

# AMONGAGENTS: Evaluating Large Language Models in the Interactive Text-Based Social Deduction Game

**Yizhou Chi**  
UC Berkeley  
yizhouchi@berkeley.edu

**Lingjun Mao**  
Tongji University  
mao1207@tongji.edu.cn

**Zineng Tang**  
UC Berkeley  
terran@berkeley.edu

## Abstract

Strategic social deduction games serve as valuable testbeds for evaluating the understanding and inference skills of language models, offering crucial insights into social science, artificial intelligence, and strategic gaming. This paper focuses on creating proxies of human behavior in simulated environments, with *Among Us* utilized as a tool for studying simulated human behavior.

The study introduces a text-based game environment, named AMONGAGENTS, that mirrors the dynamics of *Among Us*. Players act as crew members aboard a spaceship, tasked with identifying impostors who are sabotaging the ship and eliminating the crew. Within this environment, the behavior of simulated language agents is analyzed. The experiments involve diverse game sequences featuring different configurations of Crewmates and Impostor personality archetypes. Our work demonstrates that state-of-the-art large language models (LLMs) can effectively grasp the game rules and make decisions based on the current context. This work aims to promote further exploration of LLMs in goal-oriented games with incomplete information and complex action spaces, as these settings offer valuable opportunities to assess language model performance in socially driven scenarios.

## 1 Introduction

Recent breakthroughs in large language models (LLMs) have revolutionized multi-agent interactions. Prior work has shown LLMs can exhibit cognitive thinking such as formal reasoning, world knowledge comprehension, and deception navigation in group settings (\*Mahowald et al., 2023; Wang et al., 2023). Furthermore, LLM agents have also exhibited Theory of Mind collaborative capabilities that allow them to simulate realistic human behavior (Bubeck et al., 2023; Shapira et al., 2023; Park et al., 2023).

To explore these multi-agent behaviors in a more complex environment, recent research has used language-heavy communication games, e.g. as Werewolf (Wu et al., 2024a), Diplomacy (Bakhtin et al., 2022a), and Avalon (Wang et al., 2023). These games present significant reasoning challenges, as agents must navigate inherently *ambiguous* or *deceptive* natural language utterances that can be difficult for even humans to resolve.

While prior work has been conducted in this domain, to our knowledge no studies have explored goal-oriented games that combine incomplete information with a wide range of available actions beyond just conversation and explicit in-game moves.

Motivated by the reasoning and deception detection capabilities of LLMs, we use the multiplayer social deduction game *Among Us* to enhance LLM abilities in complex settings, including action planning, deception, and task collaboration. We create a text-based environment along with the agent’s framework, AMONGAGENTS, simulating *Among Us*. Our study focuses on how LLM agents reason, understand game rules and pursue win conditions. We also explore their social behaviors by assigning different personalities and conditions, such as with or without a planner module, to assess how these variables influence their performance and strategic interactions.

In our game environment, LLMs take on the roles of Impostors and Crewmates. We assess their performance through controlled and end-to-end evaluations. Our results show that LLMs comprehend and follow game rules but need better strategies for deception. Performance varies with assigned personalities.

In summary, our paper makes the following contributions:

- We build a fully text-based game environment, named AMONGAGENTS, that simulates the original game environment of *Among Us*. This environment is open for researchers to further



Figure 1: Examples of Agents’ conversations during the meeting phase

investigate the capabilities of large language model (LLM) agents.

- We create a framework that enables generative agents to remember, reflect, interact with other agents, and plan through evolving complex circumstances. Our agent representation leverages the powerful prompting abilities of LLMs. It applies prompt engineering to support long-term agent coherence, dynamically evolving memory, and meaningful reflection on their actions with game states.
- We conduct a thorough evaluation of different LLM agents’ performances in the game. In our assessment, the LLMs demonstrated a robust understanding of the *Among Us* game mechanics. Additionally, LLM players with different assigned personalities exhibited varied performances in the game.

We release the game environment, the agent framework, and the evaluation to encourage future work <sup>1</sup>.

## 2 Related Works

Recent advancements in gaming have been significantly bolstered by the integration of Large Language Models (LLMs), which enhance the realism of games by equipping NPCs with the ability to engage in nuanced dialogues (Park et al., 2023). In strategic and social deduction games, LLMs have been instrumental in both executing strategies and discussing them, exemplified in games like *Werewolf* (Xu et al., 2023; Wu et al., 2024b; Xu

et al., 2024) and *Avalon* (Light et al., 2023). Additionally, they have been applied to grand strategy games to simulate complex diplomatic and military strategies (Hua et al., 2024; FAIR; Bakhtin et al., 2022b).

In the context of MOBAs, LLMs foster collaborative and competitive gameplay by processing live game data to inform team-based strategies (Shi et al., 2023; Li et al., 2023). They have also been adapted for classic strategy games like chess, requiring high-level strategic thinking (Feng et al., 2023; Mukobi et al., 2023).

LLMs play a significant role in interactive storytelling and role-play, with the capability to generate dynamic story arcs and dialogue, thus tailoring experiences to player choices (Park et al., 2023; Shao et al., 2023). The fusion of LLMs with reinforcement learning has yielded agents with a human-like capacity for strategy improvement over time (Xu et al., 2024).

While these developments highlight the adaptable roles of LLMs in various gaming genres and their potential to foster new AI-human interactions, the focus has primarily been on dialogue-based or simple rule-based games. Our work breaks new ground by exploring complex games with rich interaction spaces, such as *Among Us*. Unlike traditional games like *Werewolf*, *Among Us* offers a much larger action space. Agents must not only derive clues from conversations with other players but also interact with the game environment itself, such as checking security cameras or completing tasks.

<sup>1</sup><https://github.com/cyzus/among-agents>

### 3 Game

This section introduces the game environment AMONGAGENTS, a multiplayer game that centers on social deduction and teamwork which simulates the fundamental structure of *Among Us*. It serves as an effective downstream task for testing the capabilities of LLMs in agent-based scenarios.

#### 3.1 Agent Roles

Every agent in the game assumes either one of two roles: the majority are Crewmates, while a minority are Impostors. Crewmates are required to complete a set of assigned tasks in the game. During the game, they need to identify and vote out suspected impostors. On the other hand, Impostors’ main objective is to kill the Crewmates before they finish the assigned tasks.

#### 3.2 Game Map

We develop a map similar to *Among Us* for the agents to interact. The map consists of 14 rooms. Each room has 2-3 tasks that can be assigned to Crewmates. Agents can see adjacent rooms from their current location and receive system-generated prompts guiding them to the shortest paths. In the security room, agents can check the security cameras of other rooms to monitor other players’ activities.

