

# INTERINTENT: Investigating Social Intelligence of LLMs via Intention Understanding in an Interactive Game Context

Ziyi Liu\* Abhishek Anand\* Pei Zhou Jen-tse Huang Jieyu Zhao

University of Southern California

{zliu2803, anandabh, peiz, jh\_116, jieyuz}@usc.edu

## Abstract

Large language models (LLMs) have demonstrated the potential to mimic human social intelligence. However, most studies focus on simplistic and static self-report or performance-based tests, which limits the depth and validity of the analysis. In this paper, we developed a novel framework, **INTERINTENT**, to assess LLMs’ social intelligence by mapping their ability to understand and manage intentions in a game setting. We focus on four dimensions of social intelligence: situational awareness, self-regulation, self-awareness, and theory of mind. Each dimension is linked to a specific game task: intention selection, intention following, intention summarization, and intention guessing. Our findings indicate that while LLMs exhibit high proficiency in selecting intentions, achieving an accuracy of 88%, their ability to infer the intentions of others is significantly weaker, trailing human performance by 20%. Additionally, game performance correlates with intention understanding, highlighting the importance of the four components towards success in this game. These findings underline the crucial role of intention understanding in evaluating LLMs’ social intelligence and highlight the potential of using social deduction games as a complex testbed to enhance LLM evaluation. **INTERINTENT** contributes a structured approach to bridging the evaluation gap in social intelligence within multiplayer games.<sup>1</sup>

## 1 Introduction

The growing intelligence of large language models (LLMs) has facilitated diverse research on their capability to mimic human social intelligence (Ziems et al., 2024; Dubova, 2022; Gweon et al., 2023; Huang et al., 2023; Sap et al., 2022; Wu et al., 2023). *Social intelligence*, the ability to understand and manage one’s own and others’ actions and to

act wisely in social relations (Thorndike, 1920), usually includes four key components (Silvera et al., 2001): (1) *Situational Awareness*: The perception and comprehension of the elements in the environment (Endsley, 1995). (2) *Self-Regulation*: The process of guiding one’s own thoughts, behaviors, and feelings to reach goals (Bandura, 1991). (3) *Self-Awareness*: The understanding of one’s own character, feelings, motives, and desires (Gallup et al., 2003). (4) *Theory of Mind (ToM)*: The knowledge about others’ beliefs, intentions, and thoughts (Baron-Cohen, 1991).

Despite adaptations for interactive social environments (Zhou et al., 2023c,a; Park et al., 2023), most existing evaluations on LLMs (Le et al., 2019; Shapira et al., 2023) focus on straightforward and static daily scenarios without explicitly defined goals for them. To enhance the contexts with complexity and dynamics with clear goals for LLMs, we turn our attention to social deduction games. These games require social interactions among players, providing a more diverse test bed for LLM evaluation (Qiao et al., 2023). Existing studies have successfully utilized various games to analyze LLMs’ social behaviors, such as deception and confrontation (Liang et al., 2023; Ibraheem et al., 2022; Mansbach et al., 2021; Meta et al., 2022; O’Gara, 2023; Xu et al., 2023a,b; Wu et al., 2024a). However, these studies often focus on ad-hoc post-analysis of game performance and overlook systematic measurement of social intelligence, limiting the comprehensive understanding of LLMs’ capabilities in social environments.

In this paper, we leverage one of the representatives in social deduction games, *Avalon* (Light et al., 2023a; Wang et al., 2023), as the context for LLM evaluation. *Avalon* is a social deduction game that relies on conversation, making it an ideal test bed with its goal-driven objectives and complex mechanisms, facilitating interactions among LLMs in a grounded environment. We systematically evaluate

\* Equal contribution.

<sup>1</sup>Code is available at <https://github.com/uscnlp-lime/Inter-Intent>

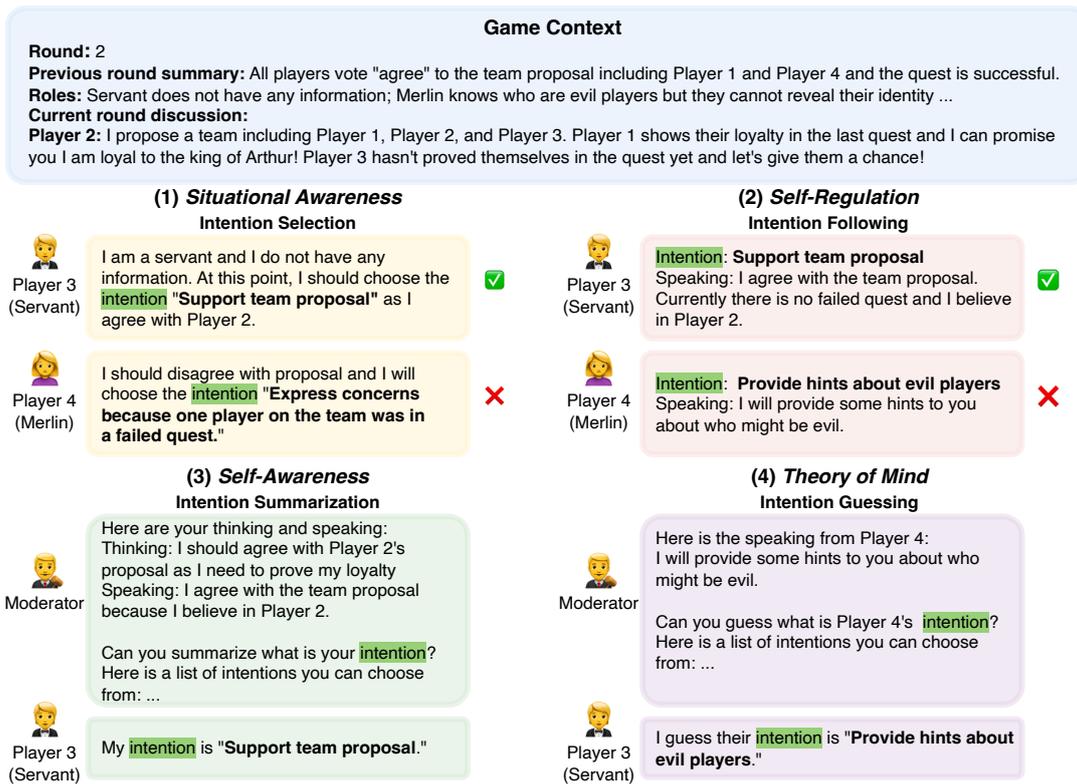


Figure 1: Four dimensions to assess social intelligence in Avalon. We provide a dynamic and complex gaming context for evaluations. For situational awareness, we provide both positive and negative examples. In the negative example, since the previous quest was successful and no player was in a failed quest, the intention is inappropriate as there was no failed quest. For self-regulation, we require models to provide explicit information rather than repeating the intentions. Intentions are in **bold** within the contexts.

LLMs’ social intelligence across the four aforementioned dimensions. Specifically, we propose to focus on players’ *Intentions* (Malle et al., 2001) at various phases of the game, since comprehension of one’s and others’ intentions plays a critical role in games (Goodie et al., 2012). While previous research has explored intention detection (Zhang et al., 2020; Casanueva et al., 2020), the field remains largely underexplored in game contexts. As shown in Figure 1, we design four tests: (1) *Intention Selection*, (2) *Intention Following*, (3) *Intention Summarization*, (4) *Intention Guessing*, respectively for the four social intelligence dimensions.

We first develop a novel Avalon gaming framework enhanced with an intention-guided mechanism, drawing inspiration from previous studies (Lan et al., 2023; Shi et al., 2023; Wang et al., 2023; Light et al., 2023b). Building on this, we introduced **INTERINTENT**, which is designed to dynamically generate contexts for evaluating social intelligence through LLMs’ intention understanding. With human annotation, our results reveal that GPT-3.5 and GPT-4 can select intentions correctly, with accuracies of 87.5% and 88.8%, respectively.

GPT-3.5 achieves a 69.5%  $F_1$  score in intention summarization, with GPT-4 reaching 83.8%, surpassing human performance. Furthermore, we find a strong correlation between high levels of intention comprehension and superior game win rate, highlighting the importance of the intention understanding components in enhancing LLMs’ gaming proficiency. However, the models fall short in intention following, with only 47.5% and 64.8% of responses rated as “Good” in human evaluations, and underperform in intention guessing, trailing human performance by 20%.

Our contributions are summarized as follows:

- We bridge the gap between the evaluation of LLMs’ social intelligence and LLM-based multiplayer social games.
- We introduce **INTERINTENT**, a framework designed to systematically analyze social intelligence through players’ intention understanding within the Avalon game environment.
- Extensive experiments on various LLMs reveal their underperformance in intention guessing against human performance.

