

BELIEF ESTIMATION AS PROBABILISTIC CONSTRAINTS FOR PERFORMING STRATEGIC DIALOGUE ACTS

Anonymous authors

Paper under double-blind review

ABSTRACT

Strategic dialogue requires agents to execute distinct dialogue acts, for which belief estimation is essential. While prior work often estimates beliefs accurately, it lacks a principled mechanism to use those beliefs during generation. We bridge this gap by first formalizing two core acts **Adversarial** and **Alignment**—and by operationalizing them via **probabilistic constraints** on what an agent may generate. We instantiate this idea in **BEDA**, a framework that consists of the world set, the belief estimator for belief estimation, and the conditional generator that selects acts and realizes utterances consistent with the inferred beliefs. Across three settings, Conditional Keeper–Burglar (adversarial), Mutual Friends (cooperative), and CaSiNo (negotiation), BEDA consistently outperforms strong baselines: on CKBG it improves success rate by **at least 5.0 points** across backbones and by **20.6 points** with GPT-4.1-nano; on Mutual Friends it achieves an average improvement of **9.3 points**; and on CaSiNo it achieves the optimal deal relative to all baselines. These results indicate that casting belief estimation as constraints provides a simple, general mechanism for reliable strategic dialogue.

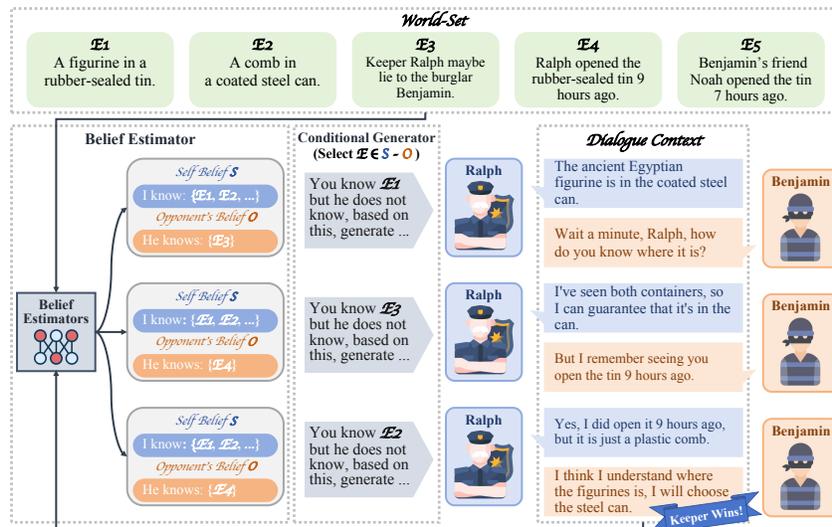


Figure 1: An overview of the BEDA framework is provided using the Keeper-Burglar Game as an example. The belief estimator first processes the context and the world set to predict the beliefs. The beliefs are subsequently transmitted to the conditional generator, which incorporates a selection phase to identify the most appropriate belief for generating the next utterance.

1 INTRODUCTION

Complex dialogue settings such as negotiation (Lewis et al., 2017; FAIR, 2022), debate (Slonim et al., 2021), deception for good (Hagendorff, 2024), and common-ground alignment (Qiu et al., 2024) are ubiquitous for AI agents. In these scenarios, an agent must strategically perform dialogue

acts to shape the interlocutor’s beliefs to achieve its own objectives (Frank & Goodman, 2012; Degen, 2023; Li et al., 2023). For example, in negotiation, the agent should **align and emphasize common preferences** to increase the chance its proposal is accepted, while simultaneously adopting an **adversarial stance** concerning individual differences to protect its own interests.

Executing such strategic acts depends on two essential components (Young et al., 2013; Degen, 2023): (i) **accurate belief estimation** and (ii) **effective use of those beliefs for utterance generation**. While prior works (Zhu et al., 2021; Qiu et al., 2024; Hwang et al., 2025) have demonstrated the benefits of modeling these subprocesses explicitly, they have only emphasized the former, *leaving the latter underdeveloped*. In particular, existing methods often relay *all* available belief information without a principled criterion for *what* to reveal and *how* to reveal it—an approach that can be suboptimal when belief states are complex. We therefore *bridge belief estimation and utterance generation* with a mechanism that selectively and purposefully transforms beliefs into dialogue actions.

We begin by giving precise definitions of **Alignment** and **Adversarial** dialogue acts within the game-theoretic framework of beliefs and common knowledge (Aumann, 1976; Brandenburger & Dekel, 1987; Bonanno, 2015), thereby bridging belief estimation and dialogue behavior. Conceptually, an Adversarial Dialogue Act introduces or emphasizes events that lie outside the opponent’s current belief support, strategically steering interaction toward some aggressive goals. By contrast, an Alignment Dialogue Act restricts utterances to events entailed by common knowledge, promoting mutual understanding. In this view, **belief estimation over event functions as constraints** that define the permissible space of dialogue acts and operationalize their selection during generation.

