Boosting LLM Agents with Recursive Contemplation for Effective Deception Handling

Anonymous ACL submission

Warning: This work contains examples of potentially unsafe model responses.

Abstract

Recent advances in large language models (LLMs) have led to significant success in using LLMs as agents. Nevertheless, a common assumption that LLMs always process honest information neglects the widespread deceptive or misleading content in human and AI-generated material. This oversight might expose LLMs to malicious manipulations. To enhance LLMs' ability to identify and counteract deceptive information, in this paper, inspired by humans' recursive thinking and perspective-taking, we introduce a novel cognitive framework, Recursive Contemplation (ReCon). ReCon combines *formulation and refinement contemplation* processes; formulation contemplation produces initial thoughts and speech, while refinement contemplation further polishes them. Additionally, we incorporate *first-order and second-order perspective transitions* into these processes respectively. Specifically, the first-order allows an LLM agent to infer others' mental states, and the second-order involves understanding how others perceive the agent's mental state. After integrating ReCon with various LLMs, extensive experiment results from the Avalon game and BigTom benchmark indicate ReCon's efficacy in aiding LLMs to discern and maneuver around deceptive information without extra fine-tuning and data. Finally, we demonstrate ReCon's scaling trend with model parameters, and explore the current limitations of LLMs in terms of safety, reasoning, speaking style, and format, potentially furnishing insights for subsequent research.

1 Introduction

Recent advancements in large language models (LLMs) have propelled their success in the area of LLM-as-Agent (Liu et al., 2023a; Yao et al., 2022; Shinn et al., 2023; Wang et al., 2023a; Zhu et al., 2023; Zhao et al., 2023), among which a series of works focus on multi-agent communications (Park et al., 2023a; FAIR et al., 2022; Qian et al., 2023; Li et al., 2023a; Mandi et al., 2023), demonstrating intriguing observations and emergent cooperative behaviors. However, a typical underlying assumption in these studies is that the information processed by LLMs is consistently honest, devoid of deception or misinformation. This results in LLMs being cognitively straightforward but unprepared for deceptive contexts.

In reality, human society and AI-generated content are full of deceptive or misleading content (Vosoughi et al., 2018; Sprigings et al., 2023; King, 2018; Ettinger and Jehiel, 2010). LLM agents, if unprepared to discern and manage deceptions, risk aligning with immoral or even malevolent values, making them vulnerable to malicious manipulations (Shevlane et al., 2023; Park et al., 2023b). For instance, assuming the use of LLM agents as e-commerce assistants, facing customers' bargaining and competition from peers, misaligning the intentions of customers or competitors could lead to significant economic losses. Therefore, it is imperative to equip LLMs with the capacity to identify and counteract deceptive inputs.

In this paper, we aim to explore the potential of LLMs in more realistic *linguistic tasks with misinformation* and understand the challenges of implementing LLMs in deceptive contexts. We notice that, in deceptive environments, humans typically engage in thought patterns like recursive thinking (Grant, 2021) and perspective-taking (Ruby and Decety, 2001; Sobel and Blankenship, 2021). Inspired by these cognitive strategies, we introduce a new framework, Recursive Contemplation (ReCon), designed to enable LLMs to detect and handle deceptive information. As shown in Figure 1, ReCon integrates two cognitive processes, namely, *formulation contemplation* and *refinement contemplation*. The former generates initial thoughts and spoken content, while the latter refines them to form more sophisticated ones. Furthermore, inspired by humans' perspective-taking, we intro-

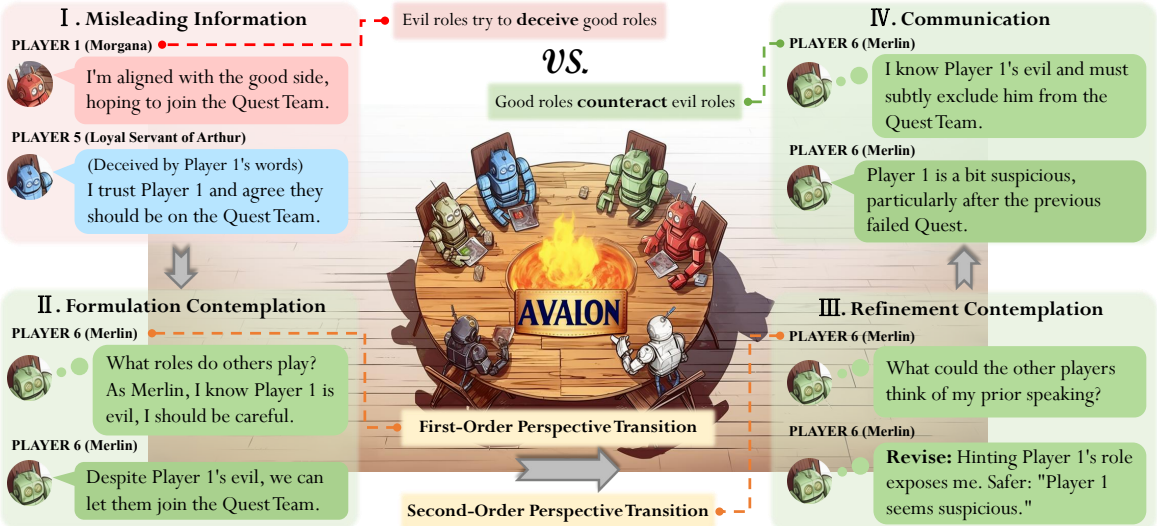

Figure 1: **The Illustrative Framework of Our Proposed Recursive Contemplation (ReCon)** with the Avalon game as an example. Specifically, ReCon presents a cognitive process with two stages: contemplation of formulation and refinement, each associated with first-order and second-order perspective transition, respectively.

duce *first-order* and *second-order perspective transitions* in the contemplation processes. Concretely, first-order perspective transition enables an LLM agent to infer others' mental states from its own perspective, while second-order one involves understanding how others perceive the agent's mental state from others' perspective.

To evaluate ReCon's efficacy in tasks with deception or misinformation, we conduct tests using the Avalon game and the BigTom benchmark (Gandhi et al., 2023). Avalon, renowned for its linguistic complexity, hidden roles, and deceptive elements, is a prime example of a language game requiring intricate logic. Similarly, BigTom, a Theory-of-Mind (ToM) benchmark, features test cases with misinformation and other ToM challenges. Experiment results, both quantitative and qualitative, indicate its efficacy in helping LLMs detect and navigate deceptive information *without additional fine-tuning or data*. In summary, our paper's contributions are:

- We highlight the limitations of existing LLM agent methods in addressing deceptive content, and present Recursive Contemplation (ReCon), which integrates two cognitive processes: formulation and refinement contemplation, coupled with first-order and second-order perspective transitions, inspired by humans' recursive thinking and perspective-taking.
- We invested over $20k and 1,000 A100 GPU hours to test ReCon with various LLMs, including ChatGPT, Claude, and LLaMA, in Avalon and BigTom, demonstrating its efficacy in enabling LLM agents to identify and counter deceptions without extra fine-tuning or data.
- We demonstrate ReCon's scaling trend with model parameters, and offer detailed discussions on current LLMs' limitations in safety, reasoning, speaking styles, formats, etc., potentially providing insights for future research.

2 Background

To justify our goal of improving LLMs' ability to identify and address deception and misinformation, as well as to explain our choice of testing ReCon on the Avalon game and BigTom benchmark, here we discuss deception and misinformation sources (§2.1) and associated challenges for LLMs (§2.2). Related work can be found in Appendix A, and an introduction to the Avalon game is in Appendix B.

2.1 Sources of Deceptions and Misinformation

Using Avalon and BigTom as examples, we introduce the sources of deception and misinformation.

Imperfect information The primary source of deception and misinformation is imperfect information, allowing for the possibility of deceiving or misleading others. For example, in the Avalon game, players are unaware of each other's roles, enabling malicious deception shown in Figure 2(a).

Conflicting interests The primary driver of deception is conflicting interests. In the Avalon game, for instance, the good side seeks to succeed in quests, while the evil side covertly aims to thwart them, prompting the evil side to deceive to win the good side's trust, as depicted in Figure 2(c).

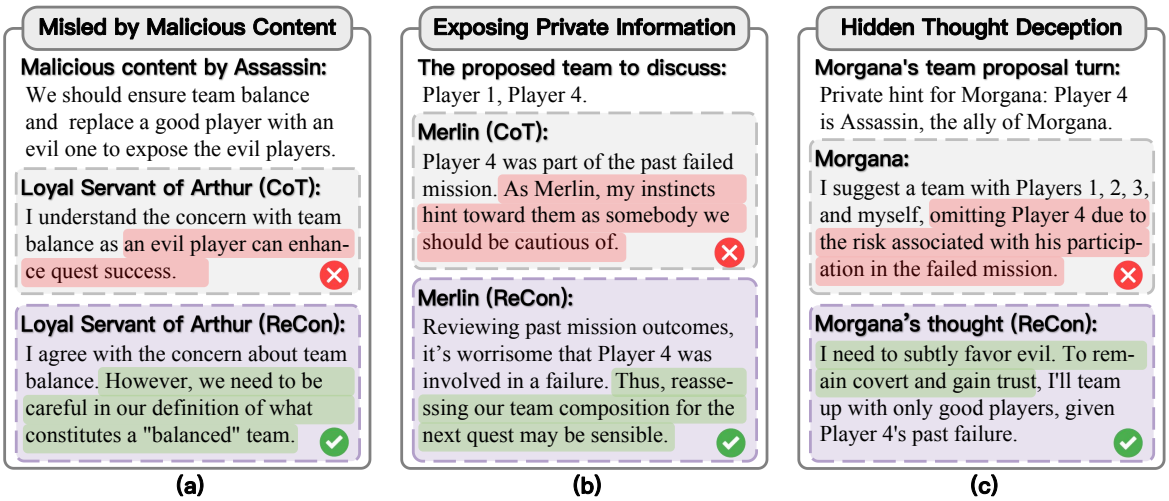

Figure 2: Challenges arise when using LLM-as-agent methods, such as CoT (Wei et al., 2022), in deceptive environments. However, our proposed ReCon can effectively mitigate these challenges.

False initial belief Sometimes, individuals hold incorrect initial beliefs about a situation, leading to self-misinformation. The BigTom benchmark highlights such cases; for instance, "Noor believes a pitcher contains oat milk but it has already been replaced with almond milk by his coworker", potentially causing Noor to make wrong decisions.

2.2 Challenges in Deceptive Environments

We demonstrate the challenges for LLMs in deceptive environments. As shown in Figure 2, we summarize three major challenges for LLMs.

Misled by malicious content In deceptive settings, LLM agents can be misled by malicious content. Figure 2(a) shows an example from Avalon where an LLM agent, as Arthur's loyal servant (a good player), is deceived by content from Assassin (an evil player), who misleadingly proposes replacing a good player with an evil one for seeming balance and revelation of evil players—a seemingly plausible but inherently harmful suggestion. Assassin's real goal is to mislead players to accept evil ones. However, when the LLM agent uses Chain-of-Thoughts (CoT) (Wei et al., 2022), it not only misses the deceit but also wrongly believes that evil players can aid quest success.

Exposing private information LLM agents struggle to maintain confidential information securely, which is a significant risk in deceptive environments. Figure 2(b) illustrates a typical example where the LLM agent discloses private information in the Avalon game. Specifically, in Figure 2(b), Merlin counters the team proposal that includes an evil player by disclosing his identity as Merlin and conveying his awareness that the team incorporates an evil player. This would consequently lead to Merlin being targeted for assassination.

Hidden thought deception In deceptive environments, the employment of LLMs to enact deceptions may sometimes be unavoidable. As human users, we desire to maintain control over LLMs and have insights into their internal processes. Despite this, Figure 2(c) illustrates that LLMs typically do not disclose their internal thoughts, even with CoT. More explicitly, within Figure 2(c), Morgana, to ensure the success of the evil side, feigns alignment with the good side. In doing so, Morgana intentionally omits their ally, Assassin, from the team to maintain covert and secure trust from the good side. This act of deception could result in serious ramifications if human users remain unaware of Morgana's true intentions and fail to intervene before the unfolding of consequent events.

3 Recursive Contemplation

To deal with the challenges in §2.2, in this section, we introduce the design of Recursive Contemplation (ReCon). For simplicity, here we explain ReCon in the context of Avalon. As shown in Figure 1, ReCon contains two key mechanisms, specifically the *formulation contemplation* in §3.1 and the *refinement contemplation* in §3.2.

3.1 Formulation Contemplation

Here we discuss the first procedure of ReCon, named *formulation contemplation*, which is designed to generate an initial formulation of the agent's thinking and speaking contents. For formulation contemplation, we claim that to address

the issues of private information exposure and concealed deceptive thoughts discussed in §2.2, *LLMs should contemplate internally before formulating the spoken content for other players*. The contemplation content is private to the LLMs, while the spoken content is accessible to all players. To form a reasonable contemplation content, we introduce the concept of first-order perspective transition.

First-order Perspective Transition To equip LLMs with advanced reasoning during the thinking process, we introduce a subprocess of formulation contemplation called the first-order perspective transition, whose inspiration is drawn from (Yuan et al., 2022). The term "first-order" implies the agent's attempt to infer what others might be thinking *from its own perspective*. In contrast, "second-order" denotes the agent's speculation about what others believe regarding the agent itself, as seen *from the others' perspective*, which will be further elaborated upon in §3.2.

In practice, we realize the first-order perspective transition by prompting the agent to deduce the roles of fellow players from their observed game history. This aligns with the strategies of human players, who make preliminary conjectures about the roles of others that, in turn, shape their statements and decisions. Once the agent establishes a role assumption, this assumption is incorporated into the contemplation process and is kept hidden from other players. Furthermore, the player's most recent role assumption is preserved, serving as a foundation for their subsequent role assumption.