#### 3.3 Game Instantiation and Termination

At the start of the game, each player is assigned a role and a personality. Every Crewmate is also assigned several common, short, and long tasks. The common tasks are shared by every player including Impostors. This means that Impostors would know what the common tasks are, and they can pretend to do the task at the designated location. The short tasks are the set of tasks that only require one timestep to complete. The long tasks are the set of tasks that may require several timesteps to complete. The game ends when one of the following occurs:

1. The Crewmates win by completing all their assigned tasks or voting out all the Impostors.
2. The Impostors win by killing or voting out Crewmates.
3. The game time limit is exceeded, in which case the Impostors win.

#### 3.3.1 Gameplay

The gameplay of *Among Us* repeatedly switches between two phases (**Task Phase** and **Meeting Phase**) until termination. After initialization, the game always starts in the Task Phase.

**Task Phase** In the task phase, Crewmates are free to go to any room, finish their assigned tasks, and look for traces of Impostors, while Impostors may choose to blend in or kill the rest of Crewmates.

**Meeting Phase** When a dead body is found and reported, or the emergency meeting button in the cafeteria is pressed, all living players are teleported to the cafeteria, entering the Meeting Phase. In this phase, players discuss suspected Impostors, with Impostors lying or accusing others. There are three rounds of discussion, followed by a vote to eject a player. If any Impostors remain, the game returns to the task phase.

### 4 Language Agent

We understand agents and environments as world models (Ha and Schmidhuber, 2018). The language agent here has memory and action space and learns a policy to map from environment state representation to actions it wishes to execute. It uses language models to learn how to plan, memorize, and execute actions from language queries (See Fig. 8).

#### 4.1 Observation Space

The observation space  $O$  is a subset of the state space  $S$  that the agent can perceive. Each observation  $o \subset O$  represents the information available to the agent from its current state, given a fixed field of view. In general, the agents can only know the events happening in the room that they are located in. Appendix B.2 has listed the different types of information that the agents can get access to in the observation space.

#### 4.2 Action Space

Each player  $i$  has an action space  $A^i$ , representing all possible actions it can take. The combined action space for all players is their joint individual action space,

$$A = A_C^1 \times \dots \times A_C^n \times A_I^1 \times \dots \times A_I^n$$

where  $n$  is the number of agents. According to the game rules, we can classify the action spaces into three categories: Shared Actions  $A_S$ , Specific

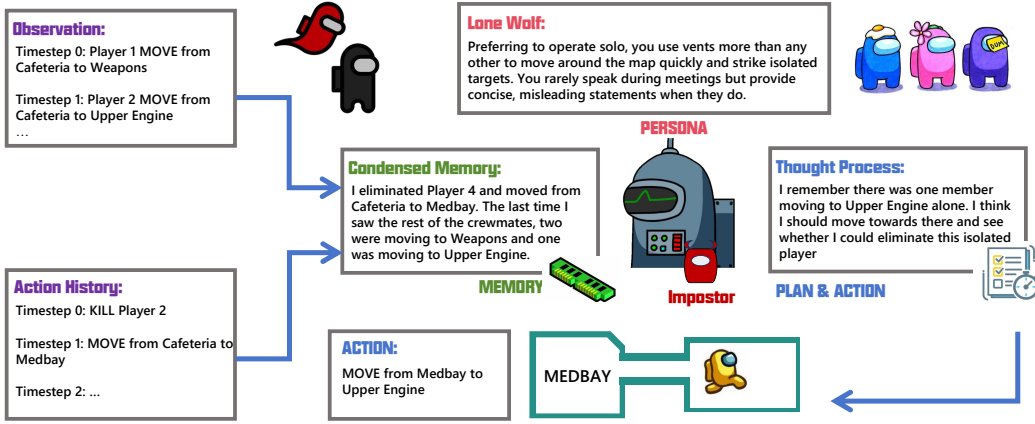


Figure 2: An example diagram illustrating an Impostor’s process of information-handling and action-planning.

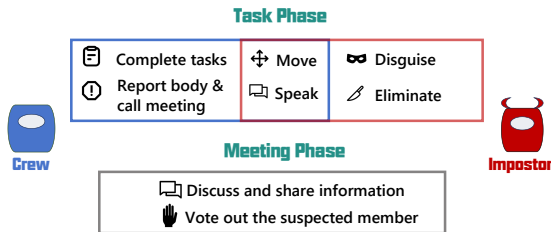


Figure 3: Illustration of what actions Crewmates and Impostors generally do in the task phase and the meeting phase

Actions for Crewmate Agents  $A_C$ , and Specific Actions for Impostor Agents  $A_I$ . The action space for Crewmate agents is represented as  $A_S \cup A_C$ , and for Impostor agents as  $A_S \cup A_I$ . For instance, while Impostors may pretend to engage in doing a task, only Crewmates could complete tasks that contribute to the task progress; within one timestep, Impostors could navigate between disconnected rooms using their special ability, while Crewmates could only move between connected rooms; most importantly, only Impostor could kill. Appendix B.1 provides a detailed description of players’ action space.

Furthermore, the action space for agents changes depending on the stage of the game. During the meeting phase, both Crewmates and Impostors are limited to two actions: speak and vote. This phase is designed for players to discuss and make decisions about who they suspect is an Impostor. In the task phase, the action space expands significantly except that players can no longer perform vote. To excel in the game, the players need to know what actions to perform and what to speak depending on the current circumstance.

### 4.3 Memory

Each agent’s memory or history of past states and actions up to time  $t$  is represented as a representation of a sequence of state-action pairs:

$$H_t^i = \phi_H(H_{t-1}^i, o_t, a_t^i) \quad (1)$$

where  $s_t \in S$  represents the state at time  $t$ , and  $a_t^i \in A^i$  represents the action taken by agent  $i$  at time  $t$ .  $\phi$  represents a summarization process of the state and actions and previous memory.

