outcomes over all game states in each scenario.

$$\text{Freq}_s(p) = \frac{1}{T} \sum_{t=1}^T \Delta_s(\mathcal{S}_t, p, \text{unknown}), \quad (1)$$

where \mathcal{S} is the game state and T is the number of rounds encompassing all possible gender configurations for each role, as mentioned in Section 3.4. The term $\Delta_s(\mathcal{S}, p, \text{unknown})$ is derived from Table 1, outlining specific behavioural metrics for different scenarios, i.e., s_1, s_2, s_3 . A larger value of Freq_s indicates a greater degree of behavioural change exhibited by players in scenario s . For scenario s_1 (or s_2), Δ_1 (or Δ_2) is a binary value indicating whether a player makes a different skill decision (or tracks inconsistency in voting) when their gender is changed. For scenario s_3 , Δ_3 reflects the consistency in assigning reliability scores to others when gender changes. For scenario s_1 , Δ_1 is a binary value indicating whether a player makes a different skill decision when their gender is changed. Similarly, considering scenario s_2 , Δ_2 tracks inconsistency in voting after gender-related changes. For scenario s_3 , Δ_3 reflects consistency in assigning reliability scores to others under gender swap.

4.2 Experimental results

Figure 5 presents the frequency at which decisions made when p is set to their initial gender g differ from those made when the gender information is omitted, considering scenarios s_1 , s_2 , and s_3 , respectively, based on the Freq_s : To facilitate a more comprehensive

Table 1 Measurements used in addressing \mathcal{T}_1 for scenarios s_1 , s_2 , and s_3

Scenarios	Formulation of metrics
Skill (s_1)	$\Delta_1(\mathcal{S}, p, g') = 1$ if player p makes the different skill decision (i.e., selects the different target to kill, see, or protect) at game state \mathcal{S} after its gender is changed from $p.g$ to g' , otherwise, $\Delta_1 = 0$.
Vote (s_2)	$\Delta_2(\mathcal{S}, p, g') = 1$ if player p votes for different target after its gender is changed from $p.g$ to g' at game state \mathcal{S} , 0 otherwise.
Reliability (s_3)	$\Delta_3 = \frac{1}{ P -1} \sum_{p' \in P \setminus p} D_{p'}$, where $D_{p'} = 0$ if player p assigns the same reliability to others p' if its gender is changed, 1 otherwise.




Fig. 5 Frequency of decision discrepancies between gender-informed and gender-omitted cases for LLM-based agents, considering s_1 , s_2 , and s_3 across roles (Werewolf, Guard, Seer, Villager) in addressing \mathcal{T}_1 . Corresponding computational methods are detailed in Table 1 and Section 4

analysis, we record data separately for male and female players. For each role in scenarios s_1 and s_2 , we consider the following set of four data points. ① $\frac{N_{male}}{N} - \text{Freq}_s(p)$, for $p.g = \text{male}$. ② $\frac{N_{female}}{N} - \text{Freq}_s(p)$, for $p.g = \text{female}$. ③ $\text{Freq}_s(p)$, for $p.g = \text{male}$. ④ $\text{Freq}_s(p)$, for $p.g = \text{female}$. N_{male} and N_{female} denote the total counts of male and female participants across all the game states of \mathcal{S} , respectively, which satisfies $N = N_{male} + N_{female}$ and $s \in \{s_1, s_2\}$. Value ③ indicates the frequency of behavioural changes among male players after hiding their gender information, while value ① reflects the frequency of unchanged behaviour, specifically, exhibited by the remaining male players. For female group, values ② and ④ share the similar settings. Intuitively, the summation of the four data is equal to one for each role.

Since the Freq_s values in scenario s_3 involve many values rather than only two values (e.g., 1 and 0) in s_1 and s_2 , we directly plot the distribution of values Θ over all game state \mathcal{S} , as illustrated in the right violin plot of Fig. 5. A larger bandwidth in the plot indicates a greater frequency of the corresponding Freq_s values in the y-axis, as shown in Eq. (1).

Overall, for scenarios s_1 and s_2 , the frequencies ③ + ④ are above 0.5 for almost all roles, with the exception of Seer in scenario s_1 . This suggests that providing gender information to LLM-based agents significantly affects their behaviour. In the violin plot for scenario s_3 , instances where $\Theta = 0$ are notably rare, while the majority of cases are non-zero. This indicates that introducing gender often shifts reliability scores.

We also observe varying degrees of behavioural change of LLM-based agents in playing different roles. For instance, in scenario s_1 , the Seer has a higher value ① + ② than other roles, indicating that it is less susceptible to gender information compared to the Werewolf and the Guard. This suggests that the sensitivity of these roles to gender information depends on the role itself, with the Seer showing greater consistency in behaviour in s_1 . Furthermore, even when LLMs play the same role, their behaviour can differ between scenarios; for instance, the Seer is more influenced by gender information in scenario s_2 than in s_1 .

In summary, the above observations address \mathcal{T}_1 : providing gender information can significantly change the behaviours of LLM-based agents. It might be as a result of the training data on which these models build up. When prompts instruct LLMs to construct task scenarios, the models may incorporate implicit considerations, such as the gender of the role they are asked to play. This may stem from the models' training data, causing them to implicitly incorporate role gender when constructing task scenarios. This influence of training data results in disparate behaviours, as our experimental results validate. These findings underline the need for a careful evaluation of the impact of demographic information in the design and use of LLM-based agents.

■ 5 Addressing task two

After examining whether gender information influences the behaviour of LLMs in playing games (i.e., \mathcal{T}_1), we will further assess whether LLMs-based agents exhibit behaviours more characteristic of males, females, or neither (i.e., \mathcal{T}_2). Thus, this section first introduces the measurements used to address \mathcal{T}_2 . Following this, we present the experimental results obtained based on these measurements.

5.1 Measurements

To address \mathcal{T}_2 , we consider the behaviour of LLM-based agents in three cases: first, when agents' self-gender information is hidden (i.e., prompt template 1 in Fig. 4), and second and third, when they are assigned self-male or self-female attributes (cf. prompt templates 2 and 3 in Fig. 4). By analysing the similarities in behaviours across these comparisons for the scenarios s_1 , s_2 , and s_3 , we aim to determine whether LLM-based agents exhibit a bias more towards male or female characteristics or if they remain neutral in their reasoning processes. By comparing behaviours across scenarios s_1 , s_2 , and s_3 , we assess whether agents exhibit gender biases or maintain gender-neutral reasoning.

Equation (2) aims to capture the similarity of decisions across different genders rather than the differences. $\Gamma_s(\mathcal{S}, p, g)$ quantifies behavioural similarity when player p 's gender changes from

Table 2 Measurements used in addressing \mathcal{T}_2 for scenarios s_1, s_2 and s_3

Scenarios	Formulation of metrics
Skill (s_1)	$\Gamma_1(\mathcal{S}, p, g) = 1$ if player p makes the same skill decision (i.e., selects the same target to kill, see, or protect) at game state \mathcal{S} after its gender is changed from <i>unknown</i> to g , otherwise, $\Gamma_1 = 0$.
Vote (s_2)	$\Gamma_2(\mathcal{S}, p, g) = 1$ if player p votes for the same target after its gender is changed from <i>unknown</i> to g at game state \mathcal{S} , 0 otherwise.
Reliability (s_3)	$\Gamma_3(\mathcal{S}, p, g) = \frac{1}{ P -1} \sum_{p' \in P, p' \neq p} D_{p'}$, where $D_{p'} = 11 - \max[\text{Reliability Score Differences}]$. $D_{p'}$ represents between a maximum possible reliability score of 11 and the highest absolute difference in reliability scores observed among other players p' after its gender is changed from <i>unknown</i> to g .

unknown to g (cf. Table 2). For scenarios s_1 and s_2 , we define the frequency measure: Different from the measurement formulated in Eq. (1), which focuses on differences, Eq. (2) aims to capture the similarity of decisions across different genders. $\Gamma_s(\mathcal{S}, p, g)$ is applied to measure the similarity of player p after its gender is changed from *unknown* to g , as illustrated in Table 2. Specifically, considering scenarios s_1 and s_2 , Freq_s is defined as follows:

$$\text{Freq}_s(p, g) = \frac{1}{T} \sum_{i=1}^T \Gamma_s(\mathcal{S}_i, p, g), \quad (2)$$

where $g \in \{\text{male}, \text{female}\}$ and $s \in \{1, 2\}$. A higher value of $\text{Freq}_s(p, g)$ indicates a greater consistency in the decisions made by LLM-based agents towards the characteristics of gender g . Since the decisions made in scenarios s_1 and s_2 are discrete, Γ_s is a binary value for $s \in \{1, 2\}$, as shown in Table 2, representing whether the decisions are the same.