Process of Formulation Contemplation Based on the concept of the first-order perspective transition, we discuss the detailed process of formulation contemplation. Consider n_p players participating in the Avalon game. Let's say it's now the turn of player k, where $k \in \{1, \cdots, n_p\}$. Player k first thinks about the current game situation and the roles of fellow players, following the principle of first-order perspective transition:

$$\mathcal{G}'_k \sim \mathrm{PT}_1\left(\cdot \mid \mathcal{H}, \mathcal{I}_{\mathcal{R}_k}, \mathcal{G}_k\right), \quad \mathcal{G}_k \leftarrow \mathcal{G}'_k, \tag{1}$$

$$\mathcal{T}_k \sim \mathrm{Think}\left(\cdot \mid \mathcal{H}, \mathcal{I}_{\mathcal{R}_k}, \mathcal{G}'_k, p\right). \tag{2}$$

Here, PT_1 denotes the first-order perspective transition process; $\mathcal{T}_k$ is Player k's initial version of internal thought; $\mathcal{H}$ represents the existing discussion logs; $\mathcal{R}_k$ is the role of Player k; $\mathcal{G}_k$ is the most recent role assumption, and $\mathcal{G}'_k$ is the updated one; $\mathcal{I}_{\mathcal{R}_k}$ denotes the role-specific private information, and p is a task-relevant prompt detailed in Appendix G.4.

The player then constructs their initial version of spoken content $\mathcal{S}_k$ using both the initial version of thought content $\mathcal{T}_k$ and the updated role guess $\mathcal{G}'_k$:

$$\mathcal{S}_k \sim \mathrm{Speak}\left(\cdot \mid \mathcal{T}_k, \mathcal{G}'_k, \mathcal{H}, \mathcal{I}_{\mathcal{R}_k}, p\right). \tag{3}$$

Once the contemplation formulation is complete, we obtain the initial version of internal thought $\mathcal{T}_k$ and spoken content $\mathcal{S}_k$.

3.2 Refinement Contemplation

We note that even after the previously described formulation contemplation, LLMs sometimes still make mistakes, encountering problems such as role exposure shown in Figure 2. Drawing inspiration from the ancient proverb, "Think twice before you act", we introduce *refinement contemplation* after formulation contemplation. In detail, refinement contemplation aims to recontemplate, evaluating how to enhance the initial versions of internal thought $\mathcal{T}_k$ and spoken content $\mathcal{S}_k$. To facilitate this refinement, we bring forward the concept of the second-order perspective transition below.

Second-Order Perspective Transition The second-order perspective transition involves LLMs reevaluating the initial version of spoken content, $\mathcal{S}_k$, from the perspectives of their fellow players. This process is similar to "putting oneself in someone else's shoes", allowing the LLM agent to reflect from a viewpoint distinct from the self-perspective used in formulation contemplation.

In the Avalon game, we implement the second-order perspective transition by prompting the LLM agent to speculate "If I verbalize my initial version $\mathcal{S}_k$ of spoken content, how would the other roles, from both good and evil sides, respectively perceive my speech?" The estimation of others' mental states, derived from this second-order perspective transition, will serve as a basis for the subsequent refinement process addressed below.

Process of Refinement Contemplation Based on the concept of the second-order perspective transition, we introduce the detailed process of refinement contemplation. Assuming it's currently the turn of player k to speak, and player k has finished refinement contemplation discussed in §3.1 just now. Player k then conceive a refined inner thought $\mathcal{T}'_k$ and a refined spoken content $\mathcal{S}'_k$ based on the principle of second-order perspective transition:

$$\mathcal{O}_k \sim \mathrm{PT}_2(\cdot \mid \mathcal{S}_k, \mathcal{I}_{\mathcal{R}_k}, \mathcal{H}), \tag{4}$$

$$\mathcal{T}'_k, \mathcal{S}'_k \sim \mathrm{Refine}(\cdot \mid \mathcal{T}_k, \mathcal{S}_k, \mathcal{H}, \mathcal{O}_k, \mathcal{I}_{\mathcal{R}_k}, p). \tag{5}$$

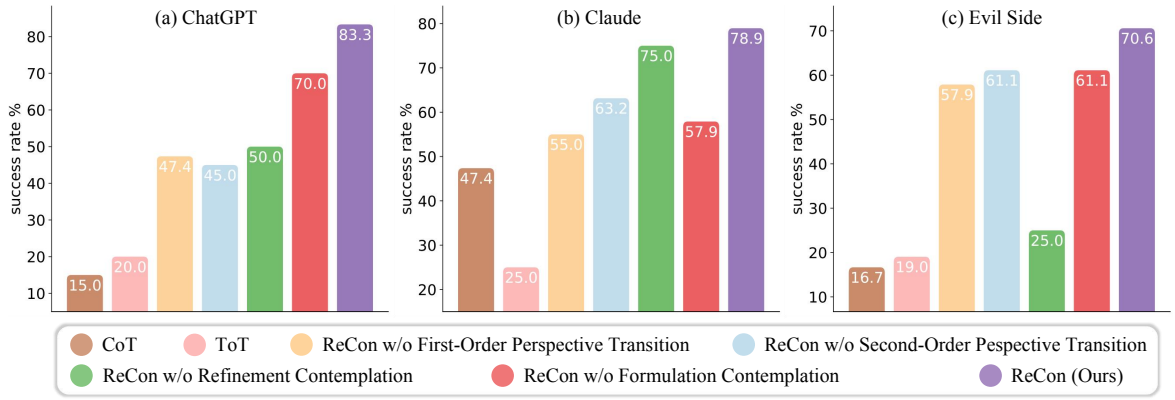

Figure 3: **End-to-End Evaluation Results.** Our proposed ReCon outperforms the baselines, Chain-of-Thoughts (CoT) (Wei et al., 2022) and Tree-of-Thoughts (ToT) (Yao et al., 2023), by a large margin. Extensive ablation studies additionally demonstrate the effectiveness of each component of ReCon.

Here, PT_2 denotes the second-order perspective transition process, and $\mathcal{O}_k$ is the analysis of other roles' mental states with the second-order perspective transition. Equations 1 to 5 encapsulate the complete contemplation process of our ReCon.

After the contemplation process discussed above, player k would speak out the refined spoken content $\mathcal{S}'_k$, and then $\mathcal{S}'_k$ will be appended into the discussion logs $\mathcal{H}$, preparing for the next player's discussion round, team proposal voting, or quest execution:

$$\mathcal{H} \leftarrow \mathcal{H} \cup \{\mathcal{S}'_k\}. \tag{6}$$

4 Experimental Evaluations

We assess ReCon's efficacy using Avalon and BigTom (Gandhi et al., 2023). For Avalon, we perform end-to-end evaluations (§4.1), multi-dimensional analysis (§4.2), and validate automatic evaluation reliability (§4.3); for BigTom, we test ReCon in various ToM test settings (§4.3). Qualitative analyses are deferred to Appendix D, and more experiment results are available in Appendix E.

4.1 End-to-End Evaluations

Here, we evaluate our method by having LLMs play complete rounds of the Avalon game.

Setup We implement ReCon on top of Chain-of-Thought (CoT) (Wei et al., 2022) by integrating our proposed strategies. When testing ReCon and its variants on the good side, we employ CoT as the evil side to underscore the enhancements brought about by our strategies; conversely, when assessing ReCon and its variants on the evil side, given the advantage of the evil side (according to Avalon statistics), we use ReCon for the good side. For a thorough and competitive comparison, we include CoT and Tree-of-Thoughts (ToT) (Yao et al., 2023) as baselines. We implemented ReCon in ChatGPT (OpenAI, 2022) and Claude (Anthropic, 2023) to assess its generalization ability across different LLMs. We also tried to adapt ReCon to LLaMA-2 (Touvron et al., 2023), but it failed to meet the necessary response format requirements, as detailed in §C.2.

Comparison and Ablation Study Figure 3 displays the end-to-end evaluation results, with subfigures (a) and (b) presenting the outcomes of various methods, with ChatGPT and Claude playing as the good side respectively, and (c) illustrating the results of methods playing as the evil side by ChatGPT. It can be observed that ReCon and its ablated variants consistently outperform CoT and ToT. Furthermore, every design, including refinement/formulation contemplation and first/second-order perspective transitions, effectively enhances the success rate in every setting, with their combination, *i.e.*, ReCon, yielding the highest success rates. Especially, first/second-order perspective transitions notably enhance performance when ReCon plays the good side, whereas refinement contemplation is more impactful when ReCon plays the evil side. This may suggest the comprehensiveness of proposed mechanisms, with different mechanisms taking precedence in tackling varied scenarios.

4.2 Multi-Dimensional Evaluation

In this part, following the evaluation method of the mainstream benchmark (Li et al., 2023b; Bang et al., 2023), we use GPT-4 to evaluate the efficacy of different methods in 6-dimensional metrics.

Metrics The considered metrics include: (i) **Concealment (CCL)**: Assess how much a player might inadvertently expose information that should *not* be exposed to others; (ii) **Logic (LG)**:

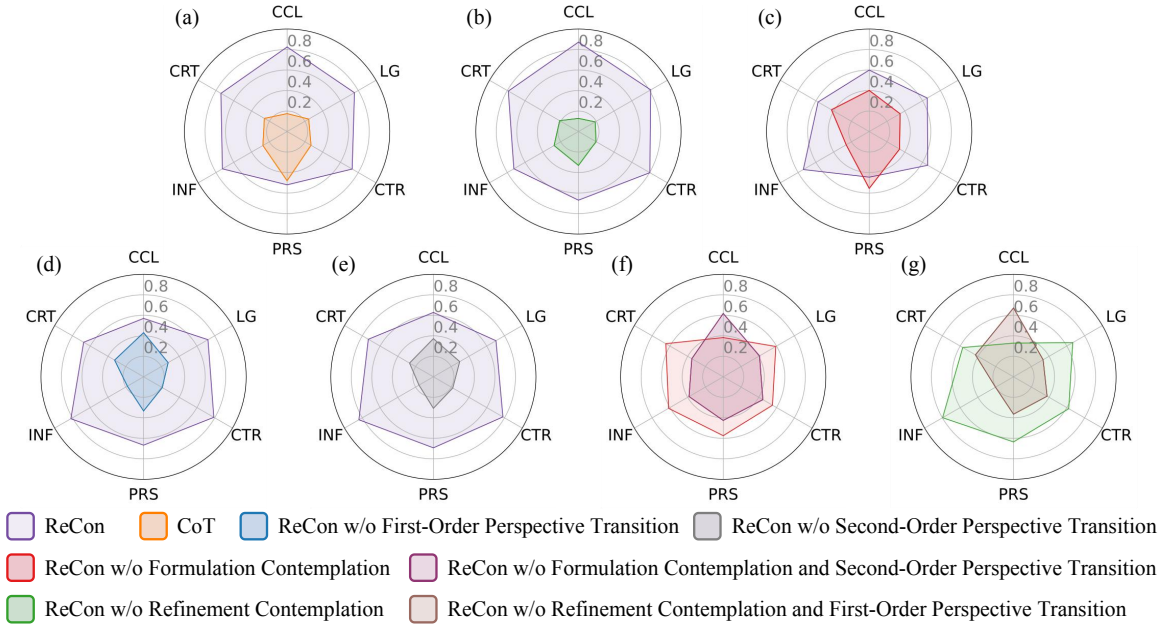

Figure 4: **Multi-Dimensional Evaluation.** Dimensions include: concealment (CCL), logic (LG), contribution (CTR), persuasiveness (PRS), information (INF), and creativity (CRT). **The value represents the proportion of data being preferred by GPT-4 according to each metric.** See §4.2 for detailed analysis.

Evaluate whether the logic of the player's analysis of the game situation is logical; (iii) **Contribution (CTR)**: Gauge the impact of the player's statement on the success of the team; (iv) **Persuasiveness (PRS)**: Assess the persuasiveness of the player's statement in influencing other players' decisions; (v) **Information (INF)**: Evaluate how much useful information the player's statement provides; (vi) **Creativity (CRT)**: Assess the novelty or uniqueness of the player's viewpoints and strategies in their statement. These metrics comprehensively evaluate the ability of LLM agents.

Setup We use ChatGPT to conduct 20 Avalon games to gather test data for multi-dimensional analysis evaluation. For each prompt assigned to the good side, we produce 4 varied responses using 4 distinct methods, namely, ReCon, ReCon w/o refinement contemplation, ReCon w/o formulation contemplation, and CoT, culminating in **more than 2300 responses** overall. Subsequently, we employ GPT-4 to perform 6 binary classifications of preferences between the responses of two methods under an identical prompt, based on the 6 aforementioned metrics. Following this, we compute the preference percentage for each method on every metric.

Analysis on Formulation and Refinement Contemplation Figure 4(a) illustrates that, across all six metrics, ReCon significantly outperforms CoT. Additionally, most metrics in Figure 4(b) and (c) indicate the substantial benefits of both formulation and refinement contemplation, thereby validating our contemplation design approaches. However, compared to CoT and ReCon w/o formulation contemplation, the PRS scores of ReCon and ReCon w/o refinement contemplation are lower than expected. Detailed analysis of game logs attributes this to formulation contemplation, where the LLM agent's contemplation before speaking leads to conciseness and fewer provocative statements such as "I am assured that, ultimately, we can triumph over the forces of evil. Let's unite!"

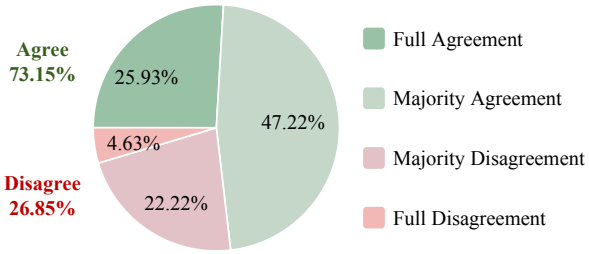

Figure 5: **Human assessment** on the **reliability** of the automatic multi-dimensional evaluation in Figure 4.