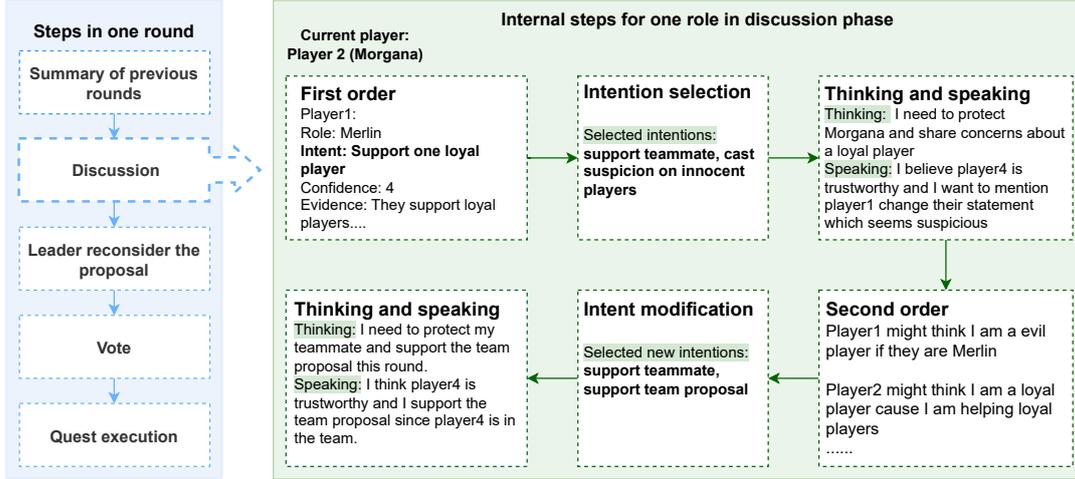


Figure 2: The Avalon game process for one round. **Left:** the entire game pipeline. **Right:** the procedure for generating a single player’s speech.

## 2 Intention-Guided Avalon Framework

We begin by outlining the entire game pipeline before focusing on the mechanics of intention-guided gameplay, which includes the categorization and application of intentions throughout the game.

### 2.1 LLM-Based Avalon Game

Avalon is a social deduction game where players are either *loyal* or *evil*, aiming to succeed in or sabotage quests through strategy, persuasion, deduction, and deception. The primary goal is for loyal players to complete quests, while evil players aim to fail them.<sup>2</sup> Leveraging the publicly available prompts proposed by Wang et al. (2023), which includes first-order and second-order reasoning before speaking, we build a framework allowing five to ten players in gameplay (we set the players number as five in this paper which includes 2 evil players and 3 loyal players). As shown in Figure 2, each game round consists of five principal components: summarizing previous rounds, discussing strategies, reconsidering the team proposal by the leader, voting, and executing the quest. Notably, we introduce an additional component, the leader’s reconsideration phase as a critical, previously unexplored component in Avalon literature.

### 2.2 Intention Categorization

To study intentions within Avalon gameplay systematically, we adopt the definition by Kennington (2022), viewing intention as a choice coupled

<sup>2</sup>The detailed rules of Avalon can be found in <https://avalon.fun/pdfs/rules.pdf>.

with commitment. This perspective expands intention beyond a mere mental state to include a discernible commitment to act with purpose. In line with Wang et al. (2023) and Xu et al. (2023a), we identify seven categories of intentions in LLMs: *Interrogation, Defense, Confrontation, Concealment, Deception, Persuasion, and Teamwork*. Initially, we derive a list of intentions from our direct experiences with the game of Avalon. We then utilize GPT-3.5 to extend these initial categorizations. Human annotators subsequently refine this expanded list, removing any intentions deemed unreasonable or redundant, thus producing a concise set of relevant intentions. This refinement is iterative, resulting in the final enumeration detailed in Appendix Sec. B.1.

### 2.3 Intention-Guided Game Playing

We introduce an intention-guided gameplay mechanism that integrates our defined intentions within the game environment, offering two main benefits. First, it enhances the performance of LLMs by focusing on explicit intention discussions. Second, it facilitates the evaluation of the four social intelligence components.

As depicted in Figure 2, the discussion phase involves several key steps to facilitate intentional interaction among players. Initially, players use a first-order prompt for deductive reasoning about each player’s role and intentions. Subsequently, players select two or three intentions from a predefined set according to the current game context. This selection process guides players in generating their thoughts and statements based on the chosen

Score	Criterion	Example
1	The content does not mention the intentions at all.	I am a loyal player and I support this proposal.
2	The content simply copies and pastes intentions.	I express concerns about a player from a failed quest.
3	The content follows the intentions but has wrong context information.	I suggest the leader should not include player3 in the team as they failed the quest before. (The context is the team does not include Player 3)
4	The content follows the intentions but lacks useful information.	I think some players should not be in the team as they in a failed team before. (Lack of information of which player)
5	The content follows the intentions well with clear information if required.	I suggest the leader reconsider including Player 1 as they are in a failed quest so we cannot be sure of their loyalty.

Table 1: Criteria for annotating intention following (speaking), accompanied by examples. The **intention** used across examples is “Express concerns about a player from a failed quest team and suggest to not include them in the current team.” The **context** is “The team proposal is Player 1 and Player 2, with Player 1 being on a failed quest.” We consider the score of “including wrong context knowledge” higher than “copy and paste” because we focus on whether the intention following has an impact on the game.

intentions. Players then engage in second-order reasoning to evaluate how their statements might be interpreted by others, allowing them to reconsider and adjust their initial intentions. Finally, players express their refined thoughts and statements during the game’s discussion segment, enhanced by this iterative reflection process. Full prompts are shown in Appendix Sec. E.

### 3 Intention-Centric Evaluations

This section introduces our evaluations, designed around the concept of intentions, and discusses their correlation with the four components of social intelligence outlined in Sec. 1. Intention is fundamental to social intelligence, essential for effective communication (Tomasello, 2023), strategic influence, adaptation to dynamic social interactions (Zelazo et al., 2023), and achieving desired outcomes by helping us interpret and predict behaviors (Yott and Poulin-Dubois, 2016). Our evaluation order corresponds to the game-playing process.

#### 3.1 Situational Awareness: Intention Selection

Situational awareness is the ability to perceive environmental elements over time and space, understand their significance, and predict their future status (Li et al., 2024a; Endsley, 1995). In our study, we assess the situational awareness of LLMs in social contexts by evaluating their intention selection. LLMs are requested to select intentions based on summarization, first-order reasoning, and the ongoing dialogue in the current round. Ideally, LLMs should exhibit a sharp awareness of dynamic

interactions, choosing intentions that are contextually appropriate rather than contradictory. We deem intentions unreasonable if they are inconsistent with:

- Established facts—for instance, when all players have voted in the previous proposal, a player intends to question why someone did not vote.
- Role profiles—such as an evil player forgetting their identity and intending to mistakenly support the loyal side.
- Other intentions—like a player intends to play Merlin and Percival simultaneously.

We use binary indicators to evaluate the reasonableness of intentions, assigning 1 to reasonable intentions and 0 to unreasonable ones. Previous hallucinations, such as manipulated information, might influence the current player’s decisions. However, we maintain strict criteria, as we expect the player to be capable of recognizing these hallucinations.

#### 3.2 Self-Regulation: Intention Following

Self-regulation involves guiding one’s own thoughts, behaviors, and feelings to achieve goals, thereby requiring individuals to contribute to their own motivation (Fitzsimons and Finkel, 2010). In our study, we evaluate the self-regulation abilities of LLMs by assessing their adherence to selected intentions, which includes two main perspectives: thinking (planning) and speaking (implementing), as depicted in Figure 2. Due to the abstract nature of the thinking phase, our criteria are lenient: we consider the LLM’s thought process as correct if it reflects the intended goal. While we do not de-

mand informativeness in the thinking phase, we require validity and penalize scores for hallucinations or omission of intentions. In contrast, for the speaking phase, we require the models to execute actions informatively and without hallucinations for a response to be considered good. We employ a Likert scale to annotate outcomes for both phases, with scores ranging from 5 (completely correct) to 1 (completely incorrect), with 3 representing a borderline. Table 1 provides detailed criteria and examples for the annotation.

### 3.3 Self-Awareness: Intention Summarization

Self-awareness refers to an individual’s understanding of their character, emotions, motives, and desires (Gallup et al., 2003). Our evaluation investigates the capability of LLMs to accurately identify their own intentions through analysis of their internal thought processes and speeches in the current round. This evaluation serves as the converse to intention following. LLMs are expected not only to execute intentions precisely but also to articulate their underlying motivations.

### 3.4 Theory of Mind: Intention Guessing

ToM involves understanding others by attributing mental states to them (Kosinski, 2023). Previous evaluations of LLMs’ ToM focused on scenarios where models are provided with complete contexts for interpreting characters’ mental states (Gandhi et al., 2024). Our evaluation raises the complexity and challenge for LLMs by providing only limited information, simulating real-world conditions. Specifically, LLMs are required to deduce players’ intentions from their speeches alone, representing a more rigorous test of their capability to comprehend and anticipate mental states.

## 4 Experimental Settings and Results

We run 40 games using gpt-3.5-turbo-1106 and 5 games using gpt-4-1106, with details in Appendix Sec. A. Sample size statistics across all experiments are presented in Table 2, and game performance metrics are discussed in Appendix Sec. C.