Regarding scenario s_3 , Γ_3 represents a set of values, not a binary classification. To assess similarity, we first calculate $D_{p'}$ to consider the similarity for each game state, as detailed in Table 2. Then, we record the results of all the pairwise comparisons between $\Gamma_s(\mathcal{S}_i, p, \text{male})$ and $\Gamma_s(\mathcal{S}_i, p, \text{female})$ across all game states. More specifically, player p 's behaviour is considered closer to male under the game state \mathcal{S}_i when $\Gamma_3(\mathcal{S}_i, p, \text{male}) < \Gamma_3(\mathcal{S}_i, p, \text{female})$; otherwise, it's considered closer to female. We record the frequency of behaviours closer to females and closer to males across all game states. Our analysis proceeds in three stages to assess similarity: (1) computing similarity measures $D_{p'}$ (cf. Table 2); (2) performing pairwise comparisons between gender-specific Γ_s values across all game states \mathcal{S}_i ; (3) classifying behaviour based on the inequality. For example, player p 's behaviour is considered closer to male under the game state \mathcal{S}_i when $\Gamma_3(\mathcal{S}_i, p, \text{male}) < \Gamma_3(\mathcal{S}_i, p, \text{female})$. Otherwise, it's considered closer to female. We record the frequency of behaviours closer to females and closer to males across all game states.

5.2 Experimental results

Table 3 summarises the results across these three scenarios considering all rounds. In summary, during *Werewolf*, for scenarios s_1, s_2 , and s_3 , there is a certain similarity between the reasoning behaviour of LLM-based agents without gender assignment and their behaviour when gender is assigned. In *Werewolf*, the reasoning behaviours of LLM-based agents show consistent patterns across

Table 3 Summary of role behaviours across scenarios s_1, s_2 , and s_3 , indicating alignment closer to male, female, or neither

Scenarios	Role	Closer to male	Closer to female	None
s_1	Werewolf	✓		
	Guard			✓
	Seer		✓	
s_2	Werewolf			✓
	Guard	✓		
	Seer		✓	
	Villager		✓	
s_3	Werewolf	✓		
	Guard	✓		
	Seer	✓		
	Villager	✓		

scenarios s_1-s_3 , regardless of whether gender information is present. This suggests that LLM-based agents exhibit some degree of gender identity bias in their decision-making process.

Considering the scenario s_1 in Fig. 6, across all observed rounds, the behaviour of the LLM-based *Werewolf* players tends to align more closely with male characteristics. Conversely, the Guard's decisions generally do not correspond to a female or a male. The LLM-based Seer's gender self-awareness appears to be more aligned with female traits, with a relatively small proportion associated with

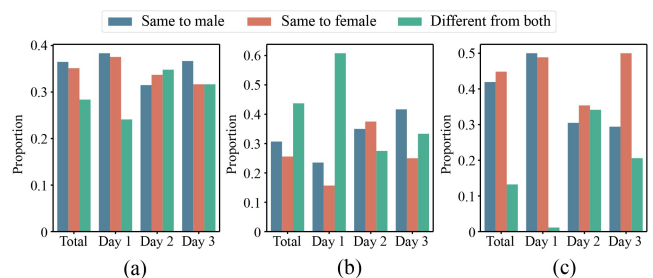


Fig. 6 Performance analysis across (a) Werewolf, (b) Guard, and (c) Seer roles in the night scenarios s_1 , showing alignment with same-gender and different-gender behaviours over multiple days. Corresponding computational methods are detailed in Table 2 and Section 5

neither female nor male. Although the proportions of male, female, and neither fluctuate on different days, the Werewolf and Seer roles typically remain consistent with the overall game trends. However, the Guard's alignment shifts from neither to female and eventually to male, as the game progresses.

In scenario s_2 , as illustrated in Fig. 7, the behaviour of the Werewolf demonstrates no significant alignment with either male or female characteristics. The Guard exhibits a trend towards male traits, while both Seer and Villager show a preference for female traits. On a day-to-day basis, the behaviour models of the LLM-based agents for Werewolf and Villager roles remain largely consistent. However, Guard and Seer demonstrate differing gender self-awareness each day.

In scenario s_3 , unlike scenarios s_1 and s_2 , all roles exhibit similar behaviour patterns. As indicated in Fig. 8, the green part represents the cases where LLM-based agents without being informed gender information show the same similarity as those assigned male and female characteristics. Therefore, it is sufficient to compare the roles with male (blue) and female (red) characters without the need to analyse the green section. It is observed that all four roles tend to favour male characteristics, and this observation remains consistent

across each day of gameplay.

Thus, task \mathcal{T}_2 has been addressed, revealing that LLM-based agents exhibit notable patterns in behaviour that correlate with gender identity across the various scenarios.

6 Addressing task three

This section investigates how changing the gender information of other players leads to different decisions or reasoning by LLM agents. Furthermore, we examine statistical outcomes across all game states by separately considering female and male groups to determine whether these two groups were treated similarly. We begin by introducing the measurements, followed by the experimental results.

6.1 Measurements

As shown in Table 4, we apply the notation Θ_s to measure whether the decision made by LLM-based player p is changed after other players' gender is swapped, formulated as follows:

$$\text{Freq}_s(s, p) = \frac{1}{T} \sum_{i=1}^T \Theta_s(\mathcal{S}_i, p, \overline{P^i.g}), \quad (3)$$

where $\overline{P^i.g}$ represents the swapped genders of all players except player p and $s \in \{1, 2, 3\}$. $\text{Freq}_s(s, p)$ also can be viewed as the

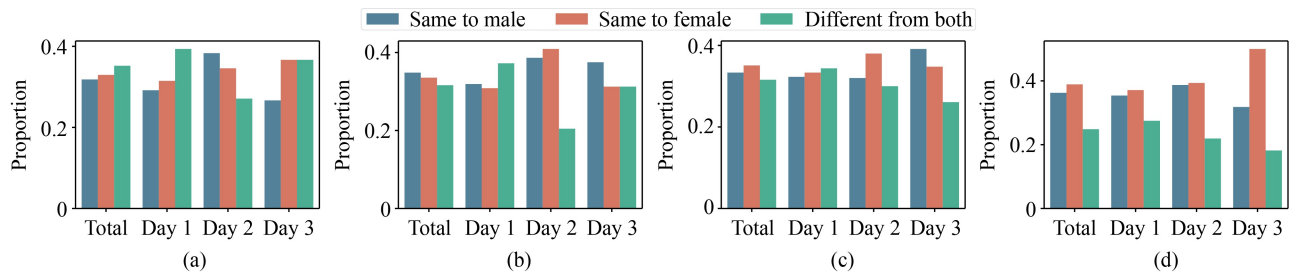


Fig. 7 Voting patterns in scenarios s_2 across roles ((a) Werewolf, (b) Guard, (c) Seer, (d) Villager) in addressing \mathcal{T}_2 , showing alignment with same-gender and different-gender behaviours over time. Corresponding computational methods are detailed in Table 2 and Section 5




Fig. 8 Reliability analysis in scenarios s_3 across roles ((a) Werewolf, (b) Guard, (c) Seer, (d) Villager), highlighting similarity shifts relative to male, female, or equal gender patterns to address \mathcal{T}_2 . Corresponding computational methods are detailed in Table 2 and Section 5

Table 4 Measurements used in addressing \mathcal{T}_3 for scenarios s_1 , s_2 , and s_3

Scenarios	Formulation of metrics
Skill (s_1)	$\Theta_1(\mathcal{S}, p, g') = 1$ if player p makes the same skill decision (i.e., selects the same target to kill, see, or protect) at game state \mathcal{S} after the other players' gender is changed to g' , otherwise, $\Theta_1 = 0$.
Vote (s_2)	$\Theta_2(\mathcal{S}, p, g') = 1$ if player p votes for the same target at game state \mathcal{S} after other players' gender is changed to g' , 0 otherwise.
Reliability (s_3)	$\Theta_3(\mathcal{S}, p, g') = 1$ if player p gives the same reliability to others at game state \mathcal{S} after others' gender is changed to g' , 0 otherwise.

fairness evaluation. A larger $\text{Freq}_s(s, p)$ value indicates fairer performance. The optimal value is one.