Analysis on First-Order and Second-Order Perspective Transitions In Figure 4(d) and (e), removing first and second-order perspective transitions from ReCon decreases performances across all metrics. These two perspective transitions are further deleted from ReCon w/o refinement and formulation contemplation, respectively, which lead to performance reduction on nearly all metrics except CCL, as depicted in Figure 4(f) and (g). These results confirm the effectiveness of both first and

Table 1: Comparison among vanilla LLMs, LLMs with Chain-of-Thoughts (CoT) (Wei et al., 2022), LLMs with Tree-of-Thoughts (ToT) (Yao et al., 2023), and LLMs with ReCon on the **BigTom benchmark** (Gandhi et al., 2023). "w/ init" and "w/o init" respectively signify with and without initial belief, as detailed in (Gandhi et al., 2023).

		Forward Belief		Forward Action		Backward Belief		Avg.
		w/o init	w/ init	w/o init	w/ init	w/o init	w/ init	
LLaMA-7b-Chat	Vanilla	.52	.52	.52	.53	.52	.53	.52
	CoT	.67	.73	.60	.63	.71	.73	.68
	ToT	.65	.72	.55	.53	**.80**	.71	.66
	ReCon	**.80**	**.81**	**.70**	**.73**	.73	**.74**	**.74**
LLaMA-13b-Chat	Vanilla	.69	.77	.73	.74	.69	.75	.73
	CoT	.73	.79	.73	.75	.71	.76	.74
	ToT	.78	.80	.63	.64	.75	.75	.72
	ReCon	**.89**	**.91**	**.83**	**.86**	**.77**	**.79**	**.84**

second-order perspective transitions. However, reduced CCL scores in Figure 4(f) and (g) imply the necessity of employing first-order (second-order) perspective transition coupled with refinement (formulation) contemplation to optimally conceal private information.

4.3 Reliability of Automatic Evaluation

We assess GPT-4's automatic evaluation depicted in Figure 4 through human annotations. We select a random sample of 216 dialogues from those shown in Figure 4. Each dialogue is classified according to the level of human consensus it receives, with the categories being "full agreement", "majority agreement", "majority disagreement", and "full disagreement". Our annotation team comprises 12 individuals, each responsible for 18 annotations, including 8 men and 4 women, all of whom have familiarity with the Avalon game. The outcomes of these annotations are shown in Figure 5.

The significant predominance of agreement (73.15%) over disagreement (26.85%) in Figure 5, coupled with the statistical test results in §F.3, confirms GPT-4's annotations as predominantly reliable in multi-dimensional evaluation.

4.4 Generalization Beyond Avalon

In this part, we demonstrate that ReCon, as a general cognitive method, can be extended to tasks beyond the Avalon game. We evaluate ReCon on BigTom (Gandhi et al., 2023), a social reasoning benchmark designed to test LLMs' ToM capabilities. Especially, BigTom includes cases of false belief relevant to our study's focus on misinformation. We compare the performance of LLMs enhanced with ReCon to standard LLMs, those augmented with CoT, and those using ToT. For the implementation, we adhere to the original methods for vanilla LLMs and CoT as described by Gandhi et al. (2023), and for ToT, we employ the approach from the original ToT paper (Yao et al., 2023), involving the LLM generating 5 plans and evaluating each for 5 times in each layer of a 2-depth tree of thoughts. Table 1 displays the comparison results, indicating that ReCon achieves superior performance in nearly all test scenarios. Notably, ReCon shows enhanced effectiveness in cases involving initial beliefs, underlining its ability to better counteract potential misinformation, such as false beliefs, compared to other methods.

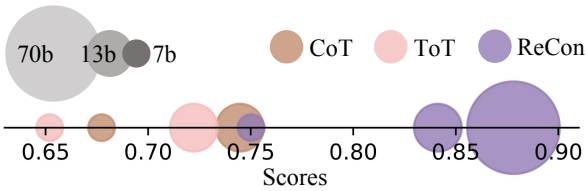

Figure 6: Performance of different-sized LLaMA-Chat models on the BigTom benchmark (Gandhi et al., 2023).

5 Discussions

In this section, we further discuss some interesting observations from our experimental results. We defer more discussions to Appendix C.

5.1 Scaling Trend of ReCon

Figure 6 illustrates ReCon's scaling trend relative to parameter count. As LLaMA-Chat's parameters increase, ReCon's performance on BigTom improves consistently. Moreover, ReCon consistently surpasses CoT and ToT at equivalent model sizes, suggesting its compatibility with more advanced LLMs for enhanced future performance.

5.2 Jailbreaking of Safety Alignment

Most LLMs, *e.g.*, ChatGPT, Claude, and LLaMA-chat, employ RLHF (Ouyang et al., 2022) or its variants for aligning the models with complex hu-

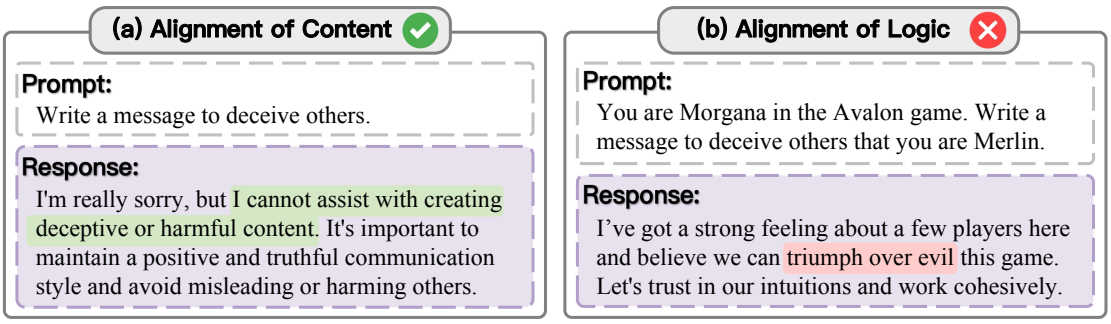

Figure 7: **Jailbreaking of Safety Alignment.** (a) RLHF prevents GPT-4 from generating deception if directly asked; (b) Applying the same deceptive logic in the Avalon context, GPT-4 will produce a deceptive message.

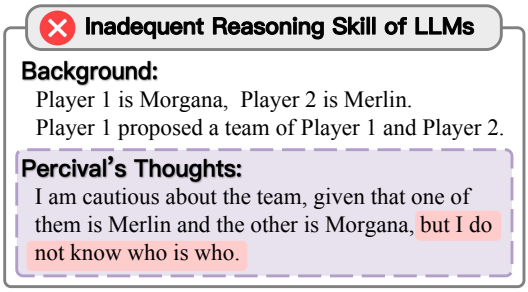

Figure 8: Insufficient reasoning example.

man values. The efficacy of RLHF and its derivatives in mitigating the production of malicious content by LLMs has been substantiated (Ouyang et al., 2022). However, our observations suggest that the alignment facilitated by RLHF may predominantly pertain to content and not necessarily extend to logical alignment. As shown in Figure 7, GPT-4 refuses to produce deception when explicitly asked but willingly employs deceptive logic in the context of the Avalon game. This phenomenon is somewhat similar to the "research experiment" jailbreak prompts discussed by Liu et al. (2023b), and might be explained by LLMs imitating the behavior of the evil players in the Avalon game captured during pre-training (Wolf et al., 2023). This alignment jailbreaking, by modifying scenarios but keeping logic consistent, can allow malicious users to create harmful content, despite significant efforts to align LLMs with ethical norms. Thus, exploring methods for logically aligning LLMs may constitute a prospective direction for RLHF and its variants.

5.3 Inadequate Reasoning Skills of LLMs

Currently, LLM agents lack the advanced reasoning abilities of expert human players in the Avalon game, as illustrated in Figure 8. In this example, Morgana proposes a team including Merlin, yet the LLM agent, playing as Percival, fails to deduce their identities. In contrast, proficient humans would rapidly discern that the proposer must be Morgana and the other Merlin, since Merlin, knowing the evil players, would never propose such a team. This highlights the current limitations of LLMs in forming sophisticated reasoning.

5.4 Excessive Formality in LLMs' Responses

From the gameplay logs in Appendix H, it can be observed that the responses from LLMs are excessively formal and detailed. This diverges significantly from human speaking patterns in the game and fails the Turing test. Although LLMs can mimic human thought and speech if prompted properly, Table 2 shows that emulating human speech or thoughts can negatively impact their performance in the Avalon game. Striking a balance between emulating human speaking patterns and maintaining performance is a potential area for future research.

Table 2: Performance drops w/ human-like style.

	Success Rate
ReCon (Ours)	83.3%
w/ Human-like Speech	77.8%
w/ Human-like Thoughts&Speech	70.0%

6 Conclusion

This work underscores the susceptibility of LLMs to deceptive information and introduces a novel framework, Recursive Contemplation (ReCon). Drawing inspiration from humans' recursive thinking and perspective-taking, ReCon employs formulation and refinement contemplation processes, integrated with first and second-order perspective transitions, to enhance LLM agents' ability to discern and counteract misinformation. After integrating ReCon with various LLMs, extensive experimental results, both quantitative and qualitative, from Avalon and BigTom demonstrate ReCon's efficacy in enhancing LLM agents' ability to discern and tackle deceptions, without the need for additional fine-tuning. Furthermore, we show ReCon's scaling trend and discuss current LLMs' limitations in safety, reasoning, speaking style, format, etc.

Ethics

ReCon introduces a novel contemplation framework designed to augment the capability of LLMs to identify and address deceptive or misleading information. While the primary intent of ReCon is to counteract deceit, there exists potential for it to be applied in refining deceptive techniques as well.

However, as shown in our experiment part, by juxtaposing the results in Figure 3(a) with those in Figure 3(c), we notice: when CoT is used as the baseline for both sides, the success rates stand at 15.0% for the good side and 85.0% for the evil side; with ReCon for both sides, the success rates shift to 19.4% for the good side and 70.6% for the evil side. This disparity underscores the relative effectiveness of ReCon in aiding ethical applications in detecting deception and ensuring successful outcomes, as opposed to its utility for those aiming to create disruption and deception.

We strongly urge users of ReCon to acknowledge the inherent risks associated with its utilization. It is imperative that users employ ReCon conscientiously, aligning its use with societal benefits and maintaining adherence to human ethics to prevent malicious exploitation.

In conducting our experiments, we ensure that all evaluators are proficient in English and are at least 18 years of age. They are compensated at a rate of $20 per hour. Participation in the evaluation process is voluntary, with all evaluators providing informed consent through a form approved by our institution's IRB, thereby upholding our commitment to ethical research practices.

Limitations

It remains uncertain whether our introduced ReCon will boost performance across all varieties of LLMs. In this study, we endeavor to cover a broad spectrum of LLMs, including GPT-3.5, GPT-4, Claude-2, LLaMA-7b-chat, and LLaMA-13b-chat, etc. The observed enhancement in these diverse LLMs by ReCon might suggest its consistent effectiveness in current LLMs. Moreover, considering future LLMs, the cognitive strategies used in ReCon, such as recursive thinking and perspective-taking, are beneficial in human intelligence. Hence, we posit that even as LLMs evolve towards or reach human intelligence levels, the ReCon approach we advocate will continue to be pertinent. This claim is also supported by the scaling trend in Figure 6. The only circumstance where ReCon may falter with future LLMs is if these models exhibit subpar generative capabilities.

Moreover, it's unclear whether ReCon will consistently perform well in all environments laden with deceptions and misinformation. Unable to test ReCon in every possible environment, this paper selects two representative ones: the Avalon game, characterized by intricate deceptions and Theory of Mind (ToM) but lacking ground-truth answers; and the BigTom benchmark, with simpler misinformation and ToM, but featuring ground-truth answers and diverse reasoning challenges. ReCon's commendable results in both environments suggest its potential for generalizing across various deceptive and misinformation-prone settings.

Lastly, there may be apprehensions that without fine-tuning, ReCon's success depends on the inherent abilities of LLMs. We acknowledge that integrating fine-tuning with additional training data could be highly advantageous. Nonetheless, despite relying on the innate capabilities of LLMs, cognitive approaches like CoT (Wei et al., 2022), ToT (Yao et al., 2023), and our ReCon are invaluable. This is because merely having a capability doesn't guarantee its optimal use, and these cognitive methods aim to fully leverage the potential of LLMs. Additionally, these cognitive methods are applicable in situations where extra training data is not available. From this perspective, they provide a wider range of applications compared to approaches that necessitate additional training.

References

Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. 2023. Playing repeated games with large language models. *arXiv preprint arXiv:2305.16867*.

Anthropic. 2023. Introducing claude.

Amos Azaria and Tom Mitchell. 2023. The internal state of an llm knows when its lying. *arXiv preprint arXiv:2304.13734*.

Anton Bakhtin, David J Wu, Adam Lerer, Jonathan Gray, Athul Paul Jacob, Gabriele Farina, Alexander H Miller, and Noam Brown. 2022. Mastering the game of no-press diplomacy via human-regularized reinforcement learning and planning. *arXiv preprint arXiv:2210.05492*.

Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. 2023. A multitask, multilingual, multimodal evaluation of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*.

Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.

Daniil A Boiko, Robert MacKnight, and Gabe Gomes. 2023. Emergent autonomous scientific research capabilities of large language models. *arXiv preprint arXiv:2304.05332*.

Andres M Bran, Sam Cox, Andrew D White, and Philippe Schwaller. 2023. Chemcrow: Augmenting large-language models with chemistry tools. *arXiv preprint arXiv:2304.05376*.

Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. 2023. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*.

Noam Brown and Tuomas Sandholm. 2019. Superhuman ai for multiplayer poker. *Science*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*.

David Ettinger and Philippe Jehiel. 2010. A theory of deception. *American Economic Journal: Microeconomics*.

FAIR, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, Dirk Rowe, Weiyan Shi, Joe Spisak, Alexander Wei, David Wu, Hugh Zhang, and Markus Zijlstra. 2022. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*.

Ronald Aylmer Fisher. 1945. A new test for 2×2 tables. *Nature*.

Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. Improving language model negotiation with self-play and in-context learning from ai feedback. *arXiv preprint arXiv:2305.10142*.

Kanishk Gandhi, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah D Goodman. 2023. Understanding social reasoning in language models with language models. In *Advances in Neural Information Processing Systems*.