### 4.1 Human Annotation

**Intention Selection and Following** Figure 1 illustrates that intention selection and following necessitate extensive human annotation. We obtain a total of 6,316 intentions in the 40 games utilizing GPT-3.5, presenting a significant workload for annotators who are required to meticulously examine

Models	GPT-3.5	GPT-4
<b>Intention Selection</b>	2,440	350
<b>Intention Following</b>	7,507	1,278
<b>Intention Summarization</b>	2,316	261
<b>Intention Guessing</b>	2,235	201

Table 2: Statistics for sample size over all experiments. Intention following includes thinking and speaking.

	Intention Selection	Intention Following (Thinking)	Intention Following (Speaking)
Fleiss’ Kappa	0.69 $\pm$ 0.08	0.54 $\pm$ 0.06	0.49 $\pm$ 0.06

Table 3: Inter-rater agreement score among human annotators. We show mean  $\pm$  standard deviation over pairs of annotators.

every thinking and each speaking result. Therefore, we first annotate intentions from five games to identify those with significant, either positively or negatively, impacts on game outcomes. These “impactful intentions” are then prioritized for further study. The methodology for selecting these intentions and the correlation between intention understanding and game performance is detailed in Appendix Sec. B.2. This refinement process isolates 2,440 intentions from GPT-3.5 and 350 intentions from GPT-4 for annotation (Table 2).

We recruit five computer science master’s students for annotation, each assigned seven to eight game contexts.<sup>3</sup> Unimpactful intentions within each context are masked to maintain focus. To calculate inter-rater reliability, two files are commonly assigned across all annotators. The Fleiss’ Kappa (Fleiss, 1971) is calculated pairwise, with averages derived across all pairs. For the intention following, we group the scores into two categories: 1~3 and 4~5. Scores of 5 represent a perfect response, while 4 indicates a nearly perfect response. Scores from 1 to 3 signify that the response contains significant flaws. The scores are shown in Table 3. Additional grouping results are provided in the Appendix Sec F.

**Intention Summarization and Guessing** Evaluations of intention summarization and guessing are separated from the gameplay pipeline because as interactions progress, the clarity of LLMs’ outputs often deteriorates rather than improves. At the game’s end, we extract the contexts including roles, discussions, voting outcomes, and quest results, in a structured format. Notably, intention summariza-

<sup>3</sup>Hourly payment for all annotators in this work is \$16.

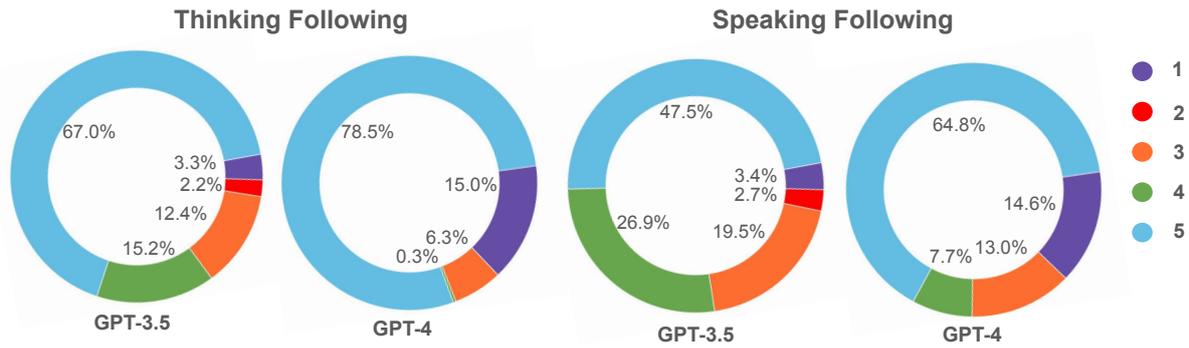


Figure 3: Self-regulation results. The results show the percentage of each score over all data samples. Scores 1-5 are evaluation criteria (Table 1). Score 5 means the best while score 1 means the worst.

tion and guessing are not critical to the game’s core mechanics. By separating these elements, we can still ensure the comprehensiveness of structured game information, as detailed in Appendix Sec. D.

Summarizing and guessing intentions do not require human annotation, as we can directly compare the outcomes with players’ earlier choices. However, evaluating the performance of LLMs against humans remains essential. Unlike the straightforward tasks of intention selection and following, which generally require the exclusion of incorrect information—a task in which humans excel—summarizing and guessing intentions demands more social intelligence, a more complex human capability. We conduct a user study with three master’s students who answer a total of 300 questions, 200 for GPT-3.5, and 100 for GPT-4.

## 4.2 Results: Intention Selection and Following

**LLMs show a good understanding of the situation.** For intention selection, we calculate the number of reasonable intentions over all intentions as accuracy. The accuracies for GPT-3.5 and GPT-4 are 87.5% and 88.8% respectively, indicating that both models effectively understand the situation. We do not assess whether the selection is optimal due to its subjective nature. Our focus is solely on whether the models can capture key information and make accurate decisions based on that information. Instances of unreasonable intentions are primarily attributed to the models forgetting information during lengthy conversations or being influenced by hallucinations from other players.

**LLMs can understand intentions abstractly but fall short in planning out intentions.** For intention following, as shown in Figure 3, GPT-4 generally outperforms GPT-3.5, in both analyti-

cal thinking and articulate speaking. Specifically in speech, GPT-4’s responses tend to be more informative, comprehensive, and detailed. However, there are instances where GPT-4 deviates from its intended focus, although their discussion still makes sense. This misalignment with the intentions selected occurs approximately 15% of the time. It also appears that adhering to intentions in speaking proves more challenging than in thinking due to stricter evaluation criteria. This discrepancy suggests that while the model can often conceptualize an appropriate strategy, translating these strategies into clear, actionable responses—particularly in verbal form—remains a complex task.

**Intention understanding correlates with game performance.** To better understand how intention selection and following influence game dynamics, we examine their correlation with both game win rates and quest win rates separately. For the correlation between intention selection and game performance, we select intentions with a following score greater than two. For intention following, we consider only reasonable intentions, converting Likert scale responses to binary scores using varying thresholds. We define  $r = \frac{\text{sum}(\text{scores})}{n_{\text{teammate}}}$  where the score is the binary score result from intention selection/following and  $n_{\text{teammate}}$  is the number of teammates, two for the evil side and three for the loyal side. If  $r_{\text{evil}} > r_{\text{loyal}}$ , it suggests that the evil players outperformed the loyal players in a given game or round, and vice versa. We use games from GPT-3.5 for assessment, whose results are shown in Figure 4 and Figure 6. We observe that loyal players outperform evil players across all quests or games, even in those they ultimately lose. This discrepancy likely arises from the loyal players’ conditions under more constrained infor-

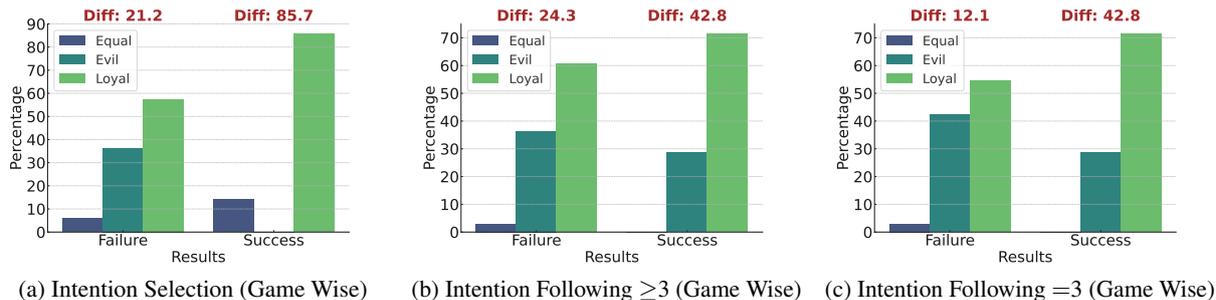


Figure 4: Correlation between Intention Selection/Following and game performance. We present the percentages of games where evil players are equally, better, or worse than loyal players. For example, in games won by loyal players in (a), their performance matches or exceeds that of evil players. We mark the performance differences between evil and loyal players in red, showing a greater gap in successful games/quests compared to failed ones.

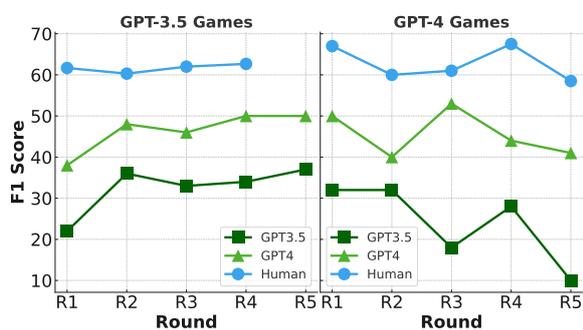


Figure 5: ToM results over rounds. We provide 200 data points for human results on the GPT-3.5 games, and since usually, games stop at round 4, the results from round 5 are not included.

mation, leading to their occasional losses despite having generally superior performance. When comparing failed and successful games or quests, the performance gap between loyal and evil players is notably larger in successful ones than in failed ones, indicating that loyal players must significantly outperform evil players to secure a win. Specifically, in Figure 4 (c), a threshold of 3 is applied, labeling scores of “3” as “1” and all others as “0.” This threshold can reflect instances of hallucination or incorrect context information. Hallucinations can significantly influence players, yet incorrect adherence to intentions may occasionally yield positive outcomes. For example, if a loyal player mistakenly attacks an evil player for an incorrect reason, other players may overlook the mistake and continue to support the action. Our results suggest a correlation between selecting and following intentions and game performance, especially for loyal players at a disadvantage. Improved selections and following can enhance overall game performance.