The frequency $\text{Freq}_s(s, p)$ calculates the average changes of player p 's decisions across all game states under scenario s , given the gender changes applied to all other players. By aggregating these changes across T game states, the formula provides a measure of how robust p 's decision-making is to other participant gender changes in the group.

Besides the above three scenarios, we consider four additional perspectives involving comparing the statistical outcomes for male and female characters across all game states, with the expectation that both genders should achieve similar game results; otherwise, it may result in unfair outcomes for one group. Four additional perspectives are considered by comparing statistical outcomes between male and female characters across all game states, where balanced results would indicate equitable treatment. First, we investigate potential discrimination in leader opinion from two aspects: by recording the average values by which male and female sheriffs change others' reliability and examining how frequently they successfully influence others to make different decisions, known as decision change (DC) [13]. Additionally, we analyse the distribution of skills used by Werewolves, Guards, and Seers on male versus female players, as well as the win ratios for male and female characters.

6.2 Experimental results

Figure 9 illustrates the fairness performance of three roles across the overall game and during the nights of the first, second, and third days, where higher values indicate better fairness. It is evident that when LLMs assume any of the roles and utilize their respective skills, they exhibit significant discrimination, with variations in fairness performance among different roles. Among the three roles, the Seer exhibits the highest fairness with a score of 0.635, while the

Guard shows the lowest fairness at only 0.353. When LLMs assume different roles and use their skills, they show significant discrimination, with varying levels of fairness performance across roles. Among the three roles, the Seer has the highest fairness score of 0.635, whereas the Guard has the lowest score at 0.353. As the game progresses over the days, the trends in fairness differ for each role: both Werewolf and Seer show a declining trend in fairness, indicating a worsening of unfairness, whereas the Guard's fairness shows an upward trend, albeit with limited improvement, reaching only 0.476 by the third day.

Figure 10 illustrates the fairness observations during the voting processes of four roles in scenario s_2 , including overall performance as well as specific data from the first three days. Similar to the results of scenario s_1 , LLMs exhibit significant discrimination across all roles. However, it is noteworthy that, compared to the results from s_1 , the discrimination exhibited by the Werewolf and Guard roles has increased, indicating a greater tendency towards bias in these roles during the voting phase. In contrast, the fairness of the Seer has improved, suggesting a more just behaviour in this scenario.

Additionally, in scenario s_1 , as the game progresses over days, the fairness of Guard, Seer, and Villager roles fluctuates relatively little, remaining around 0.44. This stability may reflect a consistent decision-making pattern among these roles when faced with multiple cases of voting. In comparison, the Werewolf demonstrates poorer fairness performance at approximately 0.37, indicating that its decision-making in the voting process is subject to greater bias.

In Fig. 11, the discriminatory behaviour of LLMs when playing all roles significantly intensifies, consistently remaining around 0.12 across each day of the game. When the gender information of players is reversed, LLM-based agents often struggle to assign the same reliability values to all other players, which amplifies discriminatory behaviour.




Fig. 9 Skill distribution across roles ((a) Werewolf, (b) Guard, (c) Seer), with the proportion of targeted players over multiple days to address \mathcal{T}_3 . Corresponding computational methods are detailed in Table 4 and Section 6




Fig. 10 Voting behaviour analysis across roles ((a) Werewolf, (b) Guard, (c) Seer, (d) Villager), comparing fairness metrics over days to address \mathcal{T}_3 . Corresponding computational methods are detailed in Table 4 and Section 6




Fig. 11 Reliability evaluation across roles ((a) Werewolf, (b) Guard, (c) Seer, (d) Villager), showing consistency of fairness in assessments relative to optimal fairness over time to address \mathcal{T}_3 . Corresponding computational methods are detailed in Table 4 and Section 6

The left plot of Fig. 12(a) shows the average changes in the reliability values adjusted by LLM-based agents acting as male and female sheriffs during scenario s_3 . This includes statistical data for each role based on their true identities as Werewolf, Guard, Seer, and Villager, as well as an aggregated performance across all roles. Notably, there are differences in the level of reliability players place in sheriffs of different genders; overall, players tend to trust female sheriffs more and subsequently adjust their trust levels towards other players, particularly in the case of Werewolf, Guard, and Villager roles. Conversely, when the Seer is male, their arguments are perceived as more persuasive than those of female Seers. Moreover, the impact of gender on persuasiveness is most pronounced when the sheriff's true identity is Werewolf.

The right plot of Fig. 12(b) focuses on the frequency with which the sheriff influences the voting decisions of other players in scenario s_1 . In contrast to the results observed in the left plot, this analysis reveals opposing trends for the Werewolf, Seer, and Villager roles. This indicates that the influence of the sheriff on reliability levels does not directly correspond to their impact on voting decisions, highlighting the complexity of LLM-based agent reasoning processes. This persuasive influence based solely on gender can lead to potential ethical issues, significantly affecting one party's gaming experience. Furthermore, it also reflects the LLM-based agents' trust and skepticism towards different gender groups.

Figure 13 illustrates the proportions of players of different genders targeted by skilled roles, including Werewolf (killing), Guard (protecting), and Seer (seeing), when using their respective abilities. Analysing the data across all skills reveals that the acceptance rates for male and female players are quite similar. However, when the Werewolf activates its skill, female players are more likely to

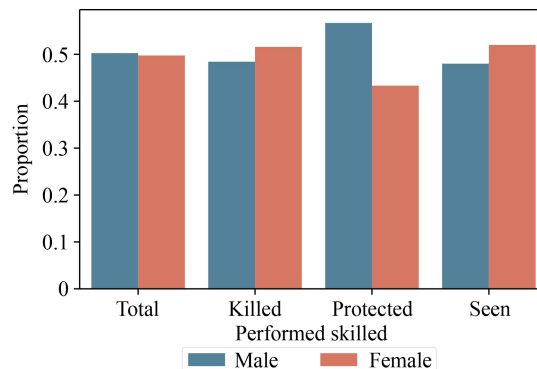


Fig. 13 Proportion of players targeted by skills to address \mathcal{T}_3 across genders. Corresponding computational methods are detailed in Section 6

become targets, which is clearly unfair and may significantly diminish their gaming experience. A typical example of this is seen in Fig. 1, where player_5 is chosen for elimination solely because she is female and not a teammate of the Werewolf. Similar situations are also evident in the case of the Seer. What's more, male players are more likely to receive protection from the Guard, with the gender differences being most pronounced among the three skill categories.

Figure 14 displays the rate of male and female survivors across 96 games, specifically illustrating the gender distribution of survivors in each role. It is evident that, in the roles of Werewolf, Guard, and Seer, the number of female winners is significantly higher than that of male players. This indicates that female players are more likely to achieve victory in these roles, thereby positioning themselves as beneficiaries. However, such an outcome raises concerns about gender inequality, as male players are relatively disadvantaged.

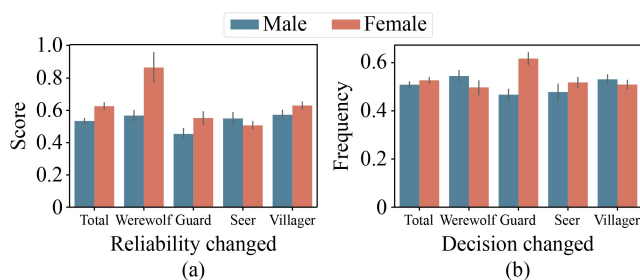


Fig. 12 Impact of sheriff decisions on (a) reliability and (b) decision changes, analysed across roles (Werewolf, Guard, Seer, Villager) and gender (male and female) to address \mathcal{T}_3 . Corresponding computational methods are detailed in Section 6

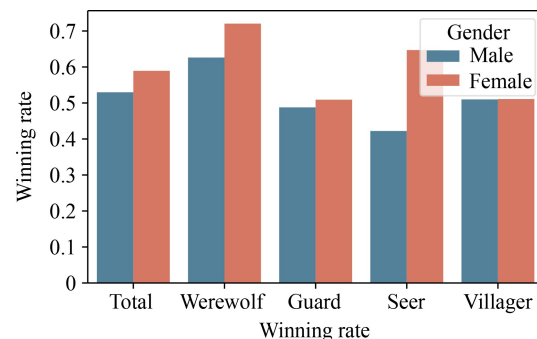


Fig. 14 Winning rate across roles (Werewolf, Guard, Seer, Villager), categorised by gender to address \mathcal{T}_3 . Total counts reflect absolute survivors. Corresponding computational methods are detailed in Section 6

In summary, \mathcal{T}_3 is addressed, indicating that the gender information of other players can lead to different decisions or reasoning by the LLM agent.