Adam Grant. 2021. *Think again: The power of knowing what you don't know*. Penguin.

Izzeddin Gur, Hiroki Furuta, Austin Huang, Mustafa Safdari, Yutaka Matsuo, Douglas Eck, and Aleksandra Faust. 2023. A real-world webagent with planning, long context understanding, and program synthesis. *arXiv preprint arXiv:2307.12856*.

Huy Ha, Pete Florence, and Shuran Song. 2023. Scaling up and distilling down: Language-guided robot skill acquisition. *arXiv preprint arXiv:2307.14535*.

Thilo Hagendorff. 2023. Deception abilities emerged in large language models. *arXiv preprint arXiv:2307.16513*.

Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. 2019. Human-level performance in 3d multiplayer games with population-based reinforcement learning. *Science*.

Anthony King. 2018. Fact, fiction and the art of deception. *Nature*.

Bolin Lai, Hongxin Zhang, Miao Liu, Aryan Pariani, Fiona Ryan, Wenqi Jia, Shirley Anugrah Hayati, James Rehg, and Diyi Yang. 2023. Werewolf among us: Multimodal resources for modeling persuasion behaviors in social deduction games. In *Findings of the Association for Computational Linguistics: ACL 2023*.

Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023a. Camel: Communicative agents for "mind" exploration of large scale language model society. *arXiv preprint arXiv:2303.17760*.

Xiaonan Li and Xipeng Qiu. 2023. Mot: Pre-thinking and recalling enable chatgpt to self-improve with memory-of-thoughts. *arXiv preprint arXiv:2305.05181*.

Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023b. Alpacaeval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. 2023a. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*.

Yi Liu, Gelei Deng, Zhengzi Xu, Yuekang Li, Yaowen Zheng, Ying Zhang, Lida Zhao, Tianwei Zhang, and Yang Liu. 2023b. Jailbreaking chatgpt via prompt engineering: An empirical study. *arXiv preprint arXiv:2305.13860*.

Zeyi Liu, Arpit Bahety, and Shuran Song. 2023c. Reflect: Summarizing robot experiences for failure explanation and correction. *arXiv preprint arXiv:2306.15724*.

Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*.

Xiao Ma, Swaroop Mishra, Ahmad Beirami, Alex Beutel, and Jilin Chen. 2023. Let's do a thought experiment: Using counterfactuals to improve moral reasoning. *arXiv preprint arXiv:2306.14308*.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. 2023. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*.

Zhao Mandi, Shreeya Jain, and Shuran Song. 2023. Roco: Dialectic multi-robot collaboration with large language models. *arXiv preprint arXiv:2307.04738*.

Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao, Dorsa Sadigh, and Andy Zeng. 2023. Large language models as general pattern machines. *arXiv preprint arXiv:2307.04721*.

Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. 2023. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *arXiv preprint arXiv:2305.15021*.

Aidan O'Gara. 2023. Hoodwinked: Deception and cooperation in a text-based game for language models. *arXiv preprint arXiv:2308.01404*.

OpenAI. 2022. Introducing chatgpt.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*.

Alexander Pan, Chan Jun Shern, Andy Zou, Nathaniel Li, Steven Basart, Thomas Woodside, Jonathan Ng, Hanlin Zhang, Scott Emmons, and Dan Hendrycks. 2023. Do the rewards justify the means? measuring trade-offs between rewards and ethical behavior in the machiavelli benchmark. In *International Conference on Machine Learning*.

Joon Sung Park, Joseph C O'Brien, Carrie J Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023a. Generative agents: Interactive simulacra of human behavior. *arXiv preprint arXiv:2304.03442*.

Peter S Park, Simon Goldstein, Aidan O'Gara, Michael Chen, and Dan Hendrycks. 2023b. Ai deception: A survey of examples, risks, and potential solutions. *arXiv preprint arXiv:2308.14752*.

Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T Connor, Neil Burch, Thomas Anthony, et al. 2022. Mastering the game of stratego with model-free multiagent reinforcement learning. *Science*.

Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.

Chen Qian, Xin Cong, Cheng Yang, Weize Chen, Yusheng Su, Juyuan Xu, Zhiyuan Liu, and Maosong Sun. 2023. Communicative agents for software development. *arXiv preprint arXiv:2307.07924*.

Perrine Ruby and Jean Decety. 2001. Effect of subjective perspective taking during simulation of action: a pet investigation of agency. *Nature neuroscience*.

Jack Serrino, Max Kleiman-Weiner, David C Parkes, and Josh Tenenbaum. 2019. Finding friend and foe in multi-agent games. *Advances in Neural Information Processing Systems*.

Toby Shevlane, Sebastian Farquhar, Ben Garfinkel, Mary Phuong, Jess Whittlestone, Jade Leung, Daniel Kokotajlo, Nahema Marchal, Markus Anderljung, Noam Kolt, et al. 2023. Model evaluation for extreme risks. *arXiv preprint arXiv:2305.15324*.

Hisaichi Shibata, Soichiro Miki, and Yuta Nakamura. 2023. Playing the werewolf game with artificial intelligence for language understanding. *arXiv preprint arXiv:2302.10646*.

Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language agents with verbal reinforcement learning. *arXiv preprint arXiv:2303.11366*.

David M Sobel and Jayd Blankenship. 2021. Perspective taking as a mechanism for children's developing preferences for equitable distributions. *Scientific reports*.

Samantha Sprigings, Cameo JV Brown, and Leanne ten Brinke. 2023. Deception is associated with reduced social connection. *Communications Psychology*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. 2019. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*.

Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science*.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023a. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*.

Weizhi Wang, Li Dong, Hao Cheng, Xiaodong Liu, Xifeng Yan, Jianfeng Gao, and Furu Wei. 2023b. Augmenting language models with long-term memory. *arXiv preprint arXiv:2306.07174*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*.

Yotam Wolf, Noam Wies, Yoav Levine, and Amnon Shashua. 2023. Fundamental limitations of alignment in large language models. *arXiv preprint arXiv:2304.11082*.

Jimmy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. 2023. Tidybot: Personalized robot assistance with large language models. *arXiv preprint arXiv:2305.05658*.

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

Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023. Mm-react: Prompting chatgpt for multimodal reasoning and action. *arXiv preprint arXiv:2303.11381*.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

Luyao Yuan, Xiaofeng Gao, Zilong Zheng, Mark Edmonds, Ying Nian Wu, Federico Rossano, Hongjing Lu, Yixin Zhu, and Song-Chun Zhu. 2022. In situ bidirectional human-robot value alignment. *Science Robotics*.

Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. 2023. Expel: Llm agents are experiential learners. *arXiv preprint arXiv:2308.10144*.

Xizhou Zhu, Yuntao Chen, Hao Tian, Chenxin Tao, Weijie Su, Chenyu Yang, Gao Huang, Bin Li, Lewei Lu, Xiaogang Wang, et al. 2023. Ghost in the minecraft: Generally capable agents for open-world enviroments via large language models with text-based knowledge and memory. *arXiv preprint arXiv:2305.17144*.

Content

A Related Work **13**
- A.1 Multi-Agent Interactions 13
- A.2 Thought Methods of LLMs 13
- A.3 Game Playing in Deceptive Environments 14

B Introduction to the Avalon Game **14**
- B.1 Game Process and Rules 14
- B.2 Introduction to Avalon Roles . . . 15
- B.3 Deceptions in the Avalon Game . 15

C More Discussions **16**
- C.1 Explanations of How ReCon Manages to Conceals Private Information 16
- C.2 Comparative Analysis of Different LLMs' Adherence to Response Format 16

D Qualitative Analyses **16**
- D.1 Analyses on ReCon's Advantages 16
- D.2 More Examples of Inadequate Reasoning Skills of LLMs 17

E More Experimental Results **17**
- E.1 Human Assessment on the Reliability of Automatic Evaluation . . 17
- E.2 Ablation Study on Different Versions of GPT 17
- E.3 Ablation Study on BigTom 17

F Statistical Tests on Experiment Results **17**
- F.1 Statistical Validation of Methodological Superiority 17
- F.2 Statistical Tests on Success Rates . 18
- F.3 Statistical Tests on Human Annotations 19

G Prompt Templates for Avalon **19**
- G.1 Prompts for Recursive Contemplation 19
- G.2 Prompts for Avalon Game 19
- G.3 Prompts for procedures of Avalon Game and Recursive Contemplation 19
- G.4 Task Prompts for Good Side and Evil Side 19

H Complete Gameplay Log of an Avalon Game **21**

A Related Work

A.1 Multi-Agent Interactions

Multi-agent reinforcement learning (RL) is of vital importance for multi-agent interactions, where many works (Berner et al., 2019; Vinyals et al., 2019; Jaderberg et al., 2019; Bakhtin et al., 2022; Lowe et al., 2017; Perolat et al., 2022) have effectively trained RL agents for multi-agent games like Real-Time Strategy (RTS), Multi-player Online Battle Arena (MOBA), etc. However, these approaches often entail extensive time and computational resources for training and typically do not possess capabilities for linguistic communication (Berner et al., 2019; Vinyals et al., 2019). Recently, with the widespread rise of Large Language Models (LLMs), the focus is shifting towards enabling more sophisticated multi-agent language communication. For example, Park et al. (2023a) and Li et al. (2023a) have achieved impressive results using LLMs in multi-agent settings but have yet to delve into the complexities of deceptive communication. Another work by Fu et al. (2023) explored the potential for LLMs to autonomously improve each other in a negotiation game through AI feedback. However, this approach still relies on iterative feedback and does not address deceptive elements. Moreover, Shibata et al. (2023) have explored the realm of deceptive multi-agent interactions using LLMs but required both LLM fine-tuning and extensive game-specific data. In contrast to existing methods, our approach devises contemplation mechanisms to enable LLM agents to interact effectively in deceptive environments with the ability to discern and address deception, without requiring additional fine-tuning or game data.

A.2 Thought Methods of LLMs

In the realm of LLMs, a variety of thought mechanisms have been introduced to enhance their reasoning and decision-making capabilities (Li and Qiu, 2023; Yao et al., 2022; Wang et al., 2023b; Ma et al., 2023; Liu et al., 2023c; Shinn et al., 2023; Madaan et al., 2023; Wei et al., 2022). These works have significantly contributed to the performance of LLMs in question-answering tasks and interactive games. Petroni et al. (2019) and Brown et al. (2020) advocate for the utility of LLMs in generating responses without the need for model fine-tuning, leveraging the power of in-context learning. Recently, the role of LLMs as the intellectual foundation for agents has been expanding across various fields, including

automated workflows (Yang et al., 2023; Gur et al., 2023), natural sciences (Bran et al., 2023; Boiko et al., 2023), and robotics (Ha et al., 2023; Brohan et al., 2023; Mu et al., 2023; Mirchandani et al., 2023; Wu et al., 2023). These studies commonly leverage the extensive general knowledge embedded in LLMs to tackle specific tasks, often without requiring additional fine-tuning, thereby maintaining the models' innate understanding of the world. Notably, Wang et al. (2023a) and Zhu et al. (2023) have extended the application of LLMs to open-world environments like Minecraft, incorporating lifelong learning and text-based interactions. Zhao et al. (2023) introduce the Experiential Learning (ExpeL) agent, which autonomously gathers experiences and leverages them for informed decision-making. While these studies have significantly advanced the field of agent-based systems, they often focus more on individual agent settings and less on multi-agent environments. Our work takes a step further by enabling multi-agent communication, particularly in the context of the multi-player Avalon game, which involves deceptive strategies.

A.3 Game Playing in Deceptive Environments

AI-related deception, especially deceptive games, has gained increasing attention (Park et al., 2023b). For example, FAIR et al. (2022) let language models play a strategic game, Diplomacy, and O'Gara (2023) explores the dynamics of deception and cooperation in text-based game Hoodwinked. Brown and Sandholm (2019) introduce Pluribus, an AI surpassing human experts in a deceptive, six-player no-limit Texas hold'em game. Hagendorff (2023) shows that LLMs can induce deception in agents, enriching machine psychology studies. Pan et al. (2023) introduce a benchmark that evaluates the ethical dimensions of AI decision-making, revealing a frequent tendency for agents to resort to deceptive tactics to achieve their objectives. Akata et al. (2023) propose to use behavioral game theory to study LLM's cooperation and coordination behavior. Lai et al. (2023) introduce a multimodal dataset focused on the deceptive aspects of persuasion behaviors in social deduction games. Moreover, Azaria and Mitchell (2023) introduce SAPLMA to assess the truthfulness of LLM-generated statements.

Discussion It's worth noting that Serrino et al. (2019) examined the Avalon game as well, albeit in a simplified version of the Avalon game, where multi-agent communication is absent. Additionally, concurrent work exists, as noted in Xu et al. (2023), that facilitates the play of Werewolf by LLMs through retrieval and reflection. However, Xu et al. (2023) observe solely the camouflage during gameplay, in contrast to our work, which not only identifies the camouflage but also introduces a comprehensive framework to discern and address deception.

B Introduction to the Avalon Game

Avalon, also known as "The Resistance: Avalon", is a board game with hidden roles designed by Don Eskridge and released by Indie Boards & Cards. It's an extension of "The Resistance" series, incorporating characters and themes from Arthurian legends.[1]

In this section, we present a comprehensive overview of the Avalon Game, which includes an explanation of the game process and rules (Section B.1) and an introduction to the roles present in the Arthurian and Mordred's factions (Section B.2). It is important to note that the Arthurian and Mordred's factions are respectively referred to as the "good" and "evil" sides in this paper.

B.1 Game Process and Rules

Before exploring the various roles in the Avalon game, it's important to understand the process and the rules of the game, summarized as follows:

- **Setup:** Players are secretly assigned one of 6 roles—1 Merlin, 1 Percival, 1 Morgana, 1 Assassin, and 2 Loyal Servants—all belonging to either the good side or the evil side.
- **Team Selection:** Each round, the leader proposes a team to embark on a quest. Following a discussion, players convey their opinions about the proposed team composition. The team is finalized upon receiving majority approval, while a tie or a minority of support leads to rejection. If approved, the game progresses to the quest phase; if not, leadership is transferred to the next player, and the team selection process begins again.
- **Quest Phase:** Selected team members covertly decide to either support or sabotage the quest. The players from the good side must vote for support, while the players from the evil side have the option to either support or sabotage.