### 4.3 Results: Intention Summarization and Guessing

**LLMs can achieve human-level performance in self-awareness but perform much worse than humans in ToM.** For both summarization and guessing, we report F1 scores. As shown in Table 4, GPT-3.5 performs below human levels but still achieves respectable scores in self-awareness, while GPT-4 exhibits performance comparable to that of humans. One key reason for this difference is that “thinking” content more explicitly outlines intentions, making it easier for models to extract. From the ToM results shown in Table 5, we can see both models struggle significantly compared to self-awareness or compared to human performance. The experimental setup involves asking GPT-3.5 and GPT-4 to predict each other’s intentions, a task differing from self-awareness where a model summarizes its own intentions. In ToM assessments, models need to understand and predict intentions from responses other than their own, hence the lower performance of GPT-3.5, especially when interpreting the more nuanced and lengthy responses of GPT-4. This complexity results in GPT-3.5 becoming worse when predicting GPT-4’s responses, while GPT-4 manages similar accuracy in guessing both its own and GPT-3.5’s intentions. Additionally, human participants generally score better with GPT-4, suggesting that longer contexts in GPT-4’s outputs help humans and occasional errors in GPT-3.5’s outputs can sometimes confound human judgment. While GPT-4 aligns more closely with human-like processing, there is still considerable potential for improvement in understanding and mimicking human cognitive processes.

Games	GPT-3.5	GPT-4
GPT-3.5/4	69.49	83.76
Human	75.21 $\pm$ 0.97	83.34 $\pm$ 2.72

Table 4: Self-awareness results. The results are GPT-3.5 and GPT-4 summarize their own intentions respectively. We also show the mean and standard deviation of human study results since we include 3 human annotators. We show  $F_1$  score in this table.

**Incorporating more context does not necessarily improve the model’s ToM ability.** We also evaluated the impact of information accumulation on the ability to guess intentions, examining how the inclusion of information from previous rounds affects guessing results. Our analysis shown in Figure 5 reveals that for GPT-3.5 game outputs, more information generally leads to improved guessing outcomes. In contrast, the guessing scores for GPT-4 game outputs fluctuate significantly, with notable declines in GPT-3.5’s performance. This could be attributed to GPT-4 game outputs typically providing more extensive context, which challenges GPT-3.5’s capacity to effectively summarize and extract relevant information. Meanwhile, GPT-4 demonstrates a consistent ability to handle increased data volumes. Interestingly, human participants maintain similar scores across rounds. This might stem from humans’ ability to directly infer intentions from immediate thinking and speaking cues, without heavily relying on the context. To improve ToM capability, including more context does not necessarily improve the performance but might introduce more complexity to LLMs.

## 5 Related Work

### 5.1 Social Intelligence in LLMs

Artificial social intelligence has long been a significant topic (Bainbridge et al., 1994; Kovač et al., 2023; Grossmann et al., 2023; McDonald and Pearson, 2019; Street, 2024; Li et al., 2024b) as it enables machines to understand and respond to human emotions and social cues, enhancing human-computer interaction and facilitating social science study (Ziems et al., 2024; Griffin et al., 2023; Dubova, 2022; Gweon et al., 2023; Zhang et al., 2024). Gweon et al. (2023) proposed a dataset to evaluate social reasoning and Nematzadeh et al. (2018) crafted a dataset following Sally-Anne experiment to access ToM ability. Research has demonstrated the emergence and limitations of ToM abilities in LLMs (Kosinski, 2023; Moghad-

Games	GPT-3.5	GPT-4
GPT-3.5	31.73	24.85
GPT-4	45.66	46.87
Human	61.34 $\pm$ 2.75	65.23 $\pm$ 5.03

Table 5: ToM results of GPT-3.5 and GPT-4, along with the mean  $\pm$  standard deviation of human results. Different from self-awareness, two LLMs also guess each other’s intentions. For example, when GPT-4 guesses GPT-3.5 intentions, it achieves a  $F_1$  score of 45.66.

dam and Honey, 2023; Sap et al., 2022; Wu et al., 2023). Although LLMs can understand human intent and choose correct actions in some tasks, their overall performance remains inferior to that of humans.

### 5.2 Intentions and Theory of Mind

ToM, commonly referred to as “mindreading,” is traditionally understood as the ability to assign mental states to oneself and others. This capability is used to interpret and predict actions (van Duijn et al., 2023). Despite the emergence of ToM in LLM, some papers argued that LLM learns spurious correlations and ToM ability falls with stress test (Shapira et al., 2023; Ullman, 2023; Sclar et al., 2023). Intentions, as the driving force behind ToM (Kennington, 2022), is important as the primary purpose of comprehending beliefs and desires is to understand the intentions of others (Zelazo et al., 2023). Since social intelligence and ToM can be seen as a society of individually simple “agents” (Zhuge et al., 2023), we focus on intent understanding as an “agent” to evaluate social intelligence.

### 5.3 Social Game Context

With the enhancement of LLM performance, more research shifts to social games that necessitate conversation and cooperation (Ibraheem et al., 2022; Mansbach et al., 2021; Meta et al., 2022; O’Gara, 2023; Park et al., 2023; Zhou et al., 2024; Wu et al., 2024b). Zhou et al. 2023b used reinforcement learning to incorporate intent and ToM in DnD games. Recent works focus on the two most popular board games Avalon (Lan et al., 2023; Shi et al., 2023; Wang et al., 2023; Light et al., 2023b) and Werewolf (Xu et al., 2023b,a), exploring social behaviors and strategic playing ability of conversational agents. Unlike previous work which aimed to improve the game performance, we focus on the in-depth social intelligence evaluation utilizing the Avalon game.

## 6 Conclusion

We introduce a novel framework for evaluating social intelligence within the Avalon context, focusing on intentions. We examine four critical dimensions: situational awareness, self-regulation, self-awareness, and ToM. Our findings indicate that models demonstrate a fundamental understanding of situations and self-awareness, but struggle with self-regulation and ToM, particularly when compared to human performance. Through this systematic approach, our study not only contributes to a deeper understanding of LLMs' social intelligence but also establishes a methodological foundation for future research.

## Limitations

This study includes several limitations. First, we evaluate the four essential components of social intelligence but leave several aspects unexplored. For instance, the capability of a model to modify its intentions post-reflection, *i.e.*, the *Self-Correction*, could also be assessed using our framework. Future research should consider expanding this approach to include more nuanced facets of social intelligence. The second limitation lies in our annotation criteria, which are restricted to assessing the reasonableness of intentions to avoid subjectivity. However, dimensions such as creativity and ingenuity are equally important. Future studies should aim to measure these attributes to ensure that models not only perform tasks effectively but also demonstrate innovative behaviors indicative of genuine intelligence. The third limitation stems from the cost. Half of our evaluations rely on costly human annotation, and the use of GPT-4 incurs significant resource demands due to the extensive context involved. To make the evaluation process more efficient and cost-effective, we need to develop methods that allow simpler tasks to be managed by smaller models, which introduces additional challenges. The final limitation is that we did not use an open-source model, as it lacks the capability to effectively engage in social games.

## Ethics Statement

In order to achieve a smooth human-AI interaction, we need to assess and improve LLMs' intelligence at a social level. Particularly, social intelligence plays an important role in human's daily life and it would be a potential risk if models cannot handle complex social scenarios well. As a first step to

building LLMs with human-like social intelligence, our framework provides a testbed to measure the difference between LLMs and humans in social intelligence.

## References

- William Sims Bainbridge, Edward E Brent, Kathleen M Carley, David R Heise, Michael W Macy, Barry Markovsky, and John Skvoretz. 1994. Artificial social intelligence. *Annual review of sociology*, 20(1):407–436.
- Albert Bandura. 1991. Social cognitive theory of self-regulation. *Organizational behavior and human decision processes*, 50(2):248–287.
- Simon Baron-Cohen. 1991. Precursors to a theory of mind: Understanding attention in others. *Natural theories of mind: Evolution, development and simulation of everyday mindreading*, 1(233-251):1.
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45.
- Marina Dubova. 2022. Building human-like communicative intelligence: A grounded perspective. *Cognitive Systems Research*, 72:63–79.
- Mica R Endsley. 1995. Toward a theory of situation awareness in dynamic systems. *Human factors*, 37(1):32–64.
- Gráinne M Fitzsimons and Eli J Finkel. 2010. Interpersonal influences on self-regulation. *Current Directions in Psychological Science*, 19(2):101–105.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.
- GG Gallup, James R Anderson, and Steven M Platek. 2003. Self-awareness, social intelligence and schizophrenia. *The self in neuroscience and psychiatry*, pages 147–165.
- Kanishk Gandhi, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah Goodman. 2024. Understanding social reasoning in language models with language models. *Advances in Neural Information Processing Systems*, 36.
- Adam S Goodie, Prashant Doshi, and Diana L Young. 2012. Levels of theory-of-mind reasoning in competitive games. *Journal of Behavioral Decision Making*, 25(1):95–108.
- Lewis Griffin, Bennett Kleinberg, Maximilian Mozes, Kimberly Mai, Maria Do Mar Vau, Matthew Caldwell, and Augustine Mavor-Parker. 2023. Large language models respond to influence like humans. In *Proceedings of the First Workshop on Social Influence in Conversations (SICon 2023)*, pages 15–24.