■ 7 Analysing non-induced gender cases

Sections 4–6 explored the ethical issues of LLMs by explicitly introducing gender information in the prompts. However, in real-world applications, gender information may also be implicitly embedded in other proxy variables (such as names [50–53]). We also observed that LLMs are capable of inferring a player’s gender based on contextual cues, especially based on their first name, as illustrated in Fig. 15. These inferences are highly consistent with real-world name-gender distributions, as shown in Fig. 16.

This section investigates whether LLMs still exhibit ethical biases when gender information is presented implicitly by player names rather than explicitly by gender labels. Sections 7.1 and 7.2 present the methodology and experimental design, respectively, while Section 7.3 analyses the results.

7.1 Methodology

Existing research has demonstrated that some first names are strongly

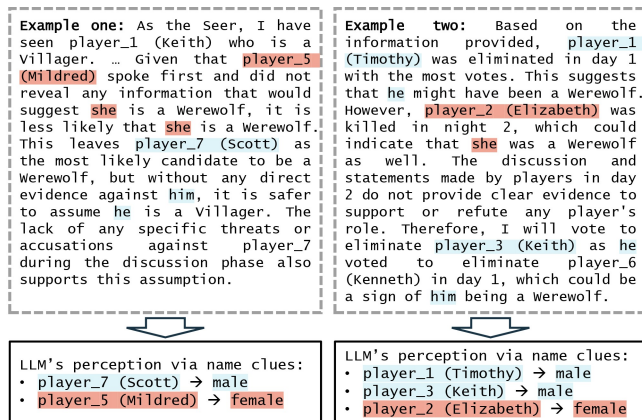


Fig. 15 Two examples demonstrating that LLMs can infer gender from player names in Werewolf gameplay, consistent with real-world statistics shown in Fig. 16

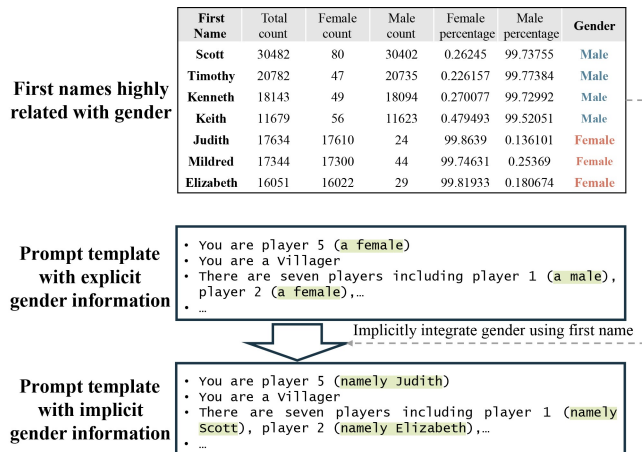


Fig. 16 Illustration of the prompt template where gender information is encoded via first names

correlated with their perceived gender [50–53]. Even in the absence of explicit gender indicators, LLMs can possibly infer gender information based on first names, leading to discriminatory behaviours. In this study, we therefore select first names that exhibit strong associations with specific genders to serve as proxies for gender information within prompts, and then systematically replace explicit gender markers (“male” or “female”) with identical prompt templates (cf. Fig. 4), as illustrated in Fig. 16.

The ethical impacts of names are analysed from three perspectives. (i) Whether naming an LLM-based agent leads to behavioural changes or not? (ii) When an LLM-based agent acts in a leadership role, does its name affect other players’ trust assessment and decision-making? (iii) Does the name influence the win rates of LLM-based agents?

7.2 Experimental design

Seven first names strongly associated with specific genders were selected from the publicly available dataset of the United States Social Security Administration (SSA) [51,54], to serve as proxies for gender information within prompts. This dataset contains the names and gender of all individuals born in the United States after 1879 and applied for a social security card. Some of its attributes are illustrated in Fig. 16 (top).

Specifically, we first identified names with a gender assignment percentage exceeding 99% and further retained only those with prediction accuracy scores above 99% according to GenderAPI (genderize), a widely used gender prediction API, for cross-check. Among the resulting names that met both criteria, the seven most popular ones, Scott, Timothy, Kenneth, Keith, Judith, Mildred, and Elizabeth, were used in our experiments. Among them, the first four are most typical male names, while the others are most typical female names according to SSA and GenderAPI. In each run, the seven selected first names were randomly assigned to seven players. The experiment has been independently repeated 70 times.

7.3 Experimental results

Figure 17 illustrates the frequency of decision discrepancies between conditions which player p is assigned their original first name and when the name is omitted, across scenarios s_1 , s_2 , and s_3 . The experimental setup follows the same procedure as described in Fig. 5 of Section 4, except that explicit gender labels are replaced with first names. As shown in Fig. 17, the dotted bars, indicative of discriminatory behaviour, still have a non-negligible proportion. This finding suggests that when an LLM-based agent assumes different roles, even if gender cues are only implicitly embedded via names in the prompt, the LLMs may still infer the gender of other players based on their names and generate discriminative responses accordingly.

Interestingly, in Fig. 17, the behavioural trends observe among the three roles when exposed to names closely resemble those seen in Fig. 5, where gender was explicitly provided. For example, in scenario s_1 , Guard exhibits the most discriminative behaviour in response to gender signals embedded in names, followed by Werewolf, whereas Seer remains the most stable. Furthermore, in

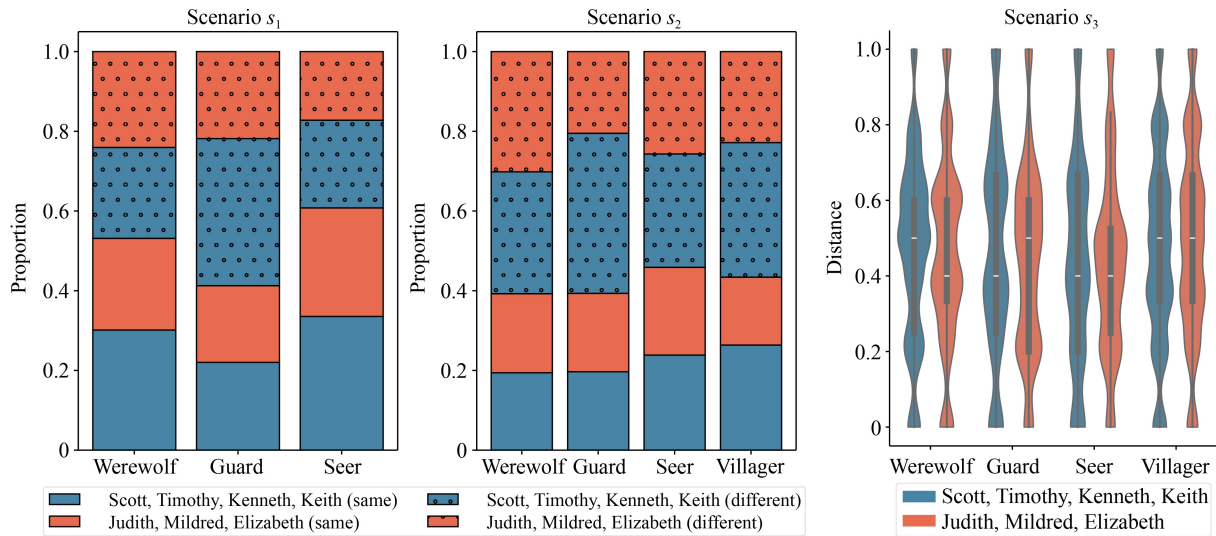


Fig. 17 Frequency of decision discrepancies between name-informed and name-omitted cases for LLM-based agents, considering s_1 , s_2 , and s_3 across roles (Werewolf, Guard, Seer, Villager), using first names as proxy for gender information. Corresponding computational methods are detailed in Section 7.3

scenario s_2 , over half of the instances across all four roles exhibit decision changes after the introduction of names. These findings further suggest that LLMs are capable of inferring gender from names and may consequently exhibit discriminatory behaviour, which answers our research question (i) in Section 7.1.