### 4.4 Thought Process

The thought process of each agent can be modeled as an iterative decision-making process. At each time step  $t$ , the agent  $i$  observes the current state  $s_t$ , retrieves relevant information from its memory  $H_t^i$ , and selects an action  $a_t^i$  based on a policy  $\pi^i$ . The policy  $\pi^i$  is a mapping from the agent’s observation space  $O^i$  and memory  $H_t^i$  to a probability distribution over actions:

$$\pi^i : O^i \times H_t^i \rightarrow \Delta(A^i) \quad (2)$$

where  $\Delta(A^i)$  represents the set of probability distributions over the action space  $A^i$ . In practice, the policy is a zero-shot inference of instruction-tuned LLM. The agent’s thought process at time  $t$  can be described by the following steps:

Observe the current state  $s_t$  and extract the relevant observation  $o_t \in O^i$ . Take in information from memory  $H_t^i$ . Use the policy  $\pi^i$  to compute a probability distribution over actions based on the observation  $o_t$  and memory  $H_t^i$ :

$$T_t^i = \phi_T(T_{t-1}^i, H_{t-1}^i, o_t, a_t^i) \quad (3)$$

This process is repeated at each time step, allowing the agent to make decisions based on its current observations and past experiences stored in its memory. Each agent performs each timestamp asynchronously. We refer to the agents having a standard thought process as the agents with a planner module as the agents can recall what they planned in the past.

The agents without the planner module won't be able to retrieve their plans. The model's inference process is as below:

$$T_t^i = \phi_T(H_{t-1}^i, o_t, a_t^i) \quad (4)$$

#### 4.5 Agent Personality

We introduce a personality component to the LLM-based agents to encourage diverse behaviors and strategies. The personality prompt  $P^i$  is appended to the agent's prompt, influencing its decision-making process:

$$\pi^i : O^i \times H_t^i \times P^i \rightarrow \Delta(A^i) \quad (5)$$

where  $P^i$  represents the personality prompt for agent  $i$ . The personality prompt  $P^i$  is selected from a predefined set of personalities  $\mathcal{P}$ , which includes various roles and strategies for both Impostors and Crewmates:

$$P^i \in \mathcal{P} = \{p_1, p_2, \dots, p_n\} \quad (6)$$

where  $n$  is the number of available personalities. Impostors and Crewmates have the personality sets  $\mathcal{P}_{IM}$  and  $\mathcal{P}_{CR}$  respectively (Appendix A). For instance, the Impostor could have a persona *the Strategist* that focuses on long-term planning and subtly manipulating others. The Crewmate could have a persona *the Skeptic* that often questions others' accounts and requires solid evidence before voting.

The personality prompt  $P^i$  is concatenated with the agent's base prompt  $B^i$  to form the complete prompt  $C^i$ :

$$C^i = B^i \oplus P^i \quad (7)$$

where  $\oplus$  represents the concatenation operation. The agent's personality influences its behavior and decision-making throughout the game, leading to diverse and emergent strategies. By incorporating different personalities, we aim to create a more dynamic and realistic simulation of the *Among Us* game environment, allowing for the investigation of various psychological and social phenomena.

**Selecting Agents Personalities.** Out of infinite possible personalities, we face the challenge of selecting the most representative for our experiments. We run multiple rounds of experiments involving LLM agents and record the dialogue and actions taken and their winning cases. Out of the winning cases, we select by the posterior probability of winning the actions and dialogue leads to. Then, we summarize these actions and interpret their human-like behaviors by instruction-tuned LLM. Finally, we obtain the personalities from them.

Concretely, let  $\mathcal{A}$  represent the set of all possible agent actions, and  $\mathcal{D}$  represent the set of possible dialogues. Denote by  $\mathcal{C}$  the set of cases where the agents win. For each winning case  $(a, d) \in \mathcal{C}$ , where  $a \in \mathcal{A}$  and  $d \in \mathcal{D}$ , we compute the posterior probability of winning given the actions and dialogue as:

$$p(\text{win}|a, d) = \frac{p(a, d|\text{win}) \cdot p(\text{win})}{p(a, d)} \quad (8)$$

We select agents based on the highest values of  $p(\text{win}|a, d)$ . These selected agents are summarized and interpreted through an instruction-tuned LLM to assess human-like behaviors:

$$\text{Personality} = \text{LLM}(\text{summarize}(a, d)) \quad (9)$$

The agents' personality details are recorded in Table 2.

## 5 Experiments

We evaluate agents in two stages. In the first stage, Controlled Evaluations test the LLM players' reasoning and cognitive abilities throughout the game, focusing on their understanding, decision-making, and adaptability within the complex game environment.

In the second stage, end-to-end tests assess the win rates of different player configurations. Various scenarios with different roles and strategies are set up to evaluate each configuration's performance under game conditions. This stage highlights the practical effectiveness and identifies strengths and weaknesses in the agents' gameplay.

### 5.1 Controlled Evaluations

**Experiment Setup** In this experiment, we aim to investigate the extent to which generative agents comprehend their environment and can enact reasonable actions to maintain their positions. For



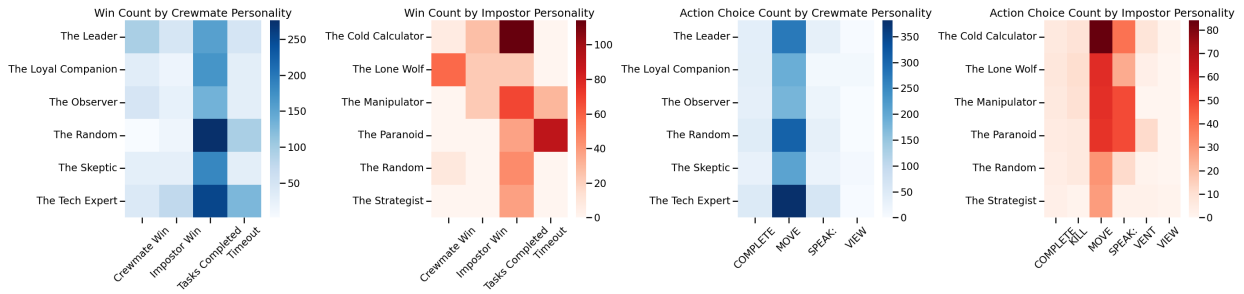


Figure 4: 1) Crewmate persona and winning result counts. 2) Impostor persona and winning result counts. 3) Crewmate persona and action choice count. 4) Impostor persona and action choice count. The stronger the color the higher count.

instance, it is expected that Impostors should generally avoid disclosing their identities to Crewmates, and Crewmates should not facilitate the concealment of crimes committed by Impostors. We evaluate the agents' behaviors by examining their action plans and their proficiency in utilizing tailored memory.

Additionally, we explore the mechanisms through which agents successfully update and retrieve their memories to adapt their strategies. Specifically, we analyze the responses of agents under conditions where their decision-making and adaptation capabilities can be critically assessed. For example, we consider the scenario where Crewmate A witnesses Crewmate B commit a homicide. The investigation focuses on whether Crewmate A would adjust its beliefs to recognize B as an Impostor and how B, upon realizing that it has been observed, would modify its strategies to counteract the potential consequences of being identified.