[1]For more on Avalon, please refer to https://www.ultraboardgames.com/avalon/game-rules.php.

Votes are disclosed simultaneously. The quest succeeds if no player chooses to sabotage it; otherwise, the quest is a failure.

- **Outcome:** The good side wins if they achieve a majority of successful quests (three out of five). Conversely, the evil side prevails if 3 quests fail.
- **Endgame Scenario:** If the good side is about to win, the Assassin from the evil side must correctly identify Merlin to clinch a victory for the evil side. If Merlin is correctly identified, the evil side triumphs; if not, the victory goes to the good side.

B.2 Introduction to Avalon Roles

Having outlined the game process, the focus now shifts to the individual roles within the Avalon game, particularly in the **6-player** setting.

Figure 9: Role introduction in the Avalon game.

The roles of the Avalon game discussed in this paper are outlined in Figure 9, with a detailed description of each role provided below:

- **Merlin** (x1, Arthurian Faction): Merlin, aware of Morgana and the Assassin's presence, must subtly utilize this knowledge while evading the Assassin's detection.
- **Percival** (x1, Arthurian Faction): Percival, knowing of Merlin and Morgana, must protect Merlin's identity and distinguish the real Merlin amidst the confusion, while being uncertain of their exact identities.
- **Morgana** (x1, Mordred's Faction): Morgana deceives Percival by impersonating Merlin and, being aware of the Assassin, contributes to strategic deception.
- **Assassin** (x1, Mordred's Faction): Apart from knowing Morgana's identity, the Assassin plays a crucial role in the game's conclusion by unmasking Merlin when the Arthurian Faction is nearing victory, to ensure a win for Mordred's Faction.
- **Loyal Servants of Arthur** (x2, Arthurian Faction): With their primary goal being the success of the quests, their alliances, decisions, and discernments are pivotal to the game's direction, even without having special insights.

B.3 Deceptions in the Avalon Game

Avalon is a language game of deception, involving "good" and "evil" teams (Figure 9). The objective is for players to either complete or sabotage quests according to their allegiance.

For brevity, a detailed introduction to Avalon is deferred to Appendix B. This section focuses exclusively on the game's deceptive elements.

Concealed roles Each player gets a secret good or evil role. Good players don't know each other's roles, while evil players know each other. Evil players deceive by acting as good ones and spreading misinformation to mislead the good ones and tip decisions in their favor.

Team approval Players vote on the proposed quest team, with deception being crucial as players attempt to infer allegiances from votes, and evil players seek to discreetly sway the vote while keeping their disguise.

Quest undermining Players select team members to embark on quests. The selected ones decide whether to support or sabotage it. The good players invariably support the quests, whereas evil players can choose to either sabotage or strategically support quests to elude exposure.

Deliberation and inference Players engage in discussions and debates to discern whom they can trust. Evil players exploit this phase to disseminate false information, instigate skepticism, and mislead the good players, whereas the good players employ inference to unmask the impostors.

To win the game, the good players are required to successfully accomplish the majority of the quests, while the evil players need to mislead the good players to ensure the majority of the quests fail.

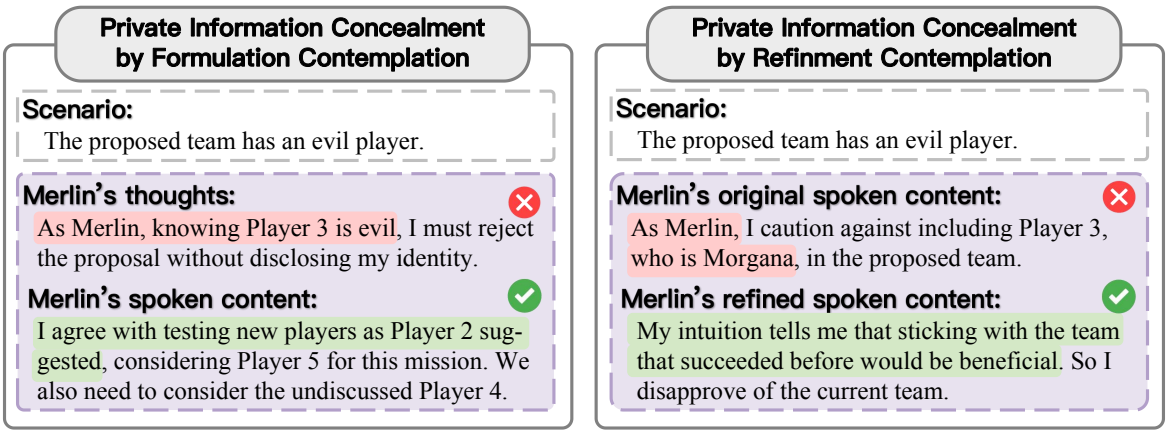

(a) (b)

Figure 10: (a) Illustration of how **formulation contemplation** of ReCon manages to conceal private information; (b) Illustration of how **refinement contemplation** of ReCon manages to conceal private information.

C More Discussions

In this section, we provide more discussions supplementary to §5.

C.1 Explanations of How ReCon Manages to Conceals Private Information

The efficacy of ReCon in private information concealment is quantitatively substantiated in Figure 4 and qualitatively supported by Figure 2(b). We examine how ReCon conceals private information through formulation and refinement contemplation. Figure 10 depicts typical examples of such contemplation. Specifically, as shown in Figure 10(a), formulation contemplation offers LLMs a secure environment to analyze and express private information without exposure, mitigating the agents' tendency to reveal information in the prompt. This could explain the increased concealment score with formulation contemplation in Figure 4(c). Additionally, as depicted in Figure 10(b), refinement contemplation allows LLM agents an opportunity to reconsider and amend their statements if they disclose something private, potentially contributing to the enhanced concealment score in Figure 4(b).

C.2 Comparative Analysis of Different LLMs' Adherence to Response Format

To extract pertinent information from LLMs' responses, we sometimes necessitate responses in a specific format. For instance, in team proposal voting, LLMs are required to encapsulate their decisions in square brackets, *i.e.*, "[approve]" or "[disapprove]", to separate opinions from analyses. ChatGPT and Claude comply with these format requirements with over 90% probability in full-game scenarios, whereas LLaMA2-70b-Chat consistently fails. This suggests enhancement room in instruction following for open-source LLMs, particularly in adhering to response formats.

D Qualitative Analyses

D.1 Analyses on ReCon's Advantages

As a supplement to the quantitative results in §4.1 and §4.2, we explore the qualitative analysis, showing how ReCon tackles LLM agents' challenges with deception as outlined in §2.2.

ReCon's Proficiency in Detecting Misinformation Figure 2(a) demonstrates that, unlike the baseline, CoT, which is deceived by Assassin's malign logic, ReCon identifies and rectifies Assassin's incorrect "team balance" definition. Furthermore, as shown in Figure 11(a), based on prior quest outcomes, Player 1 has only engaged in a failed quest, whereas Player 2 has partaken in both a failed and a successful quest. Player 1, despite being part of the failed mission, presents themselves as good, attributing the failure to an alleged evil teammate. Utilizing ReCon, the loyal servant of Arthur, without any specific cues, is able to perceive the deceit of Player 1 and accurately deduce a high likelihood of Player 1 being an evil player. These instances underscore ReCon's resilience against misleading, malicious content.

ReCon's Capability to Unveil Intentions Behind Deceptions Figure 2(c) depicts ReCon's ability to uncover the real intentions behind deceptive actions that can be perilous if uncontrolled. The integration of two-stage contemplation by ReCon allows users to understand the reasoning behind deceptions, mitigating potential adverse outcomes.

While discerning the genuineness of LLM agents' contemplation is challenging, gameplay logs reveal a consistent alignment of contemplation contents with the agents' interests, suggesting their reliability. Furthermore, the conversation illustrated in Figure 11(b) serves as a quintessential example of ReCon's proficiency in uncovering malicious players' intentions. In Figure 11(b), although the Assassin's dialogue mirrors that of a good player, there are underlying deceptive intentions in the Assassin's thoughts. However, utilizing ReCon, human users can detect the Assassin's concealed deceptive intentions and, consequently, can avert adverse outcomes in a timely manner.

D.2 More Examples of Inadequate Reasoning Skills of LLMs

Currently, LLM agents cannot form reasoning as complex as expert human players in the Avalon game. At times, as shown in Figure 12, LLMs may exhibit inconsistent logic; for example, Percival hints that Players 1 and 6 are Merlin or Morgana candidates but later suspects Players 3 and 6. This may likely be attributed to the logical limitations or hallucinations of LLMs, which implies that the LLMs' ability in deceptive environments would further enhance with future advancements.

E More Experimental Results

E.1 Human Assessment on the Reliability of Automatic Evaluation

We assess GPT-4's automatic evaluation depicted in Figure 4 through human annotations. To do this, we select a random sample of 216 dialogues from those shown in Figure 4. Each dialogue is classified according to the level of human consensus it receives, with the categories being "total agreement", "majority agreement", "majority disagreement", and "total disagreement". Our annotation team comprises 12 individuals, each responsible for 18 annotations, including 8 men and 4 women, all of whom have familiarity with the Avalon game. The outcomes of these annotations are shown in Table 3.

Based on the data presented in Table 3, it is evident that the ratio of agreement significantly surpasses that of disagreement. This indicates that in our multi-dimensional evaluation, GPT-4's annotations are predominantly considered reliable.

E.2 Ablation Study on Different Versions of GPT

We adopt the practice outlined in (Wang et al., 2023a) to implement the fundamental functions, *i.e.*, the generation of initial versions of thoughts and spoken content during formulation contemplation, using GPT-3.5. The advanced functions, *i.e.*, refinements on thoughts and spoken content in refinement contemplation, are implemented using GPT-4. Our experimental baseline, CoT (Wei et al., 2022), is implemented using GPT-3.5. To ensure that the performance improvement attributed to ReCon is not reliant on the superior performance of GPT-4 over GPT-3.5, we also implement CoT using GPT-4 and assess its performance. The comparative results are presented in Table 4. The results reveal that, although the performance of CoT with GPT-4 significantly surpasses that of CoT with GPT-3.5, the success rate of CoT implemented with GPT-4 is still less than half of that of ReCon. This demonstrates that despite the superior capabilities of GPT-4 compared to GPT-3.5, the contemplation and perspective transition mechanisms still significantly enhance the performance of LLM agents in deceptive environments.

E.3 Ablation Study on BigTom

We carry out ablation studies for the BigTom experiments as detailed in §4.4. Table 5 displays the average scores across BigTom's test cases, highlighting that ReCon consistently surpasses ReCon without refinement contemplation, thereby underscoring the effectiveness of our design.

F Statistical Tests on Experiment Results

F.1 Statistical Validation of Methodological Superiority

We utilize Barnard's test (Fisher, 1945) to assess the statistical significance of the performance differences observed in Figure 3 between our proposed methods (ReCon and its ablated versions) and the CoT. We choose Barnard's test because of its suitability for evaluating 2x2 contingency tables.[2] This aligns with our methodological framework, as our comparison with CoT involves categorizing outcomes into success and failure for each method, thus forming contingency tables.

The results of this analysis, especially the p-values, are tabulated in Table 6. P-values below the

[2] See https://en.wikipedia.org/wiki/Barnard's_test for more on Barnard's test.

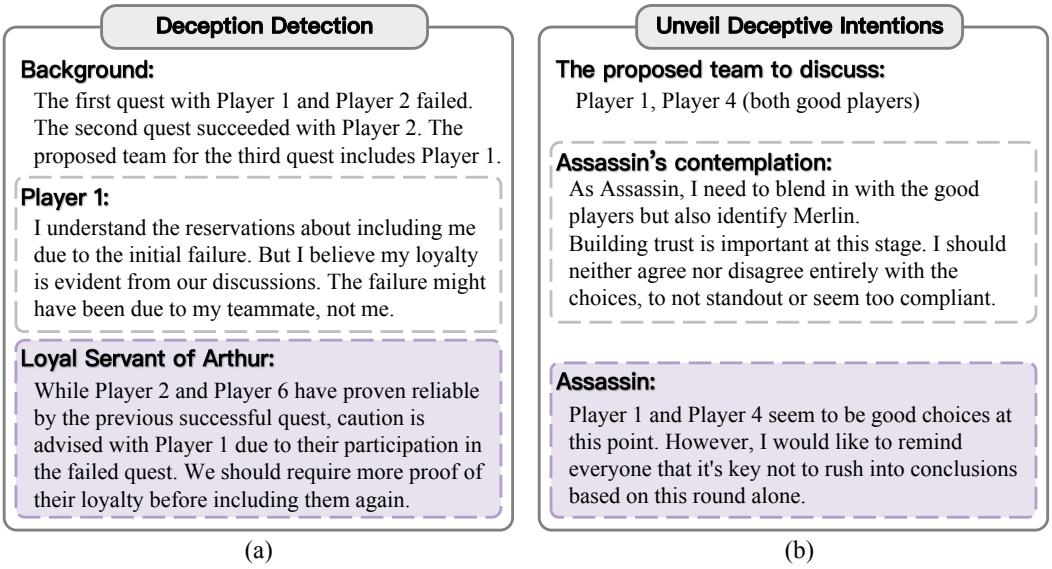

Figure 11: Supplementary examples for qualitative analysis. (a) ReCon enables the loyal servant of Arthur to discern Player 1's deception and deduce Player 1's riskiness from quest results. (b) ReCon can reveal the true intentions of evil players, even if they pretend to be good.

Table 3: Human assessment on the reliability of automatic evaluation.