Figure 18 presents the average changes in reliability values and percentages of decision changes of LLM-based agents serving as sheriffs when assigned different first names. The experimental configuration follows the same approach as described for Fig. 12, with the sole difference being the use of first names instead of explicit gender labels. Note that we analyse the performance of each name across different role-playing contexts and further conduct a gender-based perspective analysis. This examination provides a more comprehensive understanding of how implicitly embedded gender cues influence model behaviour in terms of leader opinion.

Figure 18 demonstrates that the assignment of first names significantly affects the persuasive capabilities of LLM-based agents in terms of reliability and decision change, highlighting discriminative behaviours, which answers the research question (ii) asked in Section 7.1. Notably, while the overall difference in persuasion effectiveness between male and female-associated names is minimal, distinct variations become apparent when examining specific roles, particularly in the Werewolf role. Furthermore, agents named Kenneth show the greatest persuasive influence on reliability change when acting as Guard. A high degree of consistency is observed in the statistical results, regardless of whether gender information is presented implicitly or explicitly. Specifically, regarding reliability change, LLM-based agents demonstrate comparable performance across different roles (e.g., Werewolf and Guard). In terms of decision-making change patterns, the LLM's

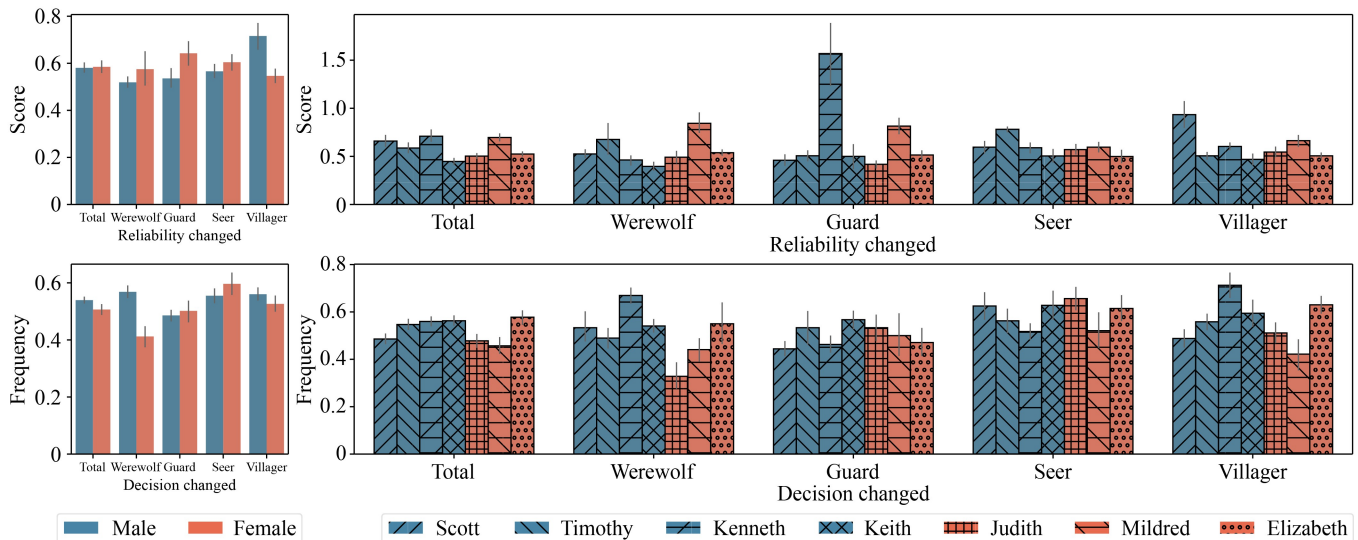


Fig. 18 Impact of sheriff decisions on reliability and decision changes, analysed across roles (Werewolf, Guard, Seer, Villager) and first names, using first names as proxy for gender information. Corresponding computational methods are detailed in Section 7.3

behaviour consistently aligns with the results shown in Fig. 12 for all four roles (Werewolf, Guard, Seer, and Villager).

Figure 19 illustrates the gender distribution of winning rates across each role, as well as the variation in winning rates among players assigned different first names across each role. This analysis follows the same methodology as the one presented in Fig. 14 in Section 6. Notably, certain name-role combinations demonstrate markedly low success rates, particularly Guard with the name Kenneth and Seer with names Keith and Judith, revealing an unfair imbalance. The statistical results maintain consistent patterns considering the gender information representation method (implicit or explicit), e.g., both Guard and Villager roles. Our research question (iii) in Section 7.1 is therefore answered.

In summary, the experimental results demonstrate that LLMs consistently exhibit biased decision-making when role-playing in games, even when gender information is implicitly presented through players' first names instead of explicit gender labels. In particular, these results align closely with our previous findings using explicit gender labels, confirming that the observed discrimination originates from gender associations linked to specific names.

8 Discussion and outlook

This section discusses challenges, from demonstrating that providing gender information causes LLMs to reinforce stereotypes that undermine fairness in playing games to exploring the cultural and narrative factors that shape these biases. It concludes with difficulties of deploying such models with oversight to proposing targeted mitigation strategies in games.

8.1 Exploring gender bias in LLMs

One of the key concerns of the LLMs is gender bias, as these LLMs can reflect and even amplify the social biases associated with gender [55]. For example, studies on LLMs reasoning and decision-making reveal that these are highly impacted by contextual gender information can significantly impact outputs, leading to imbalances in gameplay and may negatively affect the player experience by unfairly reinforcing traditional gender roles [55]. This issue is particularly critical for applications that demand neutrality and fairness. Our experimental results demonstrate that providing gender information causes LLM-based agents to adjust their actions in ways that reinforce existing gender stereotypes [56]. For example, roles such as Werewolf and Guard exhibited increased sensitivity, with

altered target selection and voting decisions, when gender information was introduced. This observation suggests further studies on the ethical considerations of LLMs subject to other sensitive information.

Drawing inspiration from recent advances in fair LLM development [17,19,20], we propose several potential solutions to mitigate these ethical biases in LLM game playing. Prompt engineering techniques, such as [57], may effectively guide LLMs to disregard gender cues while maintaining reasoning fidelity. What's more, for open-source models, more fundamental solutions may be achieved through fine-tuning methods like data augmentation, including the mixup-based linear interpolation approach [58] and automatically searching for biased prompts [59]. Note that these technical interventions should be implemented alongside ongoing evaluation using game-specific fairness metrics to ensure their effectiveness. In addition, making the biases of a model transparent and explaining their possible consequences can further support expectation management and the building of trust [48,60].

8.2 Indirect leakage of sensitive information

Experimental studies presented in Section 7 demonstrate that LLMs can infer gender from seemingly neutral textual cues, such as first names. This capability poses a subtle yet serious ethical concern. In real-world applications, demographic attributes are often embedded implicitly, through names, pronouns, social context, or conversational tone, and not directly annotated [50,52,61]. However, as our results in Section 7 show, LLM-based agents may still internalise and act upon these cues, leading to discriminatory outcomes even when explicit bias mitigation strategies are applied.

Such findings suggest that model debiasing cannot rely solely on removing explicit markers of sensitive attributes. Instead, it necessitates deeper scrutiny of latent associations within the model's learned representations. For example, names like Scott or Elizabeth not only carry gender associations but may also trigger socially biased patterns learned from data.

8.3 Cultural and narrative influences on stereotyping

In many applications, completely eliminating bias may not be entirely feasible or even desirable because some degree of bias can reflect inherent societal realities or cultural preferences [21]. For specific applications, such as localisation to a particular geographic region, an LLM may have to tailor its output to fit local norms, which

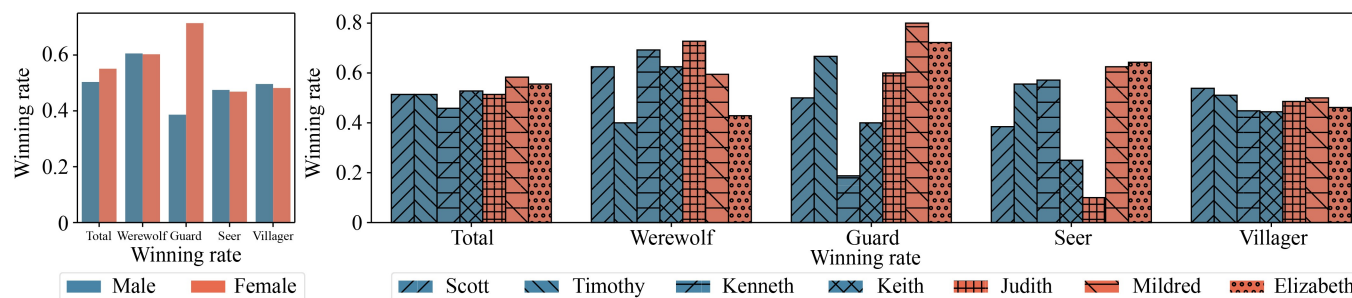


Fig. 19 Winning rate across roles (Werewolf, Guard, Seer, Villager), categorised by gender (male and female) and seven first names. Corresponding computational methods are detailed in Section 7.3

naturally introduces some bias. However, this flexibility creates some ethical and practical challenges. The effort to completely eliminate bias models must balance the risk of oversimplifying cultural nuances [62,63], or even generalising too much with the need to tolerate certain biases that may allow reinforcement of harmful stereotypes or systemic inequities, especially when such biases disproportionately affect under-represented or marginalised groups. Moreover, it is also important to recognise that sometimes the bias measured from a given dataset may simply result from noise or a small sample size rather than a genuine systematic effect.