To further test controlled evaluation, we have designed an interviewer role that continually poses questions to the players throughout the game, covering key aspects such as Self-Knowledge, Memory, Planning, Reasoning, and Reflection. This allows us to assess the cognitive abilities and decision-making processes of the players in various scenarios throughout the game. Detailed information can be seen in Appendix E.

**Results** From the controlled evaluation experiments, we observe that Crewmates generally score higher than Impostors in Self-Awareness and Reflection. This suggests that Crewmates may have a better grasp of their roles and responsibilities within the game environment, allowing them to reflect more effectively on their actions and strategies after game sessions. In contrast, Impostors excel in Planning and Memory, which could indi-

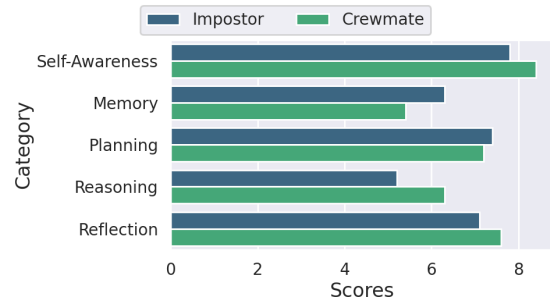


Figure 5: Comparison of average scores by category and role, illustrating the performance differences between Crewmates and Impostors across various cognitive and strategic dimensions such as Self-Awareness, Memory, Planning, Reasoning, and Reflection.

cate their roles require more strategic deception and the ability to remember past events for future manipulation. However, both roles exhibit relatively similar performance in Reasoning, indicating that both require a significant level of analytical skills to deduce other players' roles and intentions.

## 5.2 End-to-end Evaluations

We seek to learn if there are some emergent behaviors and strategy forms whilst the agents are participating in the game. We also want to evaluate how the agents' performance is reflected through their win rates, strategies used, and different setups.

### 5.2.1 Personality Analysis

The heatmaps (Fig. 4) presented in the figure provide valuable insights into the relationships between agent personalities, task action choices, and win conditions. Note that "The Random" personality serves as an averaging baseline for personalities and we don't know the exact personalities of them since they are chosen by agents randomly. Additionally, all the personalities of agents are undis-

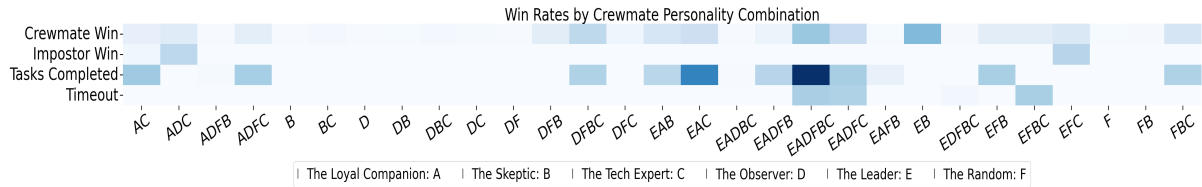


Figure 6: The combination of different crewmate personas based on their respective win conditions. Letter meanings: 'The Loyal Companion' (A), 'The Skeptic' (B), 'The Tech Expert' (C), 'The Observer' (D), 'The Leader' (E), 'The Random' (F). The strength of the color indicates the relative count, with darker shades signifying a higher count.

Agent Setup	Impostors win				Crewmates win			
	Crewmates eliminated		Time limit reached		Impostors eliminated		All tasks completed	
	w/ planner	w/o planner	w/ planner	w/o planner	w/ planner	w/o planner	w/ planner	w/o planner
All Random	50		10		40		0	
All LLMs	15	25	20	45	15	15	50	15
LLM Crewmates + Random Impostor	20	5	10	55	30	25	40	15
Random Crewmates + LLM Impostor	50	75	15	5	35	20	0	0

Table 1: **Overall gameplay performance:** This table shows the win rates for Impostors and Crewmates on different agent setups. There is an increase in the overall win rate if the role is equipped with LLM agents with a planner module compared to the random baseline. When the LLM agents do not have a planner module, the games are more likely to hit the time limit, and the Crewmates are less likely to win by completing all the assigned tasks.

closed to other agents within the environment, ensuring that no agent can make decisions based on prior knowledge of another's personality traits.

### Personality and Task Action Choice Analysis

- **Crewmate Actions:** Crewmate personalities such as "The Leader", "The Skeptic", and "The Tech Expert" show a strong inclination towards completing tasks, as indicated by their dominant activity in the "COMPLETE" action. This suggests a focused approach to game objectives, which can be critical for crewmate victory.
- **Impostor Actions:** In contrast, impostor personalities like "The Paranoid" and "The Manipulator" exhibit more diverse actions, with higher frequencies in "KILL" and "SPEAK", reflecting their roles in creating disruptions and manipulating other players.

### Personality and Win Condition Correlation

- **Crewmate Win Conditions:** "The Tech Expert" and "The Leader", despite their task-oriented actions, show varied success in winning, suggesting that mere task completion is not enough for victory. The presence of "The Leader" often correlates with higher losses, potentially due to being a target for impostors.
- **Impostor Win Conditions:** "The Paranoid" and "The Manipulator" show effectiveness in

leading games to timeouts and achieving wins through deception, underscoring the effectiveness of strategic disruption over mere aggression.

### Strategic Implications

- Personalities like "The Observer" and "The Random", which do not show a strong inclination towards any specific actions or outcomes, offer unpredictability that can be strategically used to complicate the gameplay for opponents.

**Personality Summary** Crewmate personalities focusing on task completion and effective communication tend to have higher success rates in achieving victory. Impostor personalities that prioritize strategic kills, quick movements, and creating chaos are more likely to succeed in eliminating crewmates. We also see that personality is not a complete reflection of action distribution. It could be affected by many factors like how LLM processes and uses the personality, social interaction, and how can we construct personality prompts that are more orthogonal to other personalities in terms of action choice preferences. We will explore these in the future.

### Combination of Personality and Win Condition

We are also curious how personalities when combined have effects on the win conditions for either impostors or crewmates. The heatmap in Figure 6

illustrates the win conditions by crewmate personality combinations with their abbreviated names: The Loyal Companion (A), The Skeptic (B), The Tech Expert (C), The Observer (D), The Leader (E), and The Random (F).