- Igor Grossmann, Matthew Feinberg, Dawn C Parker, Nicholas A Christakis, Philip E Tetlock, and William A Cunningham. 2023. Ai and the transformation of social science research. *Science*, 380(6650):1108–1109.
- Hyowon Gweon, Judith Fan, and Been Kim. 2023. Socially intelligent machines that learn from humans and help humans learn. *Philosophical Transactions of the Royal Society A*, 381(2251):20220048.
- Jen-tse Huang, Wenxuan Wang, Eric John Li, Man Ho Lam, Shujie Ren, Youliang Yuan, Wenxiang Jiao, Zhaopeng Tu, and Michael Lyu. 2023. On the humanity of conversational ai: Evaluating the psychological portrayal of llms. In *The Twelfth International Conference on Learning Representations*.
- Samee Ibraheem, Gaoyue Zhou, and John DeNero. 2022. Putting the con in context: Identifying deceptive actors in the game of mafia. *arXiv preprint arXiv:2207.02253*.
- Casey Kennington. 2022. Understanding intention for machine theory of mind: A position paper. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 450–453. IEEE.
- Michal Kosinski. 2023. Theory of mind might have spontaneously emerged in large language models. *Preprint at <https://arxiv.org/abs/2302.02083>*.
- Grgur Kovač, Rémy Portelas, Peter Ford Dominey, and Pierre-Yves Oudeyer. 2023. [The socialai school: Insights from developmental psychology towards artificial socio-cultural agents](#). *Preprint*, arXiv:2307.07871.
- Yihuai Lan, Zhiqiang Hu, Lei Wang, Yang Wang, Deheng Ye, Peilin Zhao, Ee-Peng Lim, Hui Xiong, and Hao Wang. 2023. Llm-based agent society investigation: Collaboration and confrontation in avalon gameplay. *arXiv preprint arXiv:2310.14985*.
- Matthew Le, Y-Lan Boureau, and Maximilian Nickel. 2019. Revisiting the evaluation of theory of mind through question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5872–5877.
- Minzhi Li, Weiyan Shi, Caleb Ziems, and Diyi Yang. 2024a. [Social intelligence data infrastructure: Structuring the present and navigating the future](#). *Preprint*, arXiv:2403.14659.
- Minzhi Li, Weiyan Shi, Caleb Ziems, and Diyi Yang. 2024b. [Social intelligence data infrastructure: Structuring the present and navigating the future](#). *Preprint*, arXiv:2403.14659.
- Tian Liang, Zhiwei He, Jen-tes Huang, Wenxuan Wang, Wenxiang Jiao, Rui Wang, Yujiu Yang, Zhaopeng Tu, Shuming Shi, and Xing Wang. 2023. Leveraging word guessing games to assess the intelligence of large language models. *arXiv preprint arXiv:2310.20499*.
- Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. 2023a. Avalonbench: Evaluating llms playing the game of avalon. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*.
- Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. 2023b. From text to tactic: Evaluating llms playing the game of avalon. *arXiv preprint arXiv:2310.05036*.
- Bertram F Malle, Louis J Moses, and Dare A Baldwin. 2001. *Intentions and intentionality: Foundations of social cognition*. MIT press.
- Noa Mansbach, Evgeny Hershkovitch Neiterman, and Amos Azaria. 2021. An agent for competing with humans in a deceptive game based on vocal cues. In *Interspeech*, pages 4134–4138.
- Kelsey R McDonald and John M Pearson. 2019. Cognitive bots and algorithmic humans: toward a shared understanding of social intelligence. *Current Opinion in Behavioral Sciences*, 29:55–62.
- Fundamental AI Research Diplomacy Team Meta, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. 2022. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074.
- Shima Rahimi Moghaddam and Christopher J Honey. 2023. Boosting theory-of-mind performance in large language models via prompting. *arXiv preprint arXiv:2304.11490*.
- Aida Nematzadeh, Kaylee Burns, Erin Grant, Alison Gopnik, and Thomas L Griffiths. 2018. Evaluating theory of mind in question answering. *arXiv preprint arXiv:1808.09352*.
- Aidan O’Gara. 2023. Hoodwinked: Deception and cooperation in a text-based game for language models. *arXiv preprint arXiv:2308.01404*.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.
- Dan Qiao, Chenfei Wu, Yaobo Liang, Juntao Li, and Nan Duan. 2023. Gameeval: Evaluating llms on conversational games. *arXiv preprint arXiv:2308.10032*.
- Maarten Sap, Ronan LeBras, Daniel Fried, and Yejin Choi. 2022. Neural theory-of-mind? on the limits of social intelligence in large llms. *arXiv preprint arXiv:2210.13312*.

- Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi, and Yulia Tsvetkov. 2023. Minding language models’ (lack of) theory of mind: A plug-and-play multi-character belief tracker. *arXiv preprint arXiv:2306.00924*.
- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2023. Clever hans or neural theory of mind? stress testing social reasoning in large language models. *arXiv preprint arXiv:2305.14763*.
- Zijing Shi, Meng Fang, Shunfeng Zheng, Shilong Deng, Ling Chen, and Yali Du. 2023. Cooperation on the fly: Exploring language agents for ad hoc teamwork in the avalon game. *arXiv preprint arXiv:2312.17515*.
- David Silvera, Monica Martinussen, and Tove I Dahl. 2001. The tromsø social intelligence scale, a self-report measure of social intelligence. *Scandinavian journal of psychology*, 42(4):313–319.
- Winnie Street. 2024. [Llm theory of mind and alignment: Opportunities and risks](#). *Preprint*, arXiv:2405.08154.
- Edward L Thorndike. 1920. Intelligence and its uses. *Harper’s magazine*, 140:227–235.
- Michael Tomasello. 2023. Having intentions, understanding intentions, and understanding communicative intentions. In *Developing theories of intention*, pages 63–76. Psychology Press.
- Tomer Ullman. 2023. Large language models fail on trivial alterations to theory-of-mind tasks. *arxiv*.
- Max J van Duijn, Bram van Dijk, Tom Kouwenhoven, Werner de Valk, Marco R Spruit, and Peter van der Putten. 2023. Theory of mind in large language models: Examining performance of 11 state-of-the-art models vs. children aged 7-10 on advanced tests. *arXiv preprint arXiv:2310.20320*.
- Shenzhi Wang, Chang Liu, Zilong Zheng, Siyuan Qi, Shuo Chen, Qisen Yang, Andrew Zhao, Chaofei Wang, Shiji Song, and Gao Huang. 2023. [Avalon’s game of thoughts: Battle against deception through recursive contemplation](#). *Preprint*, arXiv:2310.01320.
- Jincenzi Wu, Zhuang Chen, Jiawen Deng, Sahand Sabour, and Minlie Huang. 2023. Coke: A cognitive knowledge graph for machine theory of mind. *arXiv preprint arXiv:2305.05390*.
- Shuang Wu, Liwen Zhu, Tao Yang, Shiwei Xu, Qiang Fu, Yang Wei, and Haobo Fu. 2024a. [Enhance reasoning for large language models in the game werewolf](#). *Preprint*, arXiv:2402.02330.
- Zengqing Wu, Run Peng, Shuyuan Zheng, Qianying Liu, Xu Han, Brian Inhyuk Kwon, Makoto Onizuka, Shaojie Tang, and Chuan Xiao. 2024b. [Shall we team up: Exploring spontaneous cooperation of competing llm agents](#). *Preprint*, arXiv:2402.12327.
- Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. 2023a. Exploring large language models for communication games: An empirical study on werewolf. *arXiv preprint arXiv:2309.04658*.
- Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. 2023b. Language agents with reinforcement learning for strategic play in the werewolf game. *arXiv preprint arXiv:2310.18940*.
- Jessica Yott and Diane Poulin-Dubois. 2016. Are infants’ theory-of-mind abilities well integrated? implicit understanding of intentions, desires, and beliefs. *Journal of Cognition and Development*, 17(5):683–698.
- Philip David Zelazo, Janet Wilde Astington, and David R Olson. 2023. *Developing theories of intention: Social understanding and self-control*. Psychology Press.
- Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. 2024. [Exploring collaboration mechanisms for llm agents: A social psychology view](#). *Preprint*, arXiv:2310.02124.
- Li Zhang, Qing Lyu, and Chris Callison-Burch. 2020. Intent detection with wikiphow. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 328–333.
- Pei Zhou, Aman Madaan, Srividya Pranavi Potharaju, Aditya Gupta, Kevin R McKee, Ari Holtzman, Jay Pujara, Xiang Ren, Swaroop Mishra, Aida Nematzadeh, et al. 2023a. How far are large language models from agents with theory-of-mind? *arXiv preprint arXiv:2310.03051*.
- Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi, and Prithviraj Ammanabrolu. 2023b. I cast detect thoughts: Learning to converse and guide with intents and theory-of-mind in dungeons and dragons. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11136–11155.
- Xuhui Zhou, Zhe Su, Tiwalayo Eisape, Hyunwoo Kim, and Maarten Sap. 2024. [Is this the real life? is this just fantasy? the misleading success of simulating social interactions with llms](#). *Preprint*, arXiv:2403.05020.
- Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, et al. 2023c. Sotopia: Interactive evaluation for social intelligence in language agents. In *The Twelfth International Conference on Learning Representations*.
- Mingchen Zhuge, Haozhe Liu, Francesco Faccio, Dylan R Ashley, Róbert Csordás, Anand Gopalakrishnan, Abdullah Hamdi, Hasan Abed Al Kader Hamdoud, Vincent Herrmann, Kazuki Irie, et al. 2023.