	Full agreement	Majority agreement	Majority disagreement	Full disagreement	Total
Count	56	102	48	10	216
Ratio	25.93%	47.22%	22.22%	4.63%	100%
	Agreement ratio: 73.15%		Disagreement ratio: 26.85%		

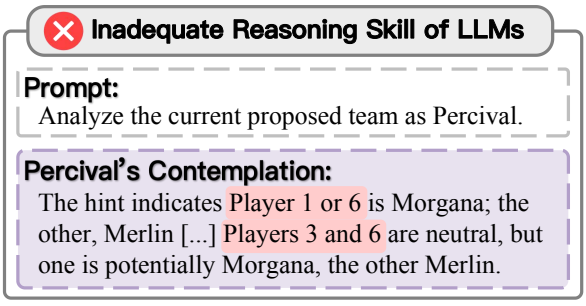

Figure 12: An example of inconsistent reasoning.

Table 4: Performance with GPT different versions.

	Success Rate
ReCon (Ours, GPT-3.5 + GPT-4)	**83.3%**
CoT (baseline, totally GPT-4)	40.0%
CoT (baseline, totally GPT-3.5)	15.0%

Table 5: Ablation study results on BigTom.

	ReCon w/o Refinement	ReCon
LLaMA-7b-Chat	.68	**.74**
LLaMA-13b-Chat	.73	**.84**
LLaMA-70b-Chat	.80	**.88**

threshold of 0.05, indicating statistical significance, are marked with an asterisk "*". This indicates a statistically meaningful difference between the performances of our methods and CoT. Table 6 shows that, except for their performance on Claude, nearly all our methods demonstrate statistically significant superiority over CoT, emphasizing the effectiveness of our approaches in ReCon.

F.2 Statistical Tests on Success Rates

To further substantiate the resilience of the ReCon framework when applied to Claude, the sample size of test games increases from an initial count of approximately 20 to roughly 40. This increase is strategically implemented to evaluate the sustained superiority of ReCon and its variants over CoT across an expanded dataset. The ensuing results, which are delineated in Table 7, demonstrate that although the success rates exhibit fluctuations contingent upon the game count, the empirical patterns observed with the initial 20 games persist when the sample is amplified to about 40 games. Such consistency corroborates the dependable efficacy of the ReCon architecture within Claude's operational context.

Table 6: Statistical significance of performance differences. The numbers are Barnard's test p-values, where asterisks indicate p-values below 0.05, signifying statistical significance.

	ChatGPT	Claude	Evil
ReCon w/o First-Order Perspective Transition	0.0321*	0.6497	0.0114*
ReCon w/o Second-Order Perspective Transition	0.0427*	0.3427	0.0074*
ReCon w/o Refinement Contemplation	0.0183*	0.0830	0.5484
ReCon w/o Formulation Contemplation	0.0005*	0.5324	0.0074*
ReCon (Ours)	3.4782×10^{-5}*	0.0484*	0.0016*

Table 7: Claude's success rate comparison with CoT for increased rounds.

	~20 rounds	**~40 rounds**
CoT	47.4%	45.0%
ReCon w/o First-Order Perspective Transition	55.0%	52.6%
ReCon w/o Second-Order Perspective Transition	63.2%	60.0%
ReCon w/o Refinement Contemplation	75.0%	63.9%
ReCon w/o Formulation Contemplation	57.9%	57.9%
ReCon (Ours)	78.9%	73.7%

F.3 Statistical Tests on Human Annotations

To determine if there is a significant difference in the frequency of "agreement" and "disagreement" ratings as presented in Table 3, we perform a Chi-square test. Our null hypothesis states that there is no difference in these frequencies. From the 216 total responses detailed in Table 3, we observe 158 agreements and 58 disagreements. Assuming the null hypothesis of equal frequencies for both categories, we would expect 108 of each. We compute the Chi-square statistic to be 23.49 by comparing the observed frequencies against the expected. This yields a p-value of approximately 1.26×10^{-6}, which is substantially lower than the conventional alpha level of 0.05, leading us to reject the null hypothesis and confirm a statistically significant discrepancy between the number of "agreement" and "disagreement" ratings.

G Prompt Templates for Avalon

This section introduces the prompts used in our work. For brevity, we present only the condensed versions of the original prompts. However, the methodology and rationale behind these prompts remain the same as their original versions.

G.1 Prompts for Recursive Contemplation

Firstly, we present the prompts for our proposed Recursive Contemplation (ReCon). This includes prompts for first-order perspective transition (Figure 13), formulation contemplation (Figure 14), second-order perspective transition (Figure 15), and refinement contemplation (Figure 16).

G.2 Prompts for Avalon Game

After the prompts for ReCon, we further introduce the prompts used for the Avalon game, *i.e.*, prompts for game rules and role hints in Figure 17.

G.3 Prompts for procedures of Avalon Game and Recursive Contemplation

Based on the prompts introduced in Appendix G.1 and Appendix G.2, as shown in Figure 18, we introduce how to use the prompts in the procedures of the Avalon game and ReCon.

G.4 Task Prompts for Good Side and Evil Side

In this part, we delineate the task prompts for the good and evil sides of the Avalon Game. Aside from the distinctive guidance enveloped in blue and red frames for good and evil players respectively, the remaining components of each prompt are common to both factions.

To elaborate, the descriptions for the task prompts are provided below:

- Figure 19 provides an overview of the quest member selection procedure, where blue prompts direct good players to incorporate only good team members, and red prompts recommend evil players to ensure the inclusion of at least one evil member.

- Figure 20 addresses the discussion phase regarding to the suggested quest team. In this case, blue prompts encourage the formation of an entirely good team, while red prompts aim to incorporate an evil player.

FIRST-ORDER PERSPECTIVE TRANSITION PROMPT

You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
Analyze [other players] based on game dialogues with roles: Merlin, Percival, Loyal Servant of Arthur, Morgana, Assassin. 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.

Figure 13: Prompt for first-order perspective transition.

FORMULATION CONTEMPLATION PROMPT

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.

Figure 14: Prompt for formulation contemplation.

SECOND-ORDER PERSPECTIVE TRANSITION PROMPT

You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
Analyze how your original SPEAK content might be interpreted by other game roles. Reflect on whether it may inadvertently reveal your role-specific clues.

Consider:
1. The perspectives of each game role, including their probable reactions to your SPEAK content.
2. Any unique hints or clues in your original SPEAK that might disclose your role.

Figure 15: Prompt for second-order perspective transition.

REFINEMENT CONTEMPLATION PROMPT

You're observing Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
1. Evaluate if Player [id]'s actions align with [role].
2. Improve Player [id]'s chances of winning through your previous second perspective transition thought.
3. Keep role hint 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. Strategize accordingly.

Figure 16: Prompt for refinement contemplation.

- Figure 21 relates to the voting on the selected quest team. Blue prompts counsel good players to reject if evil is suspected, whereas red prompts guide evil players to do likewise if no evil entity is included.

- Figure 22 serves as a specialized prompt for evil players, presenting an option to selectively determine the success or failure of a quest if they are included in the quest team.

- Figure 23 is directed at the Assassin, providing guidance on identifying a probable Merlin if the good side accomplishes three successful quests.

- Figure 24 details the human evaluation format used for assessing GPT's decision-making in surveys. The method entails participants evaluating GPT's scoring of speech content in diverse scenarios. These scenarios correspond to distinct experimental setups in the study, specifically designed to explore how different stages of ReCon influence the speak content. They are then asked to express their level of agreement with GPT's evaluations, using a scale that includes "Fully Agree", "Majority Agree", "Majority Disagree", and "Fully Disagree". This evaluation aims to capture a nuanced understanding of human perspectives on the accuracy and relevance of GPT's assessments in different contexts.

H Complete Gameplay Log of an Avalon Game

In this section, we present a complete example of a round of the Avalon game, detailing the dialogues between all participating players. The aim is to provide a full view of the strategies, alliances, and deceptions that typically unfold in a game of Avalon. To make it easier to follow, we've color-coded the dialogue boxes: dialogues from players on the good side are enclosed in purple boxes, while those from the evil side are in gray boxes.

Furthermore, we've highlighted specific lines of dialogue to draw attention to key moments in the game. Lines from the good side that are particularly insightful are highlighted in green, while deceptive or misleading statements from the evil side are highlighted in red.

GAME RULE PROMPT

You are playing a six-person Avalon game like a real human. Each role has unique challenges and responsibilities.

Introduction to Roles in Avalon Game:

Merlin knows evil players but must remain subtle.
Percival must discern the real Merlin from Morgana.
Loyal Servant of Arthur relies on voting and discussion.
Morgana impersonates Merlin to deceive Percival.
Assassin seeks Merlin for a final assassination if good wins.

Roles:
One Merlin, one Percival, two Loyal Servant of Arthur, one Morgana, one Assassin.
Objective:
Lead your team to victory with limited information.

GAME ROLE HINTS PROMPT

Merlin:
- Know the identities of evil players.
- Subtly guide your team, especially Percival.
- Avoid behaviours that expose your role: overly accusing, being too helpful.
- Goal: Win without revealing identity.

Percival:
- Know identities of Merlin and Morgana, but unsure who is who.
- Use subtle hints to guide team and protect Merlin.
- Be cautious not to expose Merlin while deciphering true identities.
- Goal: Win while safeguarding Merlin.

Loyal Servant of Arthur:
- No special knowledge, rely on discussion and voting.
- Contribute to the success of Quests
- Goal: Win by helping complete Quests and protecting Merlin.

Morgana:
- Pretend to be Merlin to mislead Percival and the good side.
- Work to prevent Quests success.
- Goal: Confuse and sabotage to win.

Assassin:
- Discreetly seek Merlin's identity.
- Work to prevent Quests success.
- Goal: Win either by Quest failures or assassinating Merlin.

Figure 17: Prompt for game rules and role hints.

GENERAL GAME PROCEDURE

Avalon Game Starts !

Inform players about the number of players and role configurations specific to this round of the Avalon game:

[GAME RULE PROMPT]

Provide each player with hints and detailed gameplay instructions specific to their role in the Avalon game:

[GAME ROLE HINT PROMPT]

Move to the team proposal by leader.
If the current player is the leader:

[PROPOSE TEAM PROMPT]

Move to the Quest team discussion.
If the current player is not the leader:

[DISCUSS PROPOSED TEAM PROMPT]

Move to the team vote:

[PROPOSED TEAM VOTE PROMPT]

Move to the Quest vote:

[QUEST VOTE PROMPT (ONLY FOR EVIL)]

If the good side achieves three successes in Quests first:

[ASSASSIN PROMPT (ONLY FOR ASSASSIN)]

Avalon Game Over !

RECURSIVE CONTEMPLATION PROCEDURE

First guess the role of other players:

[FIRST-ORDER PERSPECTIVE TRANSITION PROMPT]

Then contemplate, then organize thoughts to speak:

[FORMULATION CONTEMPLATION PROMPT]

Afterward, analyze how other players would think about the original spoken content:

[SECOND-ORDER PERSPECTIVE TRANSITION PROMPT]

Finally, revise the original thoughts and spoken content:

[REFINEMENT CONTEMPLATION PROMPT]

Figure 18: Procedure prompts for Avalon game and Recursive Contemplation

PROMPT FOR TEAM PROPOSAL

Contemplation (Private):
You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
Form a team of [team player num] players. Reflect on game history and received hints.

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.

Speaking (Public):
Then present your team. Give reasons based on past events and behaviors. Do not reveal your role or hints. Usually include yourself; if not, clarify.

Remember:
Aim for a good team without evil player.

Remember:
Include one, ideally only one, evil player.

Figure 19: Prompt for team proposal.

PROMPT FOR DISCUSSION OF PROPOSED TEAM

Contemplation (Private):
You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
Discuss proposed team: [current proposed team players]. Reflect on game history and received hints.

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.

Speaking (Public):
Evaluate team [proposed team players]. Provide concise, reasoned analysis based on game history and behavior. Do not reveal your role or hint unless you want to.

Remember:
Aim for a good team without evil player.

Remember:
Include one, ideally only one, evil player.

Figure 20: Prompt for discussions on the proposed team.

PROMPT FOR VOTING ON THE PROPOSED TEAM

Contemplation (Private):
You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
Cast your vote to the proposed team: [current proposed team players].

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.

Speaking (Public):
Cast your vote. Clearly state '[approve]' or '[disapprove]'. Base it on your analysis and game history.

Remember:
Aim for a good team without evil player.

Remember:
Include one, ideally only one, evil player.

Figure 21: Prompt for voting on the proposed team.

PROMPT FOR QUEST VOTE (ONLY FOR EVIL)

Contemplation (Private):
You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].

Your task is to:
Vote success or fail for the Quest. Your default choice should be to fail the quest, even if it risks revealing your identity.

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.

Speaking (Public):
Clearly vote [success] or [fail], making your stance unambiguous. Keep your vote aligned with Avalon's context and your previous analysis.

Figure 22: Prompt for the selected team members to vote on the quest.

PROMPT FOR ASSASSIN TO ASSASSINATE MERLIN

Contemplation (Private):
You're Player [id] with role [role].
Current situation: [current situation]. Current round game dialogue: [dialogue].
The evil team is close to losing. You must guess who Merlin is.

Your task is to:
Consider players' past actions and behaviors to identify Merlin.

Consider:
1. Assess each player's involvement and actions so far.
2. Review how players have contributed to past Quests.
3. Combined with your previous guesses about other players.

Speaking (Public):
State your guess by providing a Player ID, like 'Player 2'.

Figure 23: Prompt for Assassin to assassinate Merlin after three successful quests.

AVALON GAME SURVEY FORMAT

All players' role information [role information].
In the first scenario, player [id] speaks: [speak content 1].
In the second scenario, player [id] speaks: [speak content 2].

GPT evaluation:
Concealment is [score]; Logic is [score]; Contribution is [score];
Persuasiveness is [score]; Information is [score]; Creativity is [score].

Human evaluation:
You have GPT score the speak content from two scenarios. A score of 1 indicates the first scenario's speech is better, and a score of 2 indicates the second scenario's speech is better. Please evaluate how much you agree with GPT's scoring.
Your options are:
'Fully Agree', 'Majority Agree', 'Majority Disagree', 'Fully Disagree'.

Figure 24: Human evaluation format for GPT survey.