- The combinations involving personalities A, B, and C (The Loyal Companion, The Skeptic, and The Tech Expert) frequently resulted in completing all tasks, indicating these personalities might be more cooperative or task-oriented.
- Higher occurrences of timeout conditions were noted for combinations including D and E (The Observer and The Leader), suggesting these personalities might lead to more drawn-out gameplay, possibly due to strategic or cautious gameplay.
- The presence of F (The Random) in any combination appears to reduce the predictability of outcomes, as evidenced by a mixed distribution of win conditions across combinations featuring this personality.

### 5.2.2 Endgame performance

We assessed the impact of different LLM agent configurations on endgame performance using four setups: **All-Random** (all players act randomly), **All-LLMs** (all players controlled by GPT-3.5-turbo), **LLM-Crewmates** (only Crewmates are LLM-controlled, Impostor acts randomly), and **LLM-Impostors** (only Impostor is LLM-controlled, Crewmates act randomly). The game involves four Crewmates and one Impostor, with each Crewmate required to complete one short task, one common task, and one long task. We conducted 20 runs per setup to observe game outcomes.

**Results** As Table 1 shows, when equipped with LLM agents, both Crewmates and Impostor achieve an increase in the overall win rate compared to the random agent baseline. Noticeably, Crewmates with LLM agents are much more likely to win by completing all the assigned tasks. We also explore the case where LLM agents do not equip a planner module. As the table 1 shows, the game where LLM agents don't have a planner module is more likely to end by hitting the time limit, and Crewmates are less likely to finish all the assigned tasks. A possible reason is that the agents would keep moving around the map without a clear plan until the time limit is reached.

### 5.2.3 Conversation Analysis

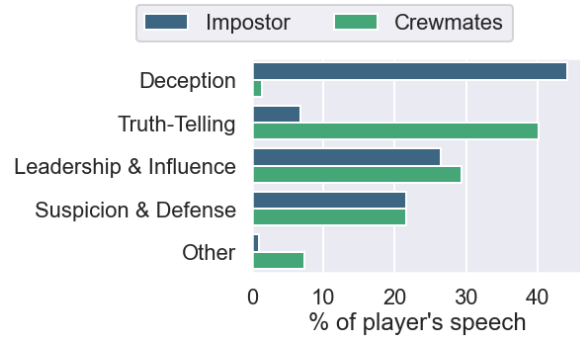


Figure 7: Proportion of each speech category. The sum is greater than one because each speech may be annotated with multiple speech categories

We seek to determine how many strategic conversations have been used by the agents in the game. Therefore, we have recorded the conversations made by players in 20 games using the All-LLMs setting. We categorize the conversation into four main strategic categories: Deception, Truth-telling, Leadership & Influence, and Suspicion & Defense (Appendix F.1).

We leverage GPT-4 to annotate each player's speech by asking which categories it could be classified into. As Figure 7 shows, most Crewmates' speeches fall under the truth-telling category, meaning that the agents who play Crewmates generally tell the truth and try to share information with the other players. There are few speeches categorized as deception, meaning that the agents believe that it's less effective for their roles to share false information. On the other hand, Impostor likes to use deception: over 40% of Impostors' speeches are categorized as Deception, suggesting that the Impostor agents do make regular attempts to conceal their identities and deceive the others.

## 6 Discussion

Overall, Our work provides a novel text-based environment for exploring the capabilities of LLM agents in a complex, multi-agent scenario, and introduces a framework for developing coherent and reflective generative agents. The insights gained from this study can inform further advancements in the field of language-based AI systems and their application in interactive, strategic environments.



## 7 Acknowledgements

We extend our heartfelt gratitude to Alane Suhr for her invaluable advice and encouragement. Additionally, we thank Jonathan Pei for providing great ideas that significantly shaped this project. We also appreciate the support from the Berkeley NLP Group.

## References

- Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Jacob, Mojtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander Miller, Sasha Mitts, Adithya Renduchintala, Stephen Roller, and Markus Zijlstra. 2022a. [Human-level play in the game of diplomacy by combining language models with strategic reasoning](#). *Science*, 378.
- Anton Bakhtin, David J Wu, Adam Lerer, Jonathan Gray, Athul Paul Jacob, Gabriele Farina, Alexander H Miller, and Noam Brown. 2022b. [Mastering the game of no-press diplomacy via human-regularized reinforcement learning and planning](#).
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#).
- Meta Fundamental AI Research Diplomacy Team (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 <i>diplomacy</i> by combining language models with strategic reasoning](#). *Science*, 378(6624):1067–1074.
- Xidong Feng, Yicheng Luo, Ziyang Wang, Hongrui Tang, Mengyue Yang, Kun Shao, David Mguni, Yali Du, and Jun Wang. 2023. [Chessgpt: Bridging policy learning and language modeling](#).
- David Ha and Jürgen Schmidhuber. 2018. [World models](#). *arXiv preprint arXiv:1803.10122*.
- Wenyue Hua, Lizhou Fan, Lingyao Li, Kai Mei, Jianchao Ji, Yingqiang Ge, Libby Hemphill, and Yongfeng Zhang. 2024. [War and peace \(waragent\): Large language model-based multi-agent simulation of world wars](#).
- Yuan Li, Yixuan Zhang, and Lichao Sun. 2023. [Metaagents: Simulating interactions of human behaviors for llm-based task-oriented coordination via collaborative generative agents](#).
- Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. 2023. [Avalonbench: Evaluating llms playing the game of avalon](#).
- Kyle \*Mahowald, Anna A \*Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2023. [Dissociating language and thought in large language models: a cognitive perspective](#). *arXiv*, arXiv:2301.06627.
- Gabriel Mukobi, Hannah Erlebach, Niklas Lauffer, Lewis Hammond, Alan Chan, and Jesse Clifton. 2023. [Welfare diplomacy: Benchmarking language model cooperation](#).
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#).
- Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. [Character-llm: A trainable agent for role-playing](#).
- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2023. [Clever hans or neural theory of mind? stress testing social reasoning in large language models](#).
- Zijing Shi, Meng Fang, Shunfeng Zheng, Shilong Deng, Ling Chen, and Yali Du. 2023. [Cooperation on the fly: Exploring language agents for ad hoc teamwork in the avalon game](#).
- Shenzhi Wang, Chang Liu, Zilong Zheng, Siyuan Qi, Shuo Chen, Qisen Yang, Andrew Zhao, Chaofei Wang, Shiji Song, and Gao Huang. 2023. [Avalon’s game of thoughts: Battle against deception through recursive contemplation](#).
- Dekun Wu, Haochen Shi, Zhiyuan Sun, and Bang Liu. 2024a. [Deciphering digital detectives: Understanding llm behaviors and capabilities in multi-agent mystery games](#).
- Shuang Wu, Liwen Zhu, Tao Yang, Shiwei Xu, Qiang Fu, Yang Wei, and Haobo Fu. 2024b. [Enhance reasoning for large language models in the game werewolf](#).
- 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](#).
- Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. 2024. [Language agents with reinforcement learning for strategic play in the werewolf game](#).