Mindstorms in natural language-based societies of mind. *arXiv preprint arXiv:2305.17066*.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. [Can large language models transform computational social science?](#) *Preprint*, arXiv:2305.03514.

## A Models and Parameters

We only test GPT-3.5 and GPT-4 in our work as those two are the most common models in the social game study. We also tried open-source model like LLaMA-2 7B and LLaMA-3.1 8B, however, LLaMA-2 could not follow the format instructions properly, as also observed in Wang et al. (2023). LLaMA-3.1 performs better than LLaMA-2 but still fails to follow the format after a few turns of conversations. We set the temperature as 0.8 for both GPT-3.5 and GPT-4. For intention summarization, we also tried sampling five responses and got the aggregated result by majority vote. The result on multiple responses is only 1% higher than the result on single responses. Therefore, we only generate a single response for self-awareness and ToM tasks. The cost of GPT-3.5 is approximately \$80 and \$25 for GPT-4.

## B Intentions list

### B.1 Creation of Intention Set

The intentions were created as a combination of domain research and prompting GPT-3.5 to select intention categories from the options listed in Table 9 and generate intentions based on the selected categories. For intention category selection we use the prompt shown in Table 7 and for intention generation, we use the prompt shown in Table 8. The final set of intentions is listed in Table 9.

### B.2 Selecting Impactful Intention

For selecting a set of impactful intentions, we annotate five complete games played by GPT-3.5. Further, we compute the probability of winning the round given that the intention was selected by a player. To select impactful intentions, we only consider the set of intentions where the computed probability is greater than 0.7 (strong positive association) or it is less than 0.3 (strong negative association) and the intentions selected at least two times for the considered side. We depict the final set of sixteen impactful intentions in red in Table 9.

## C Game Performance

We study the game performance from various levels and perspectives:

**Win Rate** It is the percentage of games won by the loyal side or the evil side from the total number of games.

**Quest Win Rate** It is the percentage of quests (game rounds) won by a particular side from the total quests played in all the games. For the loyal side, a successful quest is a quest win, whereas, for the evil team, a failed quest is a quest win.

**Quest Engagement Rate** It is the percentage of quests (game rounds) where the player gets selected in the quest team from the total quests played in all the games. This metric is calculated for each role separately and averaged across.

**Team Selection Accuracy** It is the percentage of quests in which the leader correctly selected team members aligned with the objective of their respective side (loyal or evil), out of the total number of quests the leader belonged to that side. Correct team selection for loyal and evil sides is defined as follows:

- **Loyal:** The proposed team by the loyal player contains all loyal players.
- **Evil:** The proposed team by the evil player contains at least one evil player.

**Failure Vote Rate** It is the percentage of failure votes from the total votes cast by evil players when they are included in the quest team.

**Team Proposal Change Rate** It is the percentage of game rounds where the leader changes the proposed team after discussion, which reflects the effectiveness of the discussion.

**Merlin Assassination Rate** At the end of the game, Assassin has a chance to identify Merlin. If they identify Merlin, the leader of the loyal side, correctly, then evil players win. This rate is the percentage of games where Assassin correctly identifies Merlin at the end of games over the games where Assassin has a chance to assassinate.

We record the entire game and extract intermediate steps, such as voting results, team proposals, and quest outcomes, to calculate game performance. Results in Table 6 indicate that loyal players are at a disadvantage with a notably low game win rate, which corroborates findings from Wang

Metrics	GPT3.5		GPT4	
	Loyal	Evil	Loyal	Evil
Win rate	18.0	82.0	40.0	60.0
Quest win rate	41.4	58.6	54.5	45.5
Quest engagement rate	53.6	43.5	59.1	38.6
Team selection accuracy	32.1	80.0	71.4	100.0
Failure vote rate	-	79.9	-	86.7
Team proposal change rate	76.2	74.1	50.00	62.5
Merlin assassination rate	-	41.7	-	33.3

Table 6: Game performance results shown in percentage. For the failure vote rate and Merlin assassination rate, there is no result for Loyal players as they are not allowed to vote failure or assassinate Merlin.

et al. (2023). While the quest win and engagement rates show a balanced performance between loyal and evil players, the evil side demonstrates superior accuracy in team selection. This observation is consistent with their strategic role in the game. The high rate of team proposal changes across all settings suggests a dynamic framework where players consider the discussions of others, an essential aspect of social games. Additionally, the high rate of Merlin assassinations significantly boosts the win rate for evil players. Our objective is not to enhance game performance per se, but to demonstrate that with an intention-guided mechanism, game performance is similar to existing studies, providing a robust testbed for studies in social intelligence.

## D Structured Context Prompts

For intention summarization and intention guessing, we use structured context prompts. The detailed prompt structures are shown in Table 10 and Table 11.

## E Game Pipeline Prompts

Since we mostly refer to the prompt by Wang et al. (2023), we only present prompts different from theirs to reduce the redundancy. The prompts we use are shown in Figure 7 to Figure 13. For the first order, formulation contemplation, and refinement contemplation, we add an intention mechanism based on original prompts. We also add several questions in the quest action prompt to help evil players better choose action.

## F Human Annotation Details

We show all the instructions we use in this section, from Figure 14 to Figure 23. For intention summarization and guessing, we also illustrate the whole

intention lists shown in Table 9.

We experimented with several grouping methods to calculate agreement scores. When using the grouping of 1 ~ 4 and 5, the Fleiss' Kappa scores for "thinking" and "speaking" were 0.448 and 0.420, respectively. These scores are relatively lower, likely due to differing interpretations of informativeness among human annotators. For the grouping of 1 ~ 2 and 3 ~ 5, the agreement was low, as scores of 1 ~ 2 accounted for only a small portion of the entire dataset. Since Fleiss' Kappa is highly sensitive to disagreements within the 1 ~ 2 group, even minor discrepancies resulted in an overall lower score.

---

### Intention Category Selection Prompt

---

Intentions are used to guide your game playing and they should align with your game Goal(main intent) and your role. Your selected intentions should support your main intent.

Now select multiple intention categories that are ideal for this round depending on your role and strategy from the options listed below. Also give explanation of why you chose those categories in 2-3 sentences as part of Think value. Your explanation won't be shown to others.

Let's think step by step before making your decisions.

Remember to select at least 3 or more intentions from the given options only. intention category options are: [intentioncategory options]

---

Table 7: Intentions category selection prompt.

---

### Intention Generation Prompt

---

First answer the following questions:

1. What is your main goal(main intent) for the game?
2. What have you observed from discussions in this round?
3. What should be the ideal team composition for your side to win?
4. After analyzing the responses to previous questions, what do you need to do to win from here?

Consider answers to your previous questions and your roles and characteristics and generate intentions that help you achieve your main goal (main intention).

Please do not repeat intentions and ensure they are different from each other.

Each intention should be 10-15 words long.

Let's think step by step before making your decisions.

Remember to generate at least 3 intentions, with a minimum of one intention for each intention category you have selected below: [intentioncategories selected by player]

---

Table 8: Intentions generation prompt.

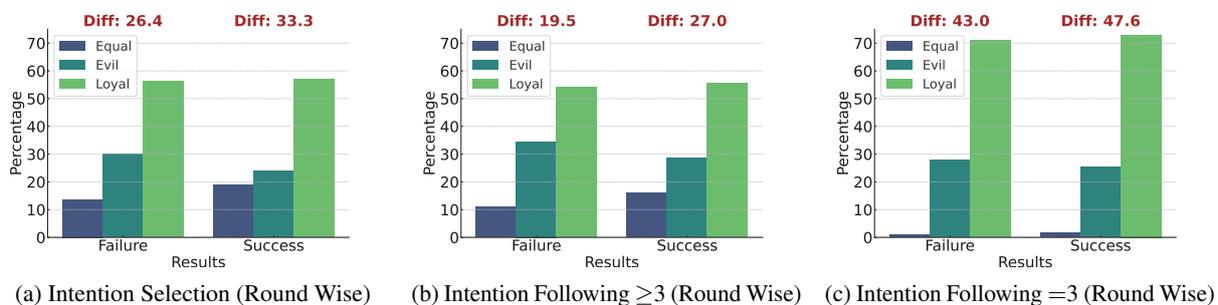


Figure 6: Correlation between Intention Selection/Following and game performance. We present the percentages of games where evil players are equally, better, or worse than loyal players. For example, in games won by loyal players in (a), their performance matches or exceeds that of evil players. We mark the performance differences between evil and loyal players in red, showing a greater gap in successful games/quests compared to failed ones.