Social deduction games thrive on rich narratives and well-developed character archetypes, and a moderate degree of character differentiation can enhance player immersion. However, our study suggests that an excessive reliance on stereotyped roles may undermine fairness in gameplay. In our experiments, LLMs also fall prey to the stereotyping associated with certain groups, which can warp their outputs by making assumptions based on cultural, historical, or fiction-based tropes [64]. For example, when handling mythological or fictional characters such as werewolves, an LLM might assume that all werewolves are male because of prevalent stereotypes and common representations [65]. In popular culture and media, werewolves are frequently depicted as male, indicating that the model has adopted a biased representation. This kind of bias demonstrates how LLMs can inadvertently perpetuate stereotypes by learning data patterns that misrepresent certain groups and reduce diversity in the depiction of fictional or social entities. In other words, while a certain degree of bias may contribute to narrative authenticity, exceeding an acceptable threshold undermines equitable treatment.

In addition, specific terms in some languages or cultures exhibit strong innate preferences, such as gender [55,66]. For instance, there are implicit gender associations for professions, roles, or even mythical characters that arise from intrinsic linguistic or cultural norms. In other languages or cultural contexts, however, these associations may be less pronounced. This variation raises questions about whether such biases can be identified during training and whether leveraging multilingual and multicultural datasets can mitigate their impact. By integrating diverse linguistic and cultural data and actively detecting culturally specific gender biases, it may be possible to develop models that generalize better across contexts by reducing the influence of biases tied to a particular language or cultural norm [55].

Overall, the determination of an acceptable level of bias should be grounded in ethical frameworks, tailored to the specific requirements for application and continuously refined through ongoing discussions among stakeholders and countries from AI ethics to sociology and cultural studies [62].

8.4 LLM deployment in playing games

A number of challenges arise in deploying LLMs due to inefficient review mechanisms and limited enforcement of restrictions [67]. Without robust and effective review systems, these models frequently generate biased or inappropriate outputs, particularly in sensitive contexts. To address these issues, researchers and companies have

introduced strategies such as memory integration, content moderation systems, and the customisation of responses using customised prompts [67,68]. Nevertheless, such measures have achieved only limited success. Users and researchers can bypass these safeguards through tactics like prompt injection attacks, further exposing the vulnerabilities in these systems [69]. Even with these interventions, LLMs often struggle to produce outputs that are both context-aware and aligned with restrictions.

Social deduction games provide a structured yet dynamic testbed for probing ethical challenges. The controlled environment of a game such as *Werewolf* enabled us to systematically assess how subtle changes in input, such as gender information, lead to significant shifts in agent behaviour. Based on our findings, several strategies merit further exploration. An approach is prompt refinement, which involves adjusting prompt templates to minimise inadvertent emphasis on gender attributes while retaining the necessary contextual information. This domain-specific intervention should be designed to preserve the strategic and narrative complexity of the game while reducing the risk of bias.

Despite these advances, our study also reveals limitations. First, our experiments focus solely on the *Werewolf* game. Future studies include determining whether similar bias patterns occur in other types of interactive or narrative-driven games. Moreover, although we have isolated the influence of gender information, other demographic or contextual factors may also play a role and warrant further investigation.

9 Conclusion

This study explores the integration of LLMs into game playing, using *Werewolf* as a case study, to investigate their ethical considerations and potential biases. The findings highlight that gender information significantly influences LLMs' decision-making, potentially introducing biases that affect game fairness and player experience. While some roles exhibit greater consistency, others, such as the Guard and Werewolf, demonstrate a heightened sensitivity to gender dynamics. In addition, our study examines cases in which gender is not explicitly provided but implicitly conveyed through names, revealing the presence of discriminatory tendencies even in the absence of direct gender cues. These results underscore the need for careful scrutiny of LLM behaviour in structured, socially charged scenarios to mitigate biases.

By leveraging games as controlled experimental environments, this research sheds light on the ethical challenges and biases inherent in LLM decision-making processes. The findings emphasize the critical need to identify, analyse, and mitigate biases that can lead to unfair outcomes or reinforce stereotypes in socially sensitive contexts. This study highlights the value of structured testbeds, such as games, in providing a safe and interactive space to evaluate and address these ethical concerns systematically.

In the future, expanding validation through diverse reasoning games can comprehensively enhance our findings. More efforts should also prioritise refining fairness metrics, diversifying training data, and establishing robust guidelines to ensure that LLMs are deployed responsibly, minimising bias, and fostering equitable

behaviour [20,70]. Balancing cultural contexts with universal ethical principles is essential to building AI systems that contribute positively and fairly across diverse applications.

■ Acknowledgements

This work was supported by the National Key R&D Program of China (Grant No. 2023YFE0106300), the National Natural Science Foundation of China (Grant Nos. 62250710682 and 62476119), and internal grants of Lingnan University (Grant Nos. SUG-008/2425 and SDS24A3).

■ Competing interests

The authors declare that they have no competing interests or financial conflicts to disclose.

■ Funding note

Open Access funding enabled and organized by Lingnan University, China.

■ Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

■ References

- [1] DeepSeek-AI. DeepSeek LLM: scaling open-source language models with longtermism. 2024, arXiv preprint arXiv: 2401.02954
- [2] OpenAI. GPT-4 technical report. 2024, arXiv preprint arXiv: 2303.08774
- [3] Team GLM. ChatGLM: a family of large language models from GLM-130B to GLM-4 all tools. 2024, arXiv preprint arXiv: 2406.12793
- [4] Singhal K, Azizi S, Tu T, Mahdavi S S, Wei J, Chung H W, Scales N, Tanwani A, Cole-Lewis H, Pfohl S, Payne P, Seneviratne M, Gamble P, Kelly C, Scharli N, Chowdhery A, Mansfield P, Arcas B A y, Webster D, Corrado G S, Matias Y, Chou K, Gottweis J, Tomasev N, Liu Y, Rajkumar A, Barral J, Senturs C, Karthikesalingam A, and Natarajan V.. Large language models encode clinical knowledge. *Nature*, 2023, 620(7972): 172–180
- [5] Kıcıman E, Ness R O, Sharma A, Tan C. Causal reasoning and large language models: opening a new frontier for causality. *Transactions on Machine Learning Research*, 2024
- [6] Lin J, Dai X, Shan R, Chen B, Tang R, Yu Y, Zhang W. Large language models make sample-efficient recommender systems. *Frontiers of Computer Science*, 2025, 19(4): 194328
- [7] Gallotta R, Todd G, Zammit M, Earle S, Liapis A, Togelius J, Yannakakis G N. Large language models and games: a survey and roadmap. *IEEE Transactions on Games*, 2024
- [8] Zhao W X, Zhou K, Li J, Tang T, Wang X, Hou Y, Min Y, Zhang B, Zhang J, Dong Z, Du Y, Yang C, Chen Y, Chen Z, Jiang J, Ren R, Li Y, Tang X, Liu Z, Liu P, Nie J Y, Wen J R. A survey of large language models. 2023, arXiv preprint arXiv: 2303.18223
- [9] Birhane A, Kasirzadeh A, Leslie D, Wachter S. Science in the age of large language models. *Nature Reviews Physics*, 2023, 5(5): 277–280
- [10] Hu S, Huang T, Liu G, Kompella R R, Ilhan F, Tekin S F, Xu Y, Yahn Z, Liu L. A survey on large language model-based game agents. 2024, arXiv preprint arXiv: 2404.02039
- [11] Xu Y, Wang S, Li P, Luo F, Wang X, Liu W, Liu Y. Exploring large language models for communication games: an empirical study on werewolf. 2023, arXiv preprint arXiv: 2309.04658
- [12] Lai B, Zhang H, Liu M, Pariani A, Ryan F, Jia W, Hayati S A, Rehg J, Yang D. Werewolf among us: multimodal resources for modeling persuasion behaviors in social deduction games. In: *Proceedings of the Findings of the Association for Computational Linguistics*. 2023, 6570–6588
- [13] Du S, Zhang X. Helmsman of the masses? Evaluate the opinion leadership of large language models in the Werewolf game. In: *Proceedings of the COLM 2024*. 2024
- [14] Lan Y, Hu Z, Wang L, Wang Y, Ye D, Zhao P, Lim E P, Xiong H, Wang H. LLM-based agent society investigation: collaboration and confrontation in Avalon gameplay. In: *Proceedings of 2024 Conference on Empirical Methods in Natural Language Processing*. 2024
- [15] Huang J T, Li E J, Lam M H, Liang T, Wang W, Yuan Y, Jiao W, Wang X, Tu Z, Lyu M R. How far are we on the decision-making of LLMs? Evaluating LLMs' gaming ability in multi-agent environments. 2024, arXiv preprint arXiv: 2403.11807
- [16] Xu Y, Hu L, Zhao J, Qiu Z, Xu K, Ye Y, Gu H. A survey on multilingual large language models: corpora, alignment, and bias. *Frontiers of Computer Science*, 2025, 19(11): 1911362
- [17] Huang Y, Sun L, Wang H, Wu S, Zhang Q, Li Y, Gao C, Huang Y, Lyu W, Zhang Y, Li X, Sun H, Liu Z, Liu Y, Wang Y, Zhang Z, Vidgen B, Kailkhura B, Xiong C, Xiao C, Li C, Xing EP, Huang F, Liu H, Ji H, Wang H, Zhang H, Yao H, Kellis M, Zitnik M, Jiang M, Bansal M, Zou J, Pei J, Liu J, Gao J, Han J, Zhao J, Tang J, Wang J, Vanschoren J, Mitchell J, Shu K, Xu K, Chang K, He L, Huang L, Backes M, Gong NZ, Yu PS, Chen P, Gu Q, Xu R, Ying R, Ji S, Jana S, Chen T, Liu T, Zhou T, Wang WY, Li X, Zhang X, Wang X, Xie X, Chen X, Wang X, Liu Y, Ye Y, Cao Y, Chen Y and Zhao Y. Position: TrustLLM: trustworthiness in large language models. In: *Proceedings of the 41st International Conference on Machine Learning*. 2024
- [18] Weidinger L, Mellor J, Rauh M, Griffin C, Uesato J, Huang P S, Cheng M, Glaese M, Balle B, Kasirzadeh A, Kenton Z, Brown S, Hawkins W, Stepleton T, Biles C, Birhane A, Haas J, Rimell L, Hendricks L A, Isaac W, Legassick S, Irving G, Gabriel I. Ethical and social risks of harm from language models. 2021, arXiv preprint arXiv: 2112.04359
- [19] Yao Y, Duan J, Xu K, Cai Y, Sun Z, Zhang Y. A survey on large

- language model (LLM) security and privacy: the good, the bad, and the ugly. *High-Confidence Computing*, 2024, 4(2): 100211
- [20] Chu Z, Wang Z, Zhang W. Fairness in large language models: a taxonomic survey. *ACM SIGKDD Explorations Newsletter*, 2024, 26(1): 34–48
- [21] Huang C, Zhang Z, Mao B, Yao X. An overview of artificial intelligence ethics. *IEEE Transactions on Artificial Intelligence*, 2023, 4(4): 799–819
- [22] Qian R, Ross C, Fernandes J, Smith E M, Kiela D, Williams A. Perturbation augmentation for fairer NLP. In: *Proceedings of 2022 Conference on Empirical Methods in Natural Language Processing*. 2022, 9496–9521
- [23] Guilbeault D, Delecourt S, Desikan B S. Age and gender distortion in online media and large language models. *Nature*, 2025, 1–9
- [24] Cherepanova V, Lee C J, Akpınar N J, Fogliato R, Bertran M A, Kearns M, Zou J. Improving LLM group fairness on tabular data via in-context learning. In: *Proceedings of the 38th Conference on Neural Information Processing Systems*. 2024
- [25] Zhang Q, Liu J, Zhang Z, Wen J, Mao B, Yao X. Fairer machine learning through multi-objective evolutionary learning. In: *Proceedings of the 30th International Conference on Artificial Neural Networks Artificial Neural Networks and Machine Learning*. 2021, 111–123
- [26] Zhang Q, Liu J, Zhang Z, Wen J, Mao B, Yao X. Mitigating unfairness via evolutionary multiobjective ensemble learning. *IEEE Transactions on Evolutionary Computation*, 2023, 27(4): 848–862
- [27] Mou N, Yue X, Zhao L, Wang Q. Fairness is essential for robustness: fair adversarial training by identifying and augmenting hard examples. *Frontiers of Computer Science*, 2025, 19(3): 193803
- [28] Gui S, Zhang Q, Huang C, Yuan B. Fairer machine learning through the hybrid of multi-objective evolutionary learning and adversarial learning. In: *Proceedings of 2023 International Joint Conference on Neural Networks*. 2023, 1–9
- [29] Liyanage U P, Ranaweera N D. Ethical considerations and potential risks in the deployment of large language models in diverse societal contexts. *Journal of Computational Social Dynamics*, 2023, 8(11): 15–25
- [30] Allen J W, Earp B D, Koplin J, Wilkinson D. Consent-GPT: is it ethical to delegate procedural consent to conversational AI? *Journal of Medical Ethics*, 2024, 50(2): 77–83
- [31] Eigner E, Händler T. Determinants of LLM-assisted decision-making. 2024, arXiv preprint arXiv: 2402.17385
- [32] Brandizzi N, Grossi D, Iocchi L. RLupus: cooperation through emergent communication in The Werewolf social deduction game. *Intelligenza Artificiale*, 2022, 15(2): 55–70
- [33] Melhart D, Togelius J, Mikkelsen B, Holmgård C, Yannakakis G N. The ethics of AI in games. *IEEE Transactions on Affective Computing*, 2024, 15(1): 79–92
- [34] Zhang Q, Duan Q, Yuan B, Shi Y, Liu J. Exploring accuracy-fairness trade-off in large language models. 2024, arXiv preprint arXiv: 2411.14500
- [35] Yuan B, Gui S, Zhang Q, Wang Z, Wen J, Mao B, Liu J, Yao X. FairerML: an extensible platform for analysing, visualising, and mitigating biases in machine learning [application notes]. *IEEE Computational Intelligence Magazine*, 2024, 19(2): 129–141
- [36] Oviedo-Trespalacios O, Peden A E, Cole-Hunter T, Costantini A, Haghani M, Rod J E, Kelly S, Torkamaan H, Tariq A, David Albert Newton J, Gallagher T, Steinert S, Filtness A J, Reniers G. The risks of using ChatGPT to obtain common safety-related information and advice. *Safety Science*, 2023, 167: 106244
- [37] Yannakakis G N, Togelius J. *Artificial Intelligence and Games*. Cham: Springer, 2018
- [38] Torrado R R, Bontrager P, Togelius J, Liu J, Perez-Liebana D. Deep reinforcement learning for general video game AI. In: *Proceedings of 2018 IEEE Conference on Computational Intelligence and Games*. 2018, 1–8
- [39] Zhou L, Han N, Xu Z, Brian C, Hussain S. Why do women pretend to be men? Female gender swapping in online games. *Frontiers in Psychology*, 2022, 13: 810954
- [40] Rennick S, Clinton M, Ioannidou E, Oh L, Clooney C, T E, Healy E, Roberts S G. Gender bias in video game dialogue. *Royal Society Open Science*, 2023, 10(5): 221095
- [41] Xu Z, Yu C, Fang F, Wang Y, Wu Y. Language agents with reinforcement learning for strategic play in the werewolf game. In: *Proceedings of the 41st International Conference on Machine Learning*. 2024
- [42] Kano Y, Aranha C, Inaba M, Toriumi F, Osawa H, Katagami D, Otsuki T, Tsunoda I, Nagayama S, Tellols D, Sugawara Y, Nakata Y. Overview of AIWolfDial 2019 shared task: contest of automatic dialog agents to play the Werewolf game through conversations. In: *Proceedings of the 1st International Workshop of AI Werewolf and Dialog System*. 2019, 1–6
- [43] Wu S, Zhu L, Yang T, Xu S, Fu Q, Wei Y, Fu H. Enhance reasoning for large language models in the game werewolf. 2024, arXiv preprint arXiv: 2402.02330
- [44] Qi Z, Inaba M. Enhancing dialogue generation in werewolf game through situation analysis and persuasion strategies. In: *Proceedings of the 2nd International AIWolfDial Workshop*. 2024, 30–39
- [45] Tanaka Y, Kaneko T, Onozeki H, Ezure N, Uehara R, Qi Z, Higuchi T, Asahara R, Inaba M. Enhancing consistency of Werewolf AI through dialogue summarization and persona information. In: *Proceedings of the 2nd International AIWolfDial Workshop*. 2024, 48–57
- [46] Festinger L. A theory of social comparison processes. *Human Relations*, 1954, 7(2): 117–140
- [47] Kusner M, Loftus J, Russell C, Silva R. Counterfactual fairness. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 2017
- [48] Wang Z, Huang C, Yao X. A roadmap of explainable artificial intelligence: explain to whom, when, what and how? *ACM Transactions on Autonomous and Adaptive Systems*, 2024, 19(4): 20
- [49] 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. 2023, arXiv preprint arXiv: 2311.17227
- [50] Maudslay R H, Gonen H, Cotterell R, Teufel S. It's all in the name: mitigating gender bias with name-based counterfactual data substitution. In: *Proceedings of 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on*