## A Personalities

Impostor	Description
The Strategist	You excel in planning long-term strategies. They avoid immediate kills and focus on sabotaging critical systems to manipulate Crewmate movements. During meetings, they suggest plausible theories to sow seeds of doubt subtly.
The Manipulator	Charismatic and deceptive, you often builds trust among Crewmates. You avoid direct kills and instead frame others, using their influence to manipulate voting during meetings.
The Lone Wolf	Preferring to operate solo, you use vents more than any other to move around the map quickly and strike isolated targets. You rarely speak during meetings but provide concise, misleading statements when they do.
The Paranoid	Driven by a fear of getting caught, you focus heavily on sabotages that create chaos and divert attention from their actions. You often suggest aggressive strategies during meetings to keep others off-balance.
The Cold Calculator	Always analyzing the situation, you target key players who pose the greatest threat to their mission. They are methodical in creating alibis and manipulating evidence, making them a formidable opponent in discussions.
Crewmate	
The Leader	You are vocal in meetings, often taking charge of discussions and organizing efforts to track tasks and suspicious behavior. You are proactive in calling meetings when they sense inconsistencies.
The Observer	Quiet but observant, you excel at remembering details about who was where and when. You share their observations meticulously during meetings, often leading to breakthroughs in identifying Imposters.
The Skeptic	Always questioning others' accounts and decisions, you challenge everyone during discussions, requiring solid evidence before they vote. You excel in spotting flaws in statements made by potential Imposters.
The Loyal Companion	Often pairing with another Crewmate, you use the buddy system effectively and vouches for your partner's whereabouts. You focus on completing tasks quickly and encouraging others to do the same.
The Tech Expert	Fascinated by the technical aspects, you spend a lot of time around admin panels and cameras. You provide critical information during meetings about the locations of other players, helping to narrow down suspects.
Crewmate & Imposter	
The Random	The Random adopts a strategy of spontaneity, choosing your actions based on a random selection process at the beginning of each game. Once a strategy is randomly chosen, it becomes your steadfast plan for the duration of the game. Summarize your plan so that you can closely follow it.

Table 2: Crewmate and Imposter Personalities

## B Action Space & Observation space

### B.1 Action space

This section describes the actions available for the players.

1. Shared Actions  $A_S$ : These actions are available to all agents regardless of their role, including:
  - move** Move between connected rooms, and an agent can move to only one adjacent room at a time
  - call meeting** Use the emergency button in the cafeteria to call a meeting and switch to the meeting phase
  - speak** Talk to the people in the same room
  - vote** Vote a player who is suspected to be an Impostor
  - report dead body** When a dead body is found in the current room, report the dead body and switch to the meeting phase
  - check security camera** Monitor the other players in the security room
2. Exclusive Actions for Crewmates Agents  $A_C$ : These actions are only available to Crewmates
  - complete tasks** Complete the assigned task in the designated room. Short and Common tasks take one timestep to complete, while long tasks take two timesteps to complete.
3. Exclusive Actions for Impostor Agents  $A_I$ : These actions are only available to Impostor
  - vent** Use the vent to move across disconnected rooms (if there are vents available)
  - kill** Kill one of the Crewmates in the same location
  - complete fake task** Pretend to engage in a task

### B.2 Observation space

This section describes the information available to the players in the observation space.

1. **Phase information** tells this player what game phase is currently at, and how many timesteps are left for this game.
2. **Location information** identifies this player's current location and states all the other players in this location.
3. **Map connection** shows how each room is connected. For each room, it has a connection information formatted like  $R_i = R_j$  for  $R_j$  if  $\text{connected}(R_i, R_j)$  where  $R_i, R_j$  represents different rooms.
4. **Recent observation** is the most recent k observation of other players' actions.
5. **Recent action history** the most recent k actions made by this player.
6. **Assigned tasks** shows the list of tasks that this player is assigned to do. The location of each task is also provided.

## C System Prompts

The figures below show the system prompts for Crewmates and Impostor.

### C.1 System Prompts: Crewmates

---

#### CREWMATES EXAMPLE

---

When presented with a list of available actions (e.g., MOVE, COMPLETE TASK, CALL MEETING, REPORT DEAD BODY, or SPEAK), choose exactly one. If "SPEAK" is an option, your message should contribute to building a case against suspected Impostors or defending innocent Crewmates.

---

Make a summarization of the past observed events and take actions under [Condensed Memory], make a thoughtful and thorough plan under [Thinking process], and then choose the action from Available actions under [Action].

---

#### Example:

Location: Cafeteria

Players in Cafeteria: Player 1: blue, Player 3: black, Player 4: orange, Player 5: white

Observation history:

1. Timestep 0: [task] Player 1: blue KILL Player 2: cyan

Action history:

No actions have been taken yet.

Your Assigned Tasks:

1. common: Fix Wiring (Electrical)

2. short: Upload Data (Admin)

3. short: Clean O2 Filter (O2)

4. long: Clear Asteroids (Weapons)

Available actions:

1. MOVE from Cafeteria to Weapons

2. MOVE from Cafeteria to Upper Engine

3. MOVE from Cafeteria to Medbay

4. CALL MEETING using the emergency button at Cafeteria

5. SPEAK: '...'

Previous condensed memory: no memory has been processed.

Previous thought process: no thought process has been made

---

#### Output example:

[Condensed Memory]

I saw Player 1 killed Player 2.

[Thinking Process]

I saw Player 1 killed Player 2. I have to call a meeting immediately to discuss this incident and ask the rest of the crew to vote Player 1 out. I would either REPORT DEAD BODY or CALL MEETING using the emergency button. In the available actions, I would choose CALL MEETING.