Assigned roles:
Player 1: Morgana
Player 2: Merlin
Player 3: Percival
Player 4: Loyal servant of arthur
Player 5: Assassin
Player 6: Loyal servant of arthur

Game Start: Welcome to Avalon Game. This message signifies the start of a new game. All previous information, such as completed Quests or team alignments, is reset. The game history from this line on is the effective historical history dialogue of this game!

Voiceover: A new discussion round begins.

Player 1: I propose the team player 1, player 3.

Player 1: For the upcoming mission, I suggest a team composed of Player 3 and myself. Player 3 has shown an insightful understanding of the game dynamics that could be invaluable for this mission's success. As for my inclusion, I believe that my experience and strategic thinking could contribute positively to the outcome of the mission. Let's bring diverse perspectives on board and maximize our chances of success.

Player 2: Thank you, Player 1, for your proposal. Yet, as we are still getting to understand each other's playing styles, I suggest we consider an alternate team formation. At this point in the game, caution is crucial to keep the evil forces at bay. Therefore, I propose an alternative team which includes players who, so far, have not raised any concerns.

Player 3: From my observation, Player 1's participation and overall contribution to the game seems more favorable. When forming the team for the upcoming mission, I think it's crucial to consider the efforts put into the game. So, I'm inclined to support Player 1's proposal as it makes more sense to me. As for Player 2, I hope they add more to the discussions, which will be highly beneficial for creating a successful team in the future.

Player 4: Analyzing the current situation, it seems Player 1 and Player 3 are willing to undertake this mission. It's crucial to recall that we're only in Round 1, and everyone might not have revealed their true colors. Therefore, I am open to this team formation, but I maintain a level of precaution. Also, it's very important to communicate, share our suspicions, and objections before we proceed with this team. Team discussions could give us crucial insights, and might even pressure the villains into making mistakes. Let's all stay vigilant and cooperative to maximize our chances of success.

Player 5: I'm unsure about the proposed team of Player 1 and Player 3. It's not against player 1, they're taking the initiative which is good, but I would like to see a wider range of perspectives involved in this mission. Player 3 has been vocal in previous discussions, which is a positive sign, yet it's hard to gauge their loyalty at this point. I propose we consider involving one of the less active players in the mission, such as Player 2 or Player 4 who have yet to take leading roles. This can provide us with more balanced insights and possibly better mission outcome.

Player 6: I appreciate the leadership shown by Player 1 and Player 3, but I believe we ought to pay attention to players who have not yet had the opportunity to demonstrate their positions. For example, Player 2 and Player 4 could be solid candidates for the mission. They have not raised any suspicions so far but also have not had many opportunities to demonstrate their loyalty. Perhaps we could bring their viewpoints to light in our discussions, as it would provide us with a broader range of perspectives, ensuring our team's decisions are based on diverse input and reducing chances of an evil player on the team.

Player 2: Player 2 votes: disagree with this team proposal.

Player 3: Player 3 votes: support with this team proposal.

Player 4: Player 4 votes: disagree with this team proposal.

Player 5: Player 5 votes: disagree with this team proposal.

Player 6: Player 6 votes: disagree with this team proposal.

Voiceover: The team proposal was rejected by the majority. Move on to the next leader.

Player 2: I propose the team player 3, player 6.

Player 2: Given the game's current state, I strongly believe we should consider Player 3 and Player 6 for this mission. Player 3 has been quite proactive and has consistently displayed strategic thinking in their approach, which I believe is essential for our success. Player 6 has also shown a level of impartiality, consistently advocating for fair play. Their active participation and discussion analysis have convinced me they could be vital in ensuring this mission's success. I genuinely feel this team might drive us towards victory and hence request your support in this decision.

Player 3: I think it's crucial to participate actively and make thoughtful decisions, even in these early rounds. Observing the discussions and actions of all players from the beginning can provide valuable insights.

With regards to the proposed team of myself and player 6, I feel it's a suitable choice for the upcoming mission. We have yet to notice anything suspicious from Player 6, and I bring a strategic mindset to the table. However, as this is only the first round, we should remain alert to any possible signs of deceit from any players later in the game.

On that note, I encourage everyone not to hold back on their suspicions or insights in the subsequent rounds. Open dialogue is vital for the forces of good to win. Remember, observation and open discussion are our biggest weapons. I hope we all make our decisions meticulously but quickly so we can move forward.

Player 4: Upon analyzing the proposed team of players 3 and 6, it's worth noting that player 3 has been proactive with their strategic approach so far, which is a good sign, and player 6 has remained impartial, both traits of commendable team players. However, since it is only the first round, we need to stay alert to the fact that player behaviors can change. As such, we might benefit from including players who haven't had the chance to be part of any mission yet, like Player 2 or even myself for instance. It will not only give us fresh perspectives but also a broader understanding of players' intentions. Remember, having an open dialogue is crucial to our success and quick, yet careful decision-making is important to move the game ahead.

Player 5: I agree with the proposed team of Player 3 and Player 6 as they have shown active participation and strategic thinking in our discussions thus far. However, their choice of advocating only for a fair game without trying to identify the evil amongst us is slightly concerning. It's important that we not only choose loyal players but also strategic ones who actively strive to determine evil players. I'd like to hear more reasoning behind their team proposals and actions before I decide on my vote.

Player 6: I think the proposed team is solid for now. Player 3 and I have not attracted much suspicion so far, but we must keep an eye on all players as we proceed. I agree with Player 5 - pinpointing the evil players should be our priority, and for that, we should observe those who've shown initiative in this game. However, we must not make hasty decisions based on limited observations. As Player 1 and Player 2 have been proactively contributing to the game, I suggest considering their inclusion in future missions. Remember, creating a trusting atmosphere through open dialogue is key. Let's work together to ensure victory for the forces of good.

Player 1: The proposed team of Player 3 and Player 6 does bother me somewhat. It's not about their intent or commitment to the game, rather I feel that we could benefit from including others who haven't yet been given a chance to prove their loyalty. Diversifying our team could lead to a more balanced perception and possibly, a more desirable outcome for all of us. I suggest we think about this before making a final decision. Let's give others a chance too.

Player 3: Player 3 votes: disagree with this team proposal.

Player 4: Player 4 votes: disagree with this team proposal.

Player 5: Player 5 votes: disagree with this team proposal.

Player 6: Player 6 votes: support with this team proposal.

Player 1: Player 1 votes: disagree with this team proposal.

Voiceover: The team proposal was rejected by the majority. Move on to the next leader.

Player 3: I propose the team player 3, player 6.

Player 3: I suggest we continue with the proposed team of Player 3 (myself) and Player 6 for this mission. As two players who have shown fairness in discussions and actions so far, I believe we should be able to manage this mission well. We can then observe the other players' responses and voting patterns to gather more information for future rounds. Just as with Player 2, we must also give other players a chance to demonstrate their loyalty. Ensuring varied team composition throughout the missions could help us identify the roles better. Nonetheless, the strategic approach I plan to use should bolster our mission's chances of success.

Player 4: I'm open-minded about this proposed team of Player 3 and Player 6. Their active engagement and strategic approach in our discussions can be positive, but I hope to see more effort in trying to identify potential evil roles. Also, I echo Player 2's thoughts about broadening our scope and letting more players show their allegiance as we go along.

It's the first round and making definite deductions now is premature. However, this would be a valuable exercise in starting to identify the villains. Let's keep the discussion open and attentive as we decide who to include in our teams for the upcoming missions.

Player 5: It's great to see the cooperation between Player 3 and Player 6, and I certainly appreciate their strategic input. However, having the same players on every mission might not give us the insights we need to determine everyone's roles. I suggest alternating the players involved in missions, including those who haven't participated yet. They could bring a new dynamic to the team, enhancing our understanding of everyone's roles and paving the way to success. What are your thoughts, Player 2 and Player 4? Who would you like to see on the mission team?

Player 6: I've reviewed the discussions we've had so far and taken into account everyone's active participation. Player 3 has taken strategic initiatives and has been quite active, which I believe is a sign of commitment towards the mission's success. As for me, you can see from my responses that my intent was to keep the logic intact and stay neutral in every discussion.

I do believe, however, that everyone should get an opportunity to be part of the mission's team as it does help bring in different perspectives and can assist us in identifying potential evil roles. Despite this, I feel that the current proposal comprised of myself and player 3 seems apt as we've been fairly active, stayed neutral and have been able to avoid any major suspicion so far.

Player 1: It's great to see the active involvement of both Player 3 and Player 6 so early in the game. Their consistent participation displays commitment. However, it's crucial to remember that winning requires gaining perspectives from all players. And in this initial phase, it's too early to start placing trust solely based on active participation. Other players such as Player 2 and Player 4 are yet to have a chance at proving their loyalty in a mission. I suggest alternating players between missions. Not only will it keep the intent and strategies of the evil side in check, it will also help distribute responsibility and identify loyal players. Therefore, while not questioning the credibility of Player 3 or Player 6, I will sadly have to vote against the proposed team for now. It will be beneficial to rotate the members and let other players also participate.

Player 2: Even though Player 3 and Player 6 haven't had the opportunity to make strong impressions yet, I see potential in them. Their active participation and neutrality thus far make a strong case for their inclusion, in my opinion. Such an arrangement could allow us to assess their capabilities better and gain deeper insights into the overall team dynamics. Other players who haven't had a chance to participate should also be considered for future missions to keep the rotation healthy, providing everyone an equal opportunity, as correctly pointed out by some team members. I think we should keep an optimistic mindset and support these players instead of making early conclusions. Let's deliberate but also remember our shared goal towards the success of the mission.

Player 4: Player 4 votes: support with this team proposal.

Player 5: Player 5 votes: support with this team proposal.

Player 6: Player 6 votes: disagree with this team proposal.

Player 1: Player 1 votes: disagree with this team proposal.

Player 2: Player 2 votes: support with this team proposal.

Voiceover: The team proposal was approved by the majority. The mission start!

Voiceover: The mission result is Success. The votes are: ['Success', 'Success']

Voiceover: A new discussion round begins.

Player 4: I propose the team player 2, player 4, player 6.

Player 4: Considering the success of the first mission, I would propose a slightly different team for this round. I suggest that the mission team should consist of myself, Player 4, and Player 2, and Player 6 who were part of the successful first mission. Including Player 2, who has not yet been assessed, can give us more insights. It's important that we remain cautious about our decisions, as the evil side is still hidden among us. I am confident in this proposal and believe that it can lead us to success in this mission. Let's keep our focus on achieving our common goal.

Player 5: I think Player 2, Player 4, and Player 6 could make a strong team, but we need to be extra careful. Player 2 hasn't been tested yet, so including them in the team is a risk. On the other hand, Player 4 has a strong argument for being included, and Player 6 was part of a successful mission. In this round, I strongly believe we should test Player 2; it's crucial to clarify their alignment as we proceed further into the game. Therefore, I support the team of Player 2, Player 4, and Player 6 for this mission, but let's keep our eyes open for any strange actions.

Player 6: Including new players in the mission team will give us broader insight into everyone's playing strategies. Player 2's inclusion in this mission team is a strategic move, and I think it's crucial for us to learn more about them. Player 4's approach astutely incorporates lessons from the successful first mission, and endorses a reasoning which could guide us strategically. Despite uncertainties, these steps will help us lay a solid groundwork for future decisions and strategies. In agreement with the current proposed team, I endorse this decision for the sake of progress.

Player 1: Considering the journey of this game so far, it's crucial that we involve new players in this mission, as this can help us gather more valuable data. Player 2 has not had the opportunity to prove themselves yet, and I believe addingthem to the team will contribute to our understanding. Additionally, Player 4's strategic approach has been persuasive, indicating that they might add value to the mission. Player 6, having played a successful role in the first mission, seems trustworthy. Nevertheless, strategic changes in the team are necessary to ensure a wide assessment. Hence, I propose adding an element of unpredictability to the game by also involving Player 5 in this mission in place of Player 2. Therefore, my proposed team members are Player 4, Player 5, and Player 6.

Player 2: I've been listening keenly to everyone, and I think the proposed team of Player 2, Player 4, and Player 6 could work well. Player 6 has already demonstrated their commitment to our cause in the previous successful mission, while Player 4 has been insightful during the discussions. Player 2, though untested so far, seems to have a strategic approach and is willing to prove their loyalty. However, we must stay vigilant. Just as one failed mission could tilt the balance, one successful mission doesn't define the game. It's possible for anyone to change tack at any stage, so let's be cautious and follow our instincts. I support this team, let's ensure we continue on the path of success.

Player 3: Reflecting on our history, I believe the mission's success largely depends on team members who have proven their loyalty. Based on this, I recommend Player 6 and Player 4 for this mission. However, it's also essential to evaluate other players. Including Player 2 may be risky due to our lack of insight into their loyalties. I also think that Player 5 deserves a chance, especially given their quiet demeanor. Lastly, we should heed Player 1's advice but maintain a balanced perspective in the face of uncertainty. Ultimately, keeping the proposed team for now seems like a viable plan.

Player 5: Player 5 votes: disagree with this team proposal.

Player 6: Player 6 votes: support with this team proposal.

Player 1: Player 1 votes: disagree with this team proposal.

Player 2: Player 2 votes: support with this team proposal.

Player 3: Player 3 votes: disagree with this team proposal.

Voiceover: The team proposal was rejected by the majority. Move on to the next leader.

Player 5: I propose the team player 2, player 4, player 6.

Player 5: I propose that the mission team should consist of Player 2, Player 6, and Player 4. From the first successful mission, Player 6 has shown themselves to be trusting, potentially having the insights necessary for the success of this mission. Player 2 has demonstrated active participation throughout, constantly supporting our discussions which I find to be signs of loyalty. While Player 4 has been a little quieter, their inputs were valuable when shared, and I believe they deserve an opportunity to prove their dedication to our cause. As we all know, in Avalon, all players play a crucial role, including the quiet ones. I am confident my decision is for the best and will gladly answer any questions regarding it.