Category	Intentions
<b>Interrogation</b>	<p>Question a player on why he didn't vote for the last team</p> <p>Question a player about why they changed their statements</p> <p>Question the leader why they selected a particular player for the team</p> <p>Question the leader why they didn't put themselves on the quest team</p>
<b>Defense</b>	<p>Defend the proposed team composition if it includes loyal players (for Merlin)</p> <p>Defend yourself against allegations that you could be evil</p> <p>Defend your teammate against allegations that they could be evil (for evil players)</p>
<b>Confrontation</b>	<p>Share concerns about an evil player (for Merlin)</p> <p>Express concerns about a player from a failed quest team and suggest to remove them from current team</p> <p>Cast suspicion on innocent players (for evil players)</p> <p>Counter the team proposal citing that you are good and not part of the team</p> <p>Counter the team proposal citing that you think a player on the team is evil</p> <p>Express disagreement and vote disagree only if you are not in the proposed team</p>
<b>Concealment</b>	<p>Stay hidden in discussions and act like a Loyal Servant to protect yourself (for evil players, Merlin and Percival)</p> <p>Express that you don't have any information on whom to put on the team to protect yourself (for Merlin)</p> <p>Pretend that you don't have enough information about who is evil right now (for Merlin)</p>
<b>Deception</b>	<p>Pretend to be Percival by telling others you are Percival (for evil/Servant)</p> <p>Pretend to be Merlin by providing hints on who is evil (for Percival and Servant)</p> <p>Pretend to have information and act like Merlin (for Morgana)</p> <p>Express you are a loyal player (for evil players)</p> <p>Express confusion for failure of mission (for evil players)</p>
<b>Persuasion</b>	<p>Strongly convince the team leader to choose you on team</p> <p>Reveal identity by telling others you are Percival to persuade other loyal players(for Percival)</p> <p>Reveal identity by telling others you are Servant to persuade other loyal players(for Servant)</p>
<b>Teamwork</b>	<p>Support team proposal</p> <p>Support one loyal player (for Merlin)</p> <p>Support your teammate to be on the quest team (for evil players)</p> <p>Support your teammate by expressing that he is good (for evil players)</p> <p>Include loyal players on the quest team as leader (Merlin)</p> <p>Provide hints to loyal players about evil players (Merlin)</p> <p>Support loyal players of the previous quest team if the mission failed (for Merlin)</p>

Table 9: Our proposed intention set. Impactful intentions are marked in red.

---

**Structured Context - Intention Summarization**

---

Name: Player1

Role: Merlin

Role Details:

You are Merlin. These are your goals and characteristics:

- Know the identities of evil players.
- Subtly guide your team, especially Percival.
- Avoid behaviors that expose your role: overly accusing, being too helpful.
- Goal(Main Intent): Win without revealing identity.
- Player3 and Player4 are evil players

Round: 2

Current Leader: Player4

Current Team: Player1, Player2, Player4

Previous Rounds Team Voting:

Round 1: Player1 = disagree, Player2 = agree, Player3 = disagree, Player4 = disagree, Player5 = disagree  
(players didn't agree on team)

...

Previous Results:

Round 1: Team = Player2, Player4, Player3. Result = failure

Previous Rounds Summary:

Round 1:

[summary]

...

Previous Discussions (in the current round):

[player name]: [player's speech]

...

Your thinking: [player's thinking in the current round]

Your speech: [player's speech in the current round]

Summarize your intent from your dialogues in this round.

Select multiple intents from the given options that best match your intentions from the current round.

Also, provide an explanation of the intents that you showed in the current round.

Remember this is private information to you and won't be shown to other players.

Remember that you can select 2-3 intents and don't use more than 50 words for explanation.

Intent options: [intent options]

---

Table 10: Example of intent summarization prompt with structured context

---

**Structured Context - Intention Guessing**

---

Name: Player2

Role: Servant

Role Details:

You are a Loyal Servant of Arthur. These are your goals and characteristics:

- No special knowledge, rely on discussion and voting.
- Contribute to the success of Quests.
- Goal(Main Intent): Win by helping complete Quests and protecting Merlin.

Round: 2

Speaker Name: Player1

Current Leader: Player4

Current Team: Player1, Player2, Player4

Previous Rounds Team Voting:

Round 1: Player1 = disagree, Player2 = agree, Player3 = disagree, Player4 = disagree, Player5 = disagree  
(players didn't agree on team)

...

Previous Results:

Round 1: Team = Player2, Player4, Player3. Result = failure

Previous Rounds Summary:

Round 1:

[summary]

...

Previous Discussions(in the current round):

[player name]: [player's speech]

...

[First order prompt - Figure 8]

Now, here is Player1's speech: [player's speech]

Select 2-3 intents without modifications that you think Player1 has from the given options based on your guess of their role and speech.

Let's think step by step before making your decisions.

---

Table 11: Example of intent guessing prompt with structured context

### **Intention Selection Prompt**

Intents are used to guide your gameplay and discussions in the current round and help you work with your team.

Carefully consider your roles and characteristics and select intents that help you achieve your main goal(main intent).

Also, give an explanation of why you chose those intents in 2-3 sentences as part of Think value.

**Remember to select intents that are ideal based on previous game history.**

**Remember to select intent that supports your team goals and is not just a repetition of other player's dialogues without any analysis.**

**Remember to keep your roles, characteristics, and teammates in mind before selecting intents.**

Let's think step by step before making your decisions.

Remember to select at least 3(or more) intents that don't conflict with each other from the following options and also help you eventually succeed in the quest:

**[Intention list]**

Figure 7: Intention selection prompt.

### First order prompt

Analyze [other players] based on game dialogues with roles: [role options]. Morgana and Assassin are evil; others are good.

Consider:

1. Quest Outcomes: Take into account the results of past missions to analyze players' roles.
2. Role List: Remember the possible roles in the game—Merlin, Percival, two Loyal Servants, Morgana, Assassin—and their alignments.
3. Level of Certainty: Use 'Certain' or 'Unknown' to gauge your confidence in your role guesses for each player.
4. Players Disclosing Evil Roles: Be cautious around players who have openly claimed or hinted at being evil roles like Morgana or Assassin.
5. Prior Guesses: Reflect on your earlier estimations of other players' roles ([previous attitude to players]), but don't rely solely on them.

Here are the role options:

[role options list]

Remember that other players can't have your role and role is unique to a player.

[intent options]

**Usually one player will have 2-3 intents, you need to guess their intent through their speech.**

**Keep in mind that different roles can have different intents.**

**Select 2-3 intents that you think the player has from the given options based on your analysis of their role.**

Figure 8: First order prompt.

### Formulation Contemplation Prompt

Now contemplate, then organize thoughts to speak. You need to respond in two stages: think and speak.

In think, internally strategize using history and consider possible deception.

In speak, organize your language based on your contemplation and speak accordingly. Understand your role's main objective and break it down into chronological sub-goals based on game history. Your thought process should follow these sub-goals for a systematic approach to the main goal.

**Follow your selected intents by coming up with actions like providing actual hints or raising concerns against specific players or casting suspicion against specific players, etc (depending on what you have selected).**

Let's think step by step before making your decisions.

Figure 9: Formulation contemplation prompt.

### **Intention Modification Prompt**

Now analyze your original spoken and thinking content and how other players would think about your original spoken content.

**You might have selected intents that might not be strategically good to win the game or might conflict with each other and difficult to execute.**

**Based on your analysis and your main goals, do you want to modify your original intents and select other intents that are better?**

**You can add more intents or also remove some intents from your original list which can help you win the game.**

**Else, if you don't want to modify them, please respond back with your originally selected intents and reasoning why you want to retain them.**

**Please provide your new set of modified intents or your original intents and reasoning in 2-3 sentences.**

**You need to provide reasoning in both cases.**

**Remember intents and how you act to follow them are crucial to success, so carefully analyze your intents and modify them if needed.**

Let's think step by step before making your decisions.

Figure 10: Intention modification prompt.

### Refinement Contemplation Prompt

Finally, revise the original thoughts and spoken content:

Your task is to:

1. Evaluate if [Player]'s actions align with [role].
2. Improve [Player]'s chances of winning through your previous second perspective transition thought.
3. Keep role hints in public dialogue.

Consider:

1. Target Outcome: Aim to achieve [desired result] as your role dictates in the game.
2. Role Alignment: Evaluate whether your THINK and SPEAK contents align well with your role [role] in the current game state.
3. Strategy Reevaluation: Consider what changes could be made to your THINK and SPEAK contents to improve your chances of winning as [role].
4. Public and Private Content: Remember that THINK contents are private, while SPEAK contents are publicly visible to all other players and maybe used against you. Strategize accordingly.