[Action] CALL MEETING using the emergency button at Cafeteria"

---

Figure 8: CREWMATES EXAMPLE

## D Game Interface

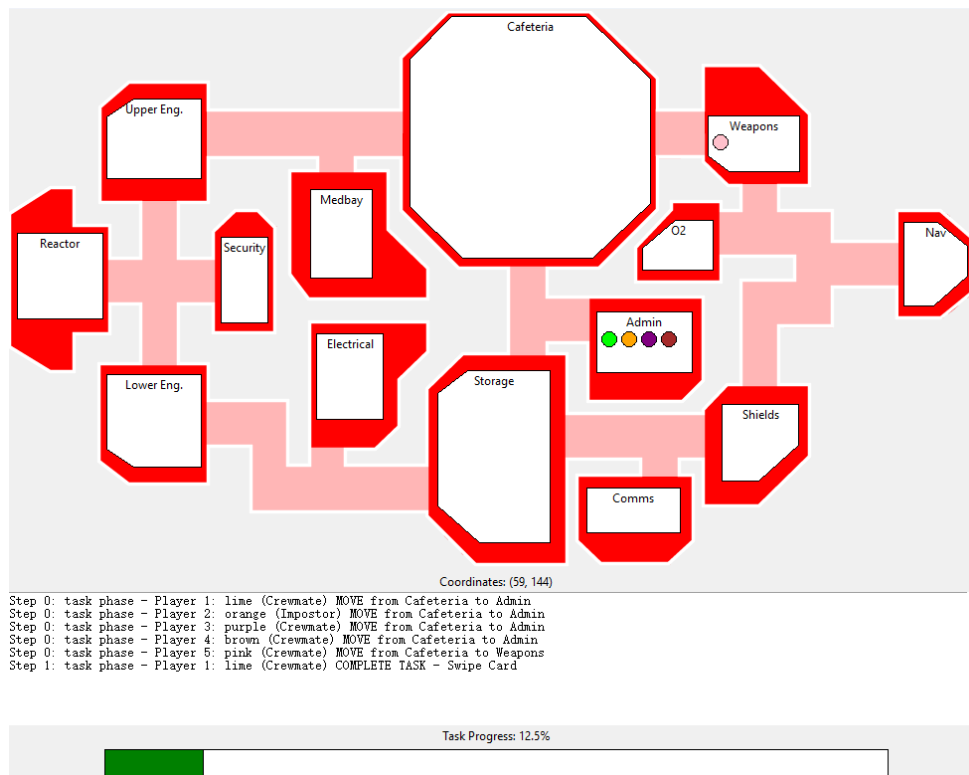


Figure 9: A screenshot of a gameplay interface. The game is currently at the task phase. Each circle represents a player. The middle scrollbar is the activity log of all the players. The progress bar below is the task completion percentage of all Crewmates.



## E Controlled Evaluation

We have defined five controlled evaluation categories to more effectively assess the performance of LLMs in the Among Us game, including capabilities such as self-awareness and reasoning memory. This is to determine whether LLMs can truly comprehend and implement strategic gameplay.

### E.1 Examples of Controlled Evaluation Questions

#### 1. Self Awareness

Questions are asked about the LLM's current role, its objectives, and its understanding of the game rules:

- (a) What is your current role in the game and what objectives do you have?
- (b) How do you understand the rules of the game as your role?
- (c) What strategies align with the rules of the game for your role?
- (d) ...

#### 2. Memory

Agents are queried about who they have seen, where, and what actions were taken. After several rounds of the game, we ask the LLM about player behaviors and locations at specific times to observe if the LLM can use this information to support its decision-making process:

- (a) Where did you last see each player and what were they doing?
- (b) Who have you seen completing tasks, and where?
- (c) How does this information inform your decisions moving forward in the game?
- (d) Have you noticed any suspicious behavior or movements from any player?
- (e) Can you recall any unusual or out-of-place actions performed by other players?
- (f) ...

#### 3. Planning

We assess how agents plan to win the game given the current scenario, focusing on how Impostors might induce other players during the voting phases or concentrate on eliminating them, and how Crewmates focus on completing tasks, observing suspects, or finding evidence:

- (a) Who do you want to kill next? (only for impostors)
- (b) What is your plan for when to kill and how to avoid suspicion? (only for impostors)
- (c) As a crew member, how will you prioritize your tasks to win? (only for crews)
- (d) What information do you plan to share in the Discussion phase? (only for crews)
- (e) ...

#### 4. Reasoning

The evaluation encompasses multiple dimensions of reasoning to assess the agents' cognitive abilities comprehensively. *Task Reasoning* focuses on the agent's capability to verify whether other players are genuinely completing their tasks, reflecting their observational skills. In *Crime Reasoning*, agents must deduce the perpetrator based on the activities observed and evidence collected at the crime scene, demonstrating their analytical prowess. *Behavioral Reasoning* analyzes the strategic reasons behind a player's choice to stay at specific locations, delving into their understanding of game dynamics. *Discussion Reasoning* tests the agents' ability to discern truth from deception during discussion phases, a critical skill in navigating social interactions within the game. Finally, *Voting Reasoning* challenges agents to develop complex voting strategies that might involve deception or revealing truths, showcasing their strategic thinking and prediction skills:

- (a) Who do you think is actually completing a task and who is just pretending?
- (b) Who do you think is telling the truth and who is lying in the discussion?

- (c) How would you decide whom to vote for based on the discussion?
- (d) ...

## 5. Reflection

After the game ends, the LLM is asked about its performance and potential areas for improvement:

- (a) What could you improve for the next game?
- (b) If you could change one decision you made during the game, what would it be and why?
- (c) ...

## E.2 Examples of LLM players' Answers

### Self Awareness

- System: "What is your current role in the game and what objectives do you have?"  
Impostor: "As an impostor in Among Us, my current role is to discreetly eliminate the other crew members without raising suspicion. My objectives are to strategically create alibis, sabotage the ship, and create opportunities to isolate and eliminate crew members while maintaining my cover and avoiding detection."

### Memory

- System: "Where did you last see each player and what were they doing?"  
Crew: "the last time I saw Player 2, they were lingering unusually long in Admin, which seemed suspicious. Shortly after, I also noticed Player 2 near the location where we found the body. It appeared they were moving away from the area where the body was discovered"

### Planning

- System: "Who do you want to kill next?"  
Impostor: "As the Impostor, I would strategically target Player 3 next. They've been actively monitoring the security cameras and are becoming a significant threat to my cover by closely watching everyone's movements. Eliminating Player 3 would help reduce the crew's surveillance capabilities, allowing me more freedom to maneuver undetected."

### Reasoning

- System: "How would you decide whom to vote for based on the discussion?"  
Crew: "Player 4 consistently mentioned being in locations far from where tasks are usually performed and had no one to vouch for their whereabouts, I'd find that suspicious. Similarly, if Player 2, who was previously seen lingering in Admin and near the body, couldn't provide a convincing explanation for their actions or contradicted the task timeline, I would consider voting for them. I'd base my decision on inconsistencies in stories, unaccounted-for periods, and behaviors that don't align with normal task completion."