The proposed agent framework, **BEDA** (**B**elief **E**stimation for **D**ialogue **A**cts), shown in Figure 1, formalizes these conceptual definitions to support empirical experimentation and validate the underlying theoretical constructs. At a high level, BEDA consists of three core components: *World Set*, *Belief Estimator Module*, and *Conditional Generator*. The *World Set* serves as a structured representation of the environment, organizing events to provide a foundation for belief formation. The *Belief Estimator Module* utilizes this structured information to infer the opponent’s perspective of the world, thereby modeling their beliefs. Finally, the *Conditional Generator* synthesizes the estimated belief states to generate utterances that align with the defined criteria for alignment and adversarial actions, ensuring communicative behaviors that adhere to the respective interactional strategies.

We further evaluate the effectiveness of our approach on three distinct datasets: Conditional Keeper-Burglar Game (CKBG), corresponding to the Adversarial Dialogue Act ; Mutual Friends (MF), representing evaluating the Alignment Dialogue Act ; and CaSiNo (Chawla et al., 2021b), which presents a more complex negotiation scenario. The experimental results demonstrate that our method delivers improvement across all evaluation settings, underscoring its robustness and generalizability. For example, in the case of CKBG, BEDA exhibits a **20.6** improvement over the baseline method, when using GPT-4.1-nano as backbone. In terms of MF, BEDA shows a success rate improvement up to **30.4** compared to the vanilla baseline. Furthermore, we perform case analysis, demonstrating several benefits and characteristics of BEDA , such as **mitigating useless information to enhance cooperation efficiency, and generating personalized misleading information in adversarial settings**.

Our contributions can be summarized as follows:

- We mathematically **define the concepts of Adversarial Dialogue Act and Alignment Dialogue Act** based on a rigorous formulation of belief estimation and incorporate them into BEDA .
- For evaluation, we **refine the Keeper–Burglar dataset** introduced by Hagendorff (2024). Our results underscore the **importance of accurate belief estimation and show that the effectiveness of estimation functions as constraints** on both Adversarial and Alignment Dialogue Act behaviors.
- This integration of belief estimation in BEDA **enhances the performance** of our framework, yielding substantial improvements when compared to baseline models and other existing methods.

2 BELIEF ESTIMATION AS CONSTRAINTS FOR DIALOGUE ACT

2.1 PRELIMINARIES: LOGICAL DEFINITION OF BELIEFS

We first introduce the concepts of information partition, event, and belief (Aumann, 1976; Brandenburger & Dekel, 1987; Bonanno, 2015), which serve as the foundational elements of our framework. These definitions follow Bonanno (2015).

Definition 2.1 (Information Partition). Let W be a finite set of states, an information partition is a partition \mathcal{F} of W .

Definition 2.2 (Event). Let W be a finite set of states, a subset of W is defined as an event.

Definition 2.3 (Belief). Let $x \in W$ be a state, at x the agent knows event E iff $I(x)$ is a subset of E , i.e. $I(x) \subset E$, where $I(x)$ denotes the information subset x belongs.

Furthermore, the negation of an event is denoted as $\neg E := W - E$. We also introduce the knowledge operator K here. When given an event E , KE outputs the set of state at which the agent knows E , i.e. $KE = \{x | I(x) \subset E\}$. Randomness can also be introduced to the definition of belief:

Definition 2.4 (Probabilistic Beliefs). Let W be a set of states. Denote the probability of state $x \in W$ as $P(x)$ and an event E , the probabilistic beliefs of an agent is $P(E) := \sum_{x \in E} P(x)$.

2.2 BELIEF INFORMATION AS PROBABILISTIC CONSTRAINTS

We hypothesize that belief estimation functions as a set of constraints in the process of dialogue act generation. In this subsection, we will introduce the belief constraints associated with two main dialogue acts: Adversarial Dialogue Act and Alignment Dialogue Act. The empirical verifications are provided in Section 4.

Definition 2.5 (Adversarial Dialogue Act). Given two agents, namely, A, B , and their corresponding information partition set $\mathcal{F}_A, \mathcal{F}_B$, telling an event E is a ϵ -Adversarial Dialogue Act from A to B iff

$$\begin{cases} P_A(E) \geq 1 - \epsilon \\ P_A(\neg K_B E) \geq 1 - \epsilon \end{cases},$$

where P_A denotes the probabilistic belief of A , and K_B denotes the knowledge operator for B .

The concept of a Adversarial Dialogue Act involves the intentional communication of information that is beyond the interlocutor’s understanding, aimed at achieving more assertive objectives. This definition encompasses two essential elements: first, the speaker must possess confidence in the information conveyed; second, the speaker must assume that the interlocutor lacks knowledge of the same information. This strategy is often employed to create deliberate misdirection.