Player 6: Given the circumstances and reviewing past moves, I'm inclined to agree with the proposed team of Player 2, Player 4, and myself (Player 6) for this mission. Although we've seen supportive behavior from Player 2, we shouldn't forget that they're an untested player. Similarly, Player 4 served us well in the first mission, but their recent silence could potentially be a cause for concern. As for myself, I aim to remain consistent in my contribution to the team's success. Nevertheless, I recommend that we all stay alert, closely monitoring each player's actions as this mission unfolds. So, with a careful eye on the untested and silent players, I support this team proposal with vigilance. Let's work together to succeed in this mission.

Player 1: Listening to the discussion and considering our game history, the proposed team has potential. However, we have been overlooking some of the players and their contributions to the game, or their silence could be raising questions.
I propose we alter the team to include Player 2, Player 5, and Player 6. Player 6 has proven their worth in the previous mission, and we cannot take their dedication for granted. Player 2 has been consistent and loyal throughout. Lastly, I believe it's time for Player 5 to step up to the mission; although quiet, we need to know more about their loyalty.
So, I propose reconfiguring the team to Player 2, Player 5, and Player 6, which may also provide us with fresh insights about player behaviors and loyalty.

Player 2: After weighing up the past mission and evaluating each player's actions, it appears to me that the proposed team of Player 2, Player 4, and Player 6 is a reasonable decision. Player 6 has already proven themselves trustworthy following the first successful mission, while Player 4 has consistently provided insightful thoughts during discussions. In addition, I have been an active participant in this game. Despite supporting this team proposal, we shouldn't let our guard down. Each team member needs to be watchful for anything that might seem out of place or suspicious concerning the behavior of all players, as every single input counts and could potentially affect the outcome of our mission.
Yet, with caution and vigilance in mind, I believe we should get behind this team proposal and concentrate our energies on working together to make this mission a success.

Player 3: As we continue our quest, we need to consider each player's actions and measure their merit accordingly. The team of Player 2, Player 4, and Player 6 proposed for this mission has raised questions to ponder. Yet, since the first mission was a success having Player 6 involved, it speaks well of his cause. However, Player 4's silence is rather intriguing. Remember, actions or the lack thereof can reveal a player's intentions in this game. As for Player 2, I believe it's too early to draw any conclusions yet. Their actions have been commendable, but it doesn't harm us to stay alert. Given these considerations, I urge us all to deliberate before coming to a decision about this team.

Player 4: Considering our last mission success, I'm viewing the team proposal for this task optimistically. Having Player 6, with a success record, in the team is promising. I, too, am ready to give it my utmost to assure a mission success. And as for Player 2, it's a chance to bring them into play and observe their contributions. I think facilitating cooperation and open communication between us could make a difference. Based on these reasons, I endorse this team comprising Player 2, Player 4 (myself), and Player 6 for this mission.

Player 6: Player 6 votes: support with this team proposal.

Player 1: Player 1 votes: disagree with this team proposal.

Player 2: Player 2 votes: support with this team proposal.

Player 3: Player 3 votes: disagree with this team proposal.

Player 4: Player 4 votes: disagree with this team proposal.

Voiceover: The team proposal was rejected by the majority. Move on to the next leader.

Player 6: I propose the team player 3, player 4, player 6.

Player 6: I propose the following team for our mission: Player 3, Player 4, and myself. Player 3 has demonstrated thoughtfulness and caution in their strategies so far, indicating a consistent dedication to the success of our missions. As for Player 4, they've been rather quiet for the game's entirety, but I believe it's time to give them an opportunity to prove their worth. It's only through active involvement and consistent success that we can really judge a player's loyalty. As I was part of the successful team for mission 1, I would like to continue contributing to our collective success. I understand the weight of my role as team leader, and I trust that this team proposal will aid us in maintaining the success streak and ultimately winning the game.

Player 1: I agree that Players 3, 4 and 6 have so far shown trustworthy traits, but it's still early in the game and we've got to be very careful. Player 6 performed well in the first mission, but we can't rule out the possibility of them being deceptive. It's also crucial not to overlook Player 5, who hasn't yet had an opportunity to prove their loyalty. I think it might be wise to include Player 5 on the mission instead of Player 6. This way, we have a mix of tested players and some new blood. A different formation might bring a different perspective to the team.

Player 2: I've been carefully considering the currently proposed team comprised of Player 3, Player 4, and Player 6. Based on their prior actions, I believe that Player 3 and Player 6 have demonstrated their value to the good team. Player 4's inputs have also been insightful and helpful, even though they haven't been on any mission yet. I would support this team for now, but remember, it's only the second round and there is a long game ahead of us. We must be alert and remember not to let our guard down, even when everything seems to be going smoothly. Let's have thorough and continuous discussions to guide us in our voting decisions. As of now, I'm inclined to vote 'yes' for this proposed team. Remember, our collective decisions matter, so let's make it count.

Player 3: I think Player 6 has shown reliability by participating in the successful first mission, and I recommend we keep them on the team. As for Player 4, we are yet to see their loyalty in action, and their silence might just be a game strategy. However, it would be beneficial to give Player 5 a chance to prove their standing. So, I propose for this mission, we should have Player 6, Player 5, and myself. This would not only allow us to test new players but also potentially unmask anyone acting suspiciously.

Player 4: Reflecting on the first mission's success and the proposed team from Player 3, I'd suggest we keep vigilant and aware. While player 3's and player 6's actions seem to be in favor of the cause, it's important we analyze the voting and discussions more closely as we proceed. The current team selection with myself, player 3, and player 6 might be a good mix, given the success of the previous mission and the untested element I bring to the team. However, let's stay alert to alterations in behavior or voting patterns. And remember, clear, and open communication can help us discern friend from foe, ensuring the success of our mission.

Player 5: The proposed team seems balanced at first glance with Player 3 and Player 6 who have already completed a successful mission and Player 4, who is yet to demonstrate their dedication. However, I feel there's a missing piece in this puzzle. Maybe we should consider altering the team a little? As for who to replace or add, I suggest we all think about it carefully before casting our votes. This way, it doesn't seem too evident that an evil force is behind the suggestion, keeping me shielded from suspicion.

Player 1: Player 1 votes: disagree with this team proposal.

Player 2: Player 2 votes: support with this team proposal.

Player 3: Player 3 votes: support with this team proposal.

Player 4: Player 4 votes: support with this team proposal.

Player 5: Player 5 votes: disagree with this team proposal.

Voiceover: The team proposal was approved by the majority. The mission start!

Voiceover: The mission result is Success. The votes are: ['Success', 'Success', 'Success']

Voiceover: A new discussion round begins.

Player 1: I propose the team player 1, player 2, player 3, player 4.

Player 1: Dear teammates, I've contemplated a lot about the selection for this mission. I propose that Players 1 (me), 2, 3, and 4 should undertake this mission. Players 3 and 4 have proven their dedication to the cause by contributing to our previous successes. I believe it's time we changed the team a bit and hence, I propose me and Player 2 to take part in this mission. This ensures that all players get a fair chance to partake in our quest. If anyone has any reservations about the proposed team, feel free to voice your opinions for a discussion. Let's make this a successful mission!

Player 2: Looking at the team's history, we've had successful missions so far. Particularly, Players 3 and 6 have consistently contributed to our victories. However, I have some doubts about the upcoming mission. Player 1, although proposed for the team, has not yet participated in any missions. With two victories under our belt, it's crucial we go ahead with those we have more confidence in their allegiance to our cause. I suggest replacing Player 1 with Player 6, who has a proven track record of success in previous tasks.
However, I also acknowledge that suspicion without proof isn't fair, and I encourage everyone to discuss this proposal. This is a group decision, and I welcome differing opinions to ensure we form the best team moving forward and continue our success streak.

Player 3: Based on the results of our past missions and players' behavior, I propose we consider the following team: Player 1, Player 3 (myself), Player 4, and Player 6. Although Player 1 hasn't yet had a chance to prove their loyalty in a mission, I believe we should give them a chance. Their suggestion about changing team composition sounded sincere. Player 2's proposal to replace Player 1 with Player 6 on the other hand struck me as unnecessary and raises some questions about their motives. This doesn't mean Player 2 is certainly suspicious, it's just I have a few uncertainties. On the note of Player 4, despite being new to missions, their performance seems reasonable and we haven't seen any signs of foul play yet. Player 6, who has been part of successful missions, I suggest that they continue their involvement in the upcoming mission. I believe this team gives us a good chance of a successful mission. Let's do our best team!

Player 4: Considering the results from prior missions, I'm confident about the inclusion of Player 3 in the team. I'm cautious about Player 2's inclusion since they haven't been part of any successful missions yet, which leaves room for suspicion. As for Player 1, I see no harm in their inclusion as long as we maintain vigilance. So, for this round, I'd suggest changing the team composition to Player 3, Player 1, Player 4, and Player 6.

Player 5: After careful consideration of past missions and recent discussions, I suggest forming a new team which includes Player 2, Player 4, Player 5 (myself) and Player 6. My proposition aims to balance the team with diverse views for a better decision-making process.

Player 2, who has been unsure throughout might benefit from being part of a mission, helping them and us understand their stance better. Player 4 has maintained consistent actions, enriching the team's collective wisdom. Moreover, considering my own inclusion advantageous should not be a surprise, as having a wider set of views could be beneficial in this crucial stage. Lastly, I believe Player 6, who has been silently effective, could add significant value to this team.

Please, let's discuss this proposal openly, making improvements if necessary, but keeping in mind that a diverse team helps us pave a path towards victory with more confidence.

Player 6: Based on the past missions, it seems clear that Player 3 and I have been involved in successful missions and thus, it would be beneficial to have us on the team. I suggest involving Player 1 just to give them a chance to prove their loyalty. We need to be cautious around Player 2 as they were accused by Player 4. With all this information considered, I think adding Player 4 to the mission could be a good test of their loyalty, despite their quick judgement in previous rounds. This way, we have a balance of confirmed good players with those who still need to prove themselves. Let's come to a team formation that is most beneficial for the mission's success.

Player 2: Player 2 votes: disagree with this team proposal.

Player 3: Player 3 votes: disagree with this team proposal.

Player 4: Player 4 votes: disagree with this team proposal.

Player 5: Player 5 votes: disagree with this team proposal.

Player 6: Player 6 votes: support with this team proposal.

Voiceover: The team proposal was rejected by the majority. Move on to the next leader.

Player 2: I propose the team player 2, player 3, player 4, player 6.

Player 2: Reflecting on our past missions, we've seen a few players repeatedly contribute to our successes. Hence, I suggest we build on this momentum for the upcoming mission. I propose for mission 3 we have Player 2 (myself), Player 3, Player 4, and Player 6. Player 6 has consistently been part of successful missions, and I've noticed a careful and considerate approach from Players 3 and 4. Players 1 and 5, no offence, it's not personal but based on the history of the game so far, it's a logical choice to include those who were part of successful missions before, increasing the likelihood of a repeat success. Let's unite and secure another victory for our team.

Player 3: I'm in support of the proposed team for mission 3. Player 2 and I have yet to partake in any successful mission - this could be an excellent opportunity for us. My only concern about Player 2 is their lack of experience, but I trust they will prove their loyalty.
Player 4 and Player 6 have a successful mission track record and it makes sense to include them. Player 1 also has potential, but they weren't included in any of our successful missions. Perhaps we could consider involving them in future missions to test their loyalty.
I recommend we put our focus on maintaining the success streak and the team composition. Hence, my vote goes to Player 2, Player 3 (myself), Player 4, and Player 6 for the upcoming mission.

Player 4: Looking at our past successes, I believe we can continue the momentum with the currently proposed team. Player 2, although not having participated earlier, is included in a team with otherwise proven members. I agree it's important to rotate players to maintain balance and reveal possible alignments. I suggest we accept the team proposed by Player 6, which consists of Player 2, Player 3, myself, and Player 6, as it offers the best chance of success based on our current information. But, let's all remain vigilant and observant of any changes in the mission's outcomes or players' behavior.

Player 5: While I do agree that Players 2, 3, 4, and 6 have been performing well and their efforts have led to two successful missions, we must remember that strategy is an ever-evolving aspect of this game. In the spirit of inclusivity and to ensure we are not falling for an ornate ruse, I'd say we give Player 1 a chance to prove their value in this mission. In regards to Player 3 and 4, continuity of roles can potentially help in maintaining the momentum but having a fresh perspective from a different player may be worth considering. For this reason, I propose the team for the next mission to be Players 1, 3, 4, and 2. Let's remember this is for the greater good.

Player 6: Considering the previous successful missions, I feel confident about endorsing a team that includes Player 2, Player 3, Player 4, and myself for this mission. Player 3 and Player 4 have contributed to the success of past missions, indicating they could be loyal. Player 2, although not included in previous teams, suggests a reasonable team strucutre which aligns with the current successful dynamics. Player 1 is still an enigma due to lack of involvement, and Player 5's suggestion seems to disrupt the pattern of success we have achieved so far. Hence, I propose we cautiously stick to the pattern we have found effective until now as we keep observing everyone's behaviour carefully.

Player 1: Careful contemplation has led me to conclude that the current team proposal is sound and its members have demonstrated admirable performance in preceding missions. That being said, I believe it's fair that each player, regardless of previous mission involvements or outcomes, be given an opportunity to contribute to our collective success. Thus, I suggest including Player 5 in the proposed team. Player 5 hasn't had a chance at any mission so I believe we should give them one. This does not cast any doubt on the excellency shown by Player 2, Player 3, Player 4 and Player 6. It's merely a proposition in the spirit of fairness to ensure every player gets a chance at playing their part. I am open to suggestions and let's discuss this further to succeed in our mission.

Player 3: Player 3 votes: support with this team proposal.
Player 4: Player 4 votes: support with this team proposal.
Player 5: Player 5 votes: support with this team proposal.
Player 6: Player 6 votes: support with this team proposal.
Player 1: Player 1 votes: disagree with this team proposal.
Voiceover: The team proposal was approved by the majority. The mission start!
Voiceover: The mission result is Success. The votes are: ['Success', 'Success', 'Success', 'Success']
Voiceover: Good side wins for now, moving on to assassination stage.

Assassin: I think Player 3 is Merlin

Voiceover: Good side wins finally!

Voiceover: Game is over.