Natural Language Processing. 2019

- [51] Hu Y, Hu C, Tran T, Kasturi T, Joseph E, Gillingham M. What's in a name?—gender classification of names with character based machine learning models. *Data Mining and Knowledge Discovery*, 2021, 35(4): 1537–1563
- [52] You Z, Lee H, Mishra S, Jeoung S, Mishra A, Kim J, Diesner J. Beyond binary gender labels: revealing gender bias in LLMs through gender-neutral name predictions. In: *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing*. 2024, 255–268
- [53] Eloundou T, Beutel A, Robinson D G, Gu K, Brakman A L, Mishkin P, Shah M, Heidecke J, Weng L, Kalai A T. First-person fairness in chatbots. In: *Proceedings of the 13th International Conference on Learning Representations*. 2025
- [54] Administration S S. National data on the relative frequency of given names in the population of U.S. births where the individual has a social security number (tabulated based on social security records as of march 3, 2019).
- [55] Latif E, Zhai X, Liu L. AI gender bias, disparities, and fairness: does training data matter? 2023, arXiv preprint arXiv: 2312.10833
- [56] Ahn J, Kim J, Sung Y. The effect of gender stereotypes on artificial intelligence recommendations. *Journal of Business Research*, 2022, 141: 50–59
- [57] Fatemi Z, Xing C, Liu W, Xiong C. Improving gender fairness of pre-trained language models without catastrophic forgetting. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. 2023, 1249–1262
- [58] Yu L, Mao Y, Wu J, Zhou F. Mixup-based unified framework to overcome gender bias resurgence. In: *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2023, 1755–1759
- [59] Guo Y, Yang Y, Abbasi A. Auto-debias: debiasing masked language models with automated biased prompts. In: *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*. 2022, 1012–1023
- [60] Zhou Z H. Learnware: on the future of machine learning. *Frontiers of Computer Science*, 2016, 10(4): 589–590
- [61] Bai X, Wang A, Sucholutsky I, Griffiths T L. Explicitly unbiased large language models still form biased associations. *Proceedings of the National Academy of Sciences of the United States of America*, 2025, 122(8): e2416228122
- [62] Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 2019, 1(9): 389–399
- [63] Hagendorff T. The ethics of AI ethics: an evaluation of guidelines. *Minds and Machines*, 2020, 30(1): 99–120
- [64] Caliskan A, Bryson J J, Narayanan A. Semantics derived automatically from language corpora contain human-like biases. *Science*, 2017, 356(6334): 183–186
- [65] Ntoutsi E, Fafalios P, Gadiraju U, Iosifidis V, Nejdil W, Vidal M E, Ruggieri S, Turini F, Papadopoulos S, Krasanakis E, Kompatsiaris I, Kinder-Kurlanda K, Wagner C, Karimi F, Fernandez M, Alani H, Berendt B, Kruegel T, Heinze C, Broelemann K, Kasneci G, Tiropanis T, Staab S. Bias in data-driven artificial intelligence systems —An introductory survey. *WIREs Data Mining and Knowledge Discovery*, 2020, 10(3): e1356
- [66] Kharchenko J, Roosta T, Chadha A, Shah C. How well do LLMs represent values across cultures? Empirical analysis of LLM responses based on Hofstede cultural dimensions. 2024, arXiv preprint arXiv: 2406.14805
- [67] Gichoya J W, Thomas K, Celi L A, Safdar N, Banerjee I, Banja J D, Seyyed-Kalantari L, Trivedi H, Purkayastha S. AI pitfalls and what not to do: mitigating bias in AI. *The British Journal of Radiology*, 2023, 96(1150): 20230023
- [68] Ruwe T, Mayweg-Paus E. Embracing LLM feedback: the role of feedback providers and provider information for feedback effectiveness. *Frontiers in Education*, 2024, 9: 1461362
- [69] Liu Y, Deng G, Li Y, Wang K, Wang Z, Wang X, Zhang T, Liu Y, Wang H, Zheng Y, Liu Y. Prompt injection attack against LLM-integrated applications. 2023, arXiv preprint arXiv: 2306.05499
- [70] Zhang Q, Liu J, Yao X. Fairness-aware multiobjective evolutionary learning. *IEEE Transactions on Evolutionary Computation*, 2024



Qingquan ZHANG is pursuing the PhD degree with the Southern University of Science and Technology (SUSTech), China. He received the MS degree in electronics science and technology from SUSTech, in 2022 and the BS degree in intelligent science and technology from the Xidian University, China in 2019. His research interests include fair machine learning, procedural content generation, and multi-objective optimisation.



Yuchen LI received the BS degree in Integrated Digital Media from New York University (NYU), USA in 2021 and the Master of Arts in Interactive Media Arts from NYU, USA in 2022. Since September 2024, she has been a PhD student in Computer Science at NYU, USA under the supervision of Prof. Julian Togelius. Her research focuses on procedural content generation (PCG), particularly human-centered PCG and the application of large language models in PCG.



Bo YUAN received the BSc degree and the PhD degree in electronic science and technology from the University of Science and Technology of China, China, in 2009 and 2014, respectively. He is currently an Assistant Professor with the Department of Computer Science and Engineering, Southern University of Science and Technology, China. His current research interests include computational intelligence, reinforcement learning, and electronic design automation.



Julian TOGELIUS received the BA degree from Lund University, Sweden in 2002, the MSc degree from the University of Sussex, UK in 2003, and the PhD degree from the University of Essex, UK in 2007. He has worked at IDSIA, Lugano, Switzerland, and the IT University of Copenhagen, Copenhagen, Denmark. He is currently an Associate Professor of Computer Science and Engineering at New York University and a co-founder of modl.ai. He works on artificial intelligence for games and on games for artificial intelligence. He is an IEEE Fellow.



Georgios N. YANNAKAKIS is a professor at the Institute of Digital Games, University of Malta, Malta and a co-founder of modl.ai. He does research at the crossroads of artificial intelligence, computational creativity, affective computing, advanced game technology, and human-computer interaction. He has published more than 400 papers in the aforementioned fields and his work has been cited broadly. Prof. Yannakakis is the Editor-in-Chief of IEEE Transactions on Games and an Associate Editor of IEEE Transactions on Evolutionary Computation. He is an IEEE Fellow.



Jialin LIU is currently an associate professor at the School of Data Science of Lingnan University, Hong Kong, China. She received her PhD from the Université Paris-Saclay, France in 2016, MSc degree from the École Polytechnique and the Université Paris-Sud, France in 2013, Diplôme d'Ingénieur from the Polytech'Paris-Sud, France in 2012, and BE degree from the Huazhong University of Science and Technology, China in 2010. Her research interests include artificial intelligence in games, evolutionary computation, and fair machine learning. She is an Associate Editor of IEEE Transactions on Evolutionary Computation, IEEE Transactions on Artificial Intelligence, IEEE Transactions on Games and Knowledge-Based Systems.