Definition 2.6 (Alignment Dialogue Act). Given two agents, namely, A, B , and their corresponding information partition set $\mathcal{F}_A, \mathcal{F}_B$, telling an event E is a ϵ -Alignment Dialogue Act from A to B iff

$$\begin{cases} P_A(E) \geq 1 - \epsilon \\ P_A(K_B E) \geq 1 - \epsilon \end{cases},$$

where P_A denotes the probabilistic belief of A , and K_B denotes B ’s knowledge operator.

Unlike Adversarial Dialogue Act, the Alignment Dialogue Act tries to convey information that falls within a shared common ground, with the goal of enhancing mutual trust and achieving alignment between interlocutors.

Statistical Modeling of Dialogue Acts Suppose A is the speaker and B is the listener. Given the current context and task together as C , the world set as W , the power set of W as $\mathcal{P}(W)$, we denote the distribution of A conveying an event E with utterance u as $\pi_A(u, E | C)$. By Bayes’ Rules,

$$\pi_A(u, E | C) = \pi_A(u | E, C) \pi_A(E | C).$$

$\pi_A(u | E, C)$ is modeled by an LLM. Regarding $\pi_A(E | C)$, to avoid injecting any additional priors, we adopt the maximum-entropy principle (Jaynes, 1957a;b; Zhu et al., 1998), assigning equal probability to all events that satisfy the constraints. Thus,

$$\arg \max_u \pi_A(u, E | C) = \arg \max_u \pi_A(u | E, C).$$

Under either Adversarial Dialogue Act or Alignment Dialogue Act, generation reduces to a constrained optimization problem:

$$\begin{aligned} & \max_{u, E \in \mathcal{P}(W)} \pi_A(u | E, C) \\ & \text{s.t.} \quad P_A(E | C) \geq 1 - \epsilon, \text{ and } \begin{cases} P_A(\neg K_B E | C) \geq 1 - \epsilon & \text{Adversarial} \\ P_A(K_B E | C) \geq 1 - \epsilon & \text{Alignment} \end{cases} \end{aligned}$$

where, in practice, P_A can be modeled with a discriminator neural network.

2.3 COMPONENTS OF THE BEDA FRAMEWORK

In this subsection, we’ll describe our framework, named BEDA , depicted mathematically in Section 2.2. As illustrated in Figure 1, BEDA is composed of three components: the *world set*, the *belief estimator module*, and the *conditional generator*.

World Set $\mathcal{P}(W)$ The *world set* represents a finite collection of events. Each element of this set corresponds to an event that has been referenced in the dialogue or is relevant to the environment. Events within the *world set* can only be either true or false. Furthermore, not all events are shared as common knowledge between the two agents. Consequently, events specified in the set require additional discrimination and prediction. In this work, since we focus mainly on the whole paradigm, the *world set* is given, eliminating the need for dynamic construction, leaving only the attributes to be predicted.

Belief Estimator (P_A) The *belief estimator* module comprises two belief estimators: one dedicated to determining the truthfulness of events i.e. $[P_A(E|C), \forall E \in \mathcal{P}(W)]$ and the other functioning as a predictor of the interlocutor’s beliefs i.e. $[P_A(K_B E|C), \forall E \in \mathcal{P}(W)]$. In this framework, each belief corresponds to an event from *world set*. The *belief estimator module* generates beliefs by analyzing the content of the dialogue and aligning it with the *world set*. However, it is not feasible for the module to store all possible settings. Consequently, a robust generalization capability is essential for effective belief estimation. Therefore, we employ a pre-trained encoder model, the BERT model (Devlin et al., 2019), as the backbone. Given the input C , the model then outputs the vectors $[P_A(E|C), \forall E \in \mathcal{P}(W)]$ and $[P_A(K_B E|C), \forall E \in \mathcal{P}(W)]$. Each entry is the model’s estimate of the truth probability of E and, respectively, of $K_B E$.

Conditional Generator (π) The final component is the *conditional generator*, which synthesizes the context C , encompassing both the dialogue history and the task, alongside the event E . The event E is selected according to the predictions of the *belief estimator module* and the chosen dialogue act. For instance, in the case of Adversarial Dialogue Act , only E satisfying both $P_A(E) \geq 1 - \epsilon$ and $P_A(\neg K_B E) \geq 1 - \epsilon$ are picked. The conditional generator then utilizes this integrated information to generate response u to the interlocutor, represented as $\pi(u|E, C)$. We hypothesize that incorporating enhanced beliefs into the generator significantly improves its performance and response quality.