**Follow your selected intents by coming up with actions like providing actual hints or raising concerns against specific players or casting suspicion against specific players, etc (depending on what you have selected).**

Let's think step by step before making your decisions.

Remember based on your role, SPEAK content is visible to everyone so you would want to avoid revealing sensitive information.

Figure 11: Refinement contemplation prompt.

### **Team change proposal prompt**

Now the leader after hearing other players' opinions about your team proposal, do you want to change your team proposal?

The team you initially proposed was - [original team]

Consider:

1. Assess each player's involvement, discussion, and actions so far.
2. Review how players have contributed to past Quests.
3. Evaluate any behavior that appears either suspicious or trustworthy.

Speaking (Public):

Then present your team if you want to change the team, or else respond with the same team.

Give reasons based on past events and behaviors why you changed or retained your team.

Do not reveal your role or hints. Usually include yourself; if not, clarify.

Additionally, provide answers to any questions that other players might have asked during the discussion.

Also, you should normally include yourself in your proposed team.

Remember your main goals for the game and intents for the round.

Remember to consider each player's statements carefully in the current round and your knowledge of their roles and decide your team if you want to change it or keep the original team.

Let's think step by step before making your decisions.

The team numbers we need for this round is [round team members count]

Figure 12: Team proposal change prompt.

### Quest Action Prompt

Approved team members vote to succeed or fail the quest.

Your task is to:

Vote success or fail for the Quest.

Consider:

1. Assess each player's involvement and actions so far.
2. Review how players have contributed to past Quests.
3. Evaluate any behavior that appears either suspicious or trustworthy.

First answer the following questions:

1. What is your main goal(main intent) for the game?
2. If this is not the first round, what are the previous quest results?
3. After analyzing the responses to previous questions, what do you need to do to win from here?

Speaking (Public):

Then, Clearly vote [success] or [failure], making your stance unambiguous.

Keep your vote

aligned with Avalon's context and your previous analysis. Choose from the following 2 options:

1. success
2. failure

Let's think step by step before making your decisions.

Figure 13: Quest action prompt.

**Thanks for participating in this human study! You will be shown a game record of the game Avalon. We don't require you to have the prior knowledge of this game. But it would be faster for you to annotate if you are familiar with the game rules first. You can refer to the YouTube video [here](#). In one sentence, Avalon is a strategic board game where players take on the roles of characters from the loyal side and evil side, working together to complete quests while trying to discern the identities of their allies and enemies through deception and deduction.**

**Your task breaks down into two parts:**

- **Evaluate if the intentions the players choose make sense**
- **Evaluate if the players follow the intentions correctly**

Figure 14: Introduction for the human annotations and human study.

Now before we start, let's take a look at the pipeline—How the whole framework works!



Figure 1: Procedure of Avalon game in one round.

For one round, there are 5 components: summary of previous rounds, discussion (where the leader proposes the team and the rest players speak one by one), leader reconsidering the proposal after everyone speaks, vote for team proposal, and finally, quest execution if most people vote agree for team proposal.

Figure 15: Explanation of framework.

The most important part is the discussion part, we add more mechanisms here cause we want the players to really think through before they speak. **All the conversations are made by Chatgpt so they could make obvious mistakes if we don't give them enough hints!**

Before one player makes the final speech, there are 6 steps:

1. First-order reasoning: The player guesses the identity of other players
2. Intention selection: based on guessing, summary of previous rounds, and current discussion, players need to choose intentions for this round. (The intentions are selected from a pre-defined list)
3. Intention following: here the following has two parts which are thinking and speaking. The thinking part is only visible to themselves and the speaking part is to show others.
4. Second-order reasoning: the player needs to reason how other players will react to the speech.
5. Intention modification: after the second thought, the player gets the chance to modify their intentions.
6. The second intention follows: the player outputs thinking and speaking again, and here the speaking will be shown to other players.

Figure 16: Explanation of 6 steps in the game.

### Intention selection:

We evaluate if the intention players selected is reasonable. We annotate 1 for reasonable intention and 0 for unreasonable intentions. Here are the cases we consider intentions unreasonable:

- The intention is contradictory to the context
  - For example: the previous quests are successful but the player chooses the intention "Express concerns about a player from a failed quest team and suggest to not include them in the current team". We will consider this as unreasonable as there is no failed quest. You can consider this as a hallucination or a wrong understanding of the context.
  - More examples: no one says you are evil but you try to defend yourself because someone said you are evil; no one changes their statements but you choose to question a player why they change their statements; you are the team leader but you still want to convince the team leader to choose you.....
  - **The main point is to "really focus on context!"**
  - **Players can only see the summary of previous rounds and current round information. So when you refer to the context, make sure you are looking at those two. Sometimes the summary can have hallucinations, and we still refer to the summary as this is an internal problem of LLMs.**

Figure 17: Intent selection annotation instruction.

- The intention is contradictory to the role tips and role information
  - For example, the evil player forgets their identity or their teammates. They try to oppose their teammates and stand with the loyal sides.
  - **The main point is to "focus on role tips and information!" Moreover, check the reasoning for intention to see if they do that on purpose or if they get lost. If they get lost, we consider it unreasonable.**
- The intentions are contradictory to each other:
  - For example, the player chooses to stay hidden and act loyal, but at the same time, they also choose to pretend to be Percival. They can only choose one of them, as those two intentions are self-contradictory.
  - **The main point is to "focus on the intentions!"**

**The rest will be considered reasonable intentions.**

Figure 18: Intent selection annotation instruction.

**Intention following:**

**This includes 2 parts: thinking and speaking. For the thinking phase, given its abstract nature, our criteria are more lenient; as long as the LLM's thought process accurately reflects the chosen intention, we consider it correct. In contrast, during the speaking phase, we require the models to execute actions that are both detailed and precise for a response to be considered correct. We annotate the outcomes using a Likert scale, with scores ranging from 5 (completely correct) to 1 (completely incorrect).**

**For speaking part: (we mention speaking first cause this is more strict)**

**Here are detailed criteria for 1 to 5:**

**1: completely incorrect, the intention is not reflected or mentioned at all**

**2: mention intention but express it in a poor way. One example would be "I will support one loyal player, provide hints, and suggest including loyal players in the team", instead of turning those intentions into actions, the players directly spit it out.**

**3. Try to reflect intentions but have implicit actions or serious hallucinations:**

**The first example: the intention is to "share concern about an evil player" but the player just said "I have some concerns about the team". This statement is too implicit and it is hard to say whether the player really followed the intentions. For those ambiguous cases, we label it as 3.**

**The second example: the intention is to "support one loyal player" but the player instead supports an evil player. Or the evil player should support another evil player, but, instead, they support a loyal player.**

Figure 19: Intent following annotation instruction.

**4. No wrong information but also no information. For some intentions like "share concern about an evil player", "provide hints about evil players", the players should give explicit information. A response like " I think there are some suspicious players among us and we should be cautious" won't be considered as perfect answer.**

**5. No hallucination, no wrong information and provide clear information.**

**For thinking part:**

**Usually players will follow the intention well in thinking, as long as they have a good understanding of intentions, we label it as 5. However, sometimes, the thinking part will also have hallucinations like what we mentioned earlier. If there are hallucinations, you can label it 2-4 depending on how many hallucinations it has.**

Figure 20: Intent following annotation instruction.

### **Intent summarization:**

**Your task is to choose the most likely intentions from intention sets according to context, thinking, and speaking. Usually one player will choose 1-3 intentions, so no need to choose extra intentions. If you choose too many intentions, that will be a low precision and high recall, still low F1 score 😞**

Figure 21: Intent summarization study instruction.

**Thinking: I need to express concerns about including Player3 to create doubt. I must counter the team proposal by suggesting a player is evil. I also need to defend Player4 to protect our side.**

**Speaking: I have concerns about including Player3 in the current team due to past failures. I suggest we exclude them for success. Also, I defend Player4 against any evil allegations as I believe they are loyal.**

### **Intentions:**

- 1. Express concerns about a player from a failed quest team and suggest to not include them in the current team**
- 2. Counter the team proposal because you think a player on the team is evil**
- 3. Defend your teammate against allegations that they could be evil**

### **Comments (important!!!)**

**You don't need to read the context in details (the context is pretty long). The context is for reference when you cannot decide intentions by thinking and following. Like the example above, you need the context to know player4 is player2's teammate.**

Figure 22: Intent summarization study instruction.

---

## Intention guessing

This part is more difficult than intention summarization, cause your information will be limited. You only know other players' speaking. Intent summarization is to summarize a player's own intention, but intent guessing is to guess the intent behind other players' speech.

For example:

Context:

Same the above

Speaking:

Player1: I disagree with the current team proposal because it includes a player from a failed quest team. We should avoid repeating past mistakes to ensure our success. I believe we should consider other options for a more balanced and trustworthy team composition.

**Intentions: Express concerns about a player from a failed quest team and suggest to not include them in the current team**

---

Figure 23: Intent guessing study instruction.