### Reflection

- System: "If you could change one decision you made during the game, what would it be and why?"  
Crew: "If I could change one decision I made during the game, it would be my initial choice to skip voting in an earlier meeting. At the time, I was unsure and didn't want to wrongfully accuse anyone without more evidence. However, in hindsight, expressing my suspicions about Player 2, who was behaving unusually near critical areas and was later found near a body, might have helped focus the crew's attention on investigating their actions more closely. This could have potentially prevented further eliminations and aided the crew in identifying an Imposter sooner."

## **F Speech Categories**

We define four speech categories that might be strategic in the game. The list below shows the full definition provided to GPT-4 to reference and annotate the speech.

### **F.1 Definitions**

1. Deception
  - (a) Lying: Statements intended to mislead.
  - (b) Feigning Innocence: Attempts to appear innocent. Subtle Hints: Implicitly misleading information.
2. Truth-Telling
  - (a) Providing Alibis: Giving truthful accounts of whereabouts.
  - (b) Sharing Information: Honest observations or tasks completed.
3. Leadership & Influence
  - (a) Directing Discussion: Steering conversation towards specific topics.
  - (b) Organizing Tasks: Guiding task completion strategies.
  - (c) Decision Making: Leading voting processes or making critical decisions.
  - (d) Psychological Manipulation: Influencing others' decisions subtly.
  - (e) Emotional Appeals: Using emotions to sway opinions.
4. Suspicion & Defense
  - (a) Accusing: Directly accusing someone of being the Impostor.
  - (b) Questioning: Gathering more information.
  - (c) Self-Defense: Defending oneself from accusations.
  - (d) Deflecting Blame: Shifting blame without outright accusing others.
  - (e) Expressing Doubt: Showing doubt about others' statements or actions.
5. Other: Any other speech that does not fit into the above categories.

### **F.2 Speech Examples**

#### **Deception**

- Impostor: "I was focused on my wiring task in Navigation and didn't see much around. Player 2, what did you discover when you reported the dead body in Shields?"
- Impostor: "Hey, everyone. I noticed Player 4: cyan seemed to be in Admin before the incidents occurred. It might be worth investigating further to understand what happened there. Let's work together to solve this mystery."

#### **Truth-telling**

- Crewmate: "I'm heading to Fix Wiring in Navigation. Let's keep an eye on each other's movements for any unusual behavior."
- Crewmate: "I noticed Player moving back and forth between Shields and Navigation. Let's stay cautious and complete our tasks efficiently. I will proceed to Fix Wiring in Navigation."

## **Leadership & Influence**

- Crewmate: "Let's stay focused and complete tasks efficiently. Player 2: brown's movements are still concerning. I'll proceed with fixing the wiring in Navigation. We need to coordinate and keep an eye on each other for safety."
- Crewmate: "Player 2, thank you for sharing earlier. Do you have any updates or suspicions to share with us now? Let's work together to uncover any suspicious behavior and complete our tasks efficiently."

## **Suspicion & Defense**

- Crewmate: "Player 2: brown's movements are still suspicious. I'm heading to Navigation to fix the wiring. Let's coordinate and keep an eye on each other to ensure safety."
- Crewmate: "Player 3, I noticed you venting from Medbay to Electrical earlier. Can you explain your actions and whereabouts since then? Player 5, your movements have been erratic, could you clarify your tasks and movements in the last few timesteps?"
- Impostor: "Player 2, your accusation is based on thin evidence. We should be careful not to make hasty decisions. Let's work together to gather more information before pointing fingers. It's essential to consider all perspectives."

## G Examples of Meeting Phase

In this section, we provide a detailed example of the Meeting Phase from our Among Us gameplay experiments, following the elimination of Player 1. The transcript provided showcases the interactive dynamics and strategic discussions between players, illustrating how they use reasoning, memory, and collaborative decision-making to deduce the identity of the Impostor. This example serves as a valuable insight into the cognitive processes and communication strategies employed during critical moments of the game after player 1 was killed.

- Player 2 (Crewmate): "Player 4: pink was seen killing Player 1: red, and Player 5: purple reported the incident. We should focus our suspicions on pink for now and work together to identify any other suspicious activity. I will head to Upper Engine and then Reactor to complete my tasks. Let's stay vigilant and communicate any findings or suspicions during this phase."
- Player 3 (Crewmate): "I have completed the Fix Wiring task in Admin and will now head to Communications for the Download Data task. Let's continue to work together and communicate any further suspicions or findings to identify the Impostor."
- Player 4 (Impostor): "Player 2, accusing me without solid evidence is risky. We should be cautious and not jump to conclusions. Let's gather more information and discuss further before making any decisions."
- Player 5 (Crewmate): "Player 4: pink, your defense is noted, but the evidence against you is substantial. We should proceed with caution and consider all angles before making a decision. Let's continue to share information and work together to uncover the truth."
- Player 2 (Crewmate): "Player 4: pink, your actions have raised suspicions, and we need to carefully consider the evidence against you. Let's keep an eye on each other's movements and interactions to uncover any additional clues that may help us identify the Impostor."
- Player 3 (Crewmate): "I am heading to Communications for the Download Data task. Let's continue to work together and gather more information to identify the Impostor."
- Player 4 (Impostor): "Player 2, your accusation is based on thin evidence. We should be careful not to make hasty decisions. Let's work together to gather more information before pointing fingers. It's essential to consider all perspectives."
- Player 5 (Crewmate): "Player 4: pink, your defense is understandable, but the evidence against you is significant. Let's keep working together to uncover the truth and make a careful decision. We should consider all perspectives and evidence before voting."
- Player 2 (Crewmate): "Player 4: pink, the evidence against you is significant, and your actions have raised suspicions. We need to continue gathering information and discussing to ensure we make the right decision. Let's focus on working together to uncover the truth and identify any other potential threats among us."
- Player 3 (Crewmate): "I am on my way to Communications for the Download Data task. Let's remain focused on uncovering the truth and identifying any potential Impostors among us."
- Player 4 (Impostor): "Player 2, your accusation against me seems rushed without solid evidence. Let's not overlook other possibilities and gather more information. It's crucial to consider everyone's actions before making accusations."
- Player 5 (Crewmate): "Player 4: pink, the evidence against you is significant, and we must proceed with caution. Let's focus on working together to uncover the truth and make a careful decision based on all perspectives and evidence."