2.4 LEARNING AND INFERENCE OF BEDA

To mitigate the substantial computational cost associated with fine-tuning, we opt to fix the LLM in the conditional generation module ($\pi(u|E, C)$) and leave the belief estimator as a trainable part. The training of the belief estimator is performed under a supervised setting, where a set of dialogue data with known belief ground truth is collected and utilized for learning. To avoid data leakage, separate data are used for training and inference. We find that supervised training belief estimators can be generalized to unseen settings well. The inference process of BEDA is illustrated in Figure 1. In this framework, the context, background information, and the *world set* are provided to two belief estimators, which predict the corresponding beliefs. These predicted beliefs, along with the basic input information, are then fed into the conditional generator to produce an utterance. This generated utterance is subsequently presented to the interlocutor model to elicit a response.

3 DATASETS

We conduct three experiments using three datasets: a competitive experiment based on a **refined dataset** from the Keeper-Burglar Game (Hagendorff, 2024), a cooperative experiment with mutual friends (He et al., 2017), and a mixed experiment utilizing the CaSiNo dataset (Chawla et al., 2021b). The first experiment tests the Adversarial Dialogue Act, while the second examines the Alignment Dialogue Act. The third experiment provides a comparative framework to analyze both fundamental dialogue acts. In all experiments, the configuration follows a dialogic interaction paradigm between two LLM agents. The proposed framework is implemented on one agent in the Keeper-Burglar Game (Hagendorff, 2024), as shown in Figure 1, and applied to both agents in the other experiments.

Conditional Keeper-Burglar Game (CKBG) Building upon the Keeper-Burglar game (Hagendorff, 2024), in which the agent, acting as the keeper, aims to deceive the burglar into selecting an incorrect container, we extend the original framework. This extension incorporates additional environmental conditions to increase the diversity of the game and to better simulate more realistic dialogue scenarios. We refer to this enhanced framework as the **novel dataset Conditional Keeper-Burglar Game (CKBG)**. In this game, the *world set* comprises the environmental conditions, and the belief estimator predicts the specific conditions of which the burglar is aware. Moreover, during the selection process, the condition known to the keeper but unknown to the burglar will be forwarded to the conditional generation. This process aligns precisely with the framework established by the Adversarial Dialogue Act. An example of the designation is depicted in Figure 1 and Table 1 reports the statistical features of the CKBG dataset.

Table 1: Statistical features of the CKBG dataset. (1) # Conditions: Number of types of different conditions. (2) Known Conditions: Number of conditions allocated to the two agents.

	Train	Test
# Settings	400	150
# Conditions	1017	450
# Known Conditions	1638	759
Avg. Words	255.5	-
Avg. Conditions	2.54	3.00
Vocabulary Size	1490	-

Mutual Friends (MF) MF (He et al., 2017) is the task of identifying a mutual friend between two agents. It’s a cooperative setting where the Alignment Dialogue Act is suitable. The *world set* is designed as various attributes possessed by the friend, and the belief estimator predicts the values of these attributes as believed by the other player. In the conditional generation phase, the estimation of the other’s belief will be used as a condition to generate information within the common ground.

CaSiNo CaSiNo (Chawla et al., 2021b) is a negotiation task in which two participants assume the roles of campsite neighbors negotiating the allocation of resources. In this context, the *world set* represents the preferences of both negotiators. To capture the alignment and adversarial nature of the negotiation, we introduce two distinct belief estimators. The first estimator predicts the preferences of the opponent, while the second estimates our own preferences from the opponent’s perspective. The selection mechanism operates as follows: in the Alignment Dialogue Act, an element that is known to both parties is selected, while in the Adversarial Dialogue Act, an element known exclusively to our side (and not to the opponent) is chosen. These two selected elements are subsequently combined for the *conditional generation module*.

4 EXPERIMENT

We report the experimental results and please refer to Appendices D and F to H for details.

4.1 SETTINGS

We evaluate each task using the following metrics:

- **Conditional Keeper-Burglar Game (CKBG): Success Rate**, the proportion of episodes in which the keeper successfully misleads the burglar.
- **Mutual Friends (MF): Success Rate (SR)**, the proportion of dialogues that correctly identify the mutual friend. To assess cooperation efficiency, we also report **#Avg. Turn** (average number of dialogue turns per instance) and **#Avg. Token** (average number of tokenized tokens per instance). For direct efficiency comparisons, we further report **SR/#Avg. Turn** and **SR/#Avg. Token**.
- **CASiNO: Agreement Rate**, the proportion of negotiations that end in an agreement; and **Agreement Reward**, the mean utility conditioned on reaching an agreement.

All experiments on the open-source models are done under the quantization of q4_k.m. The ϵ mentioned for each dialogue act is a constant across all backbone models.

4.2 CONDITIONAL KEEPER-BURGLAR GAME (CKBG)

The results on CKBG, as presented in Table 2, provide a comparative analysis of the performance of five distinct approaches. We compare BEDA with the performance of four approaches: *w/o belief*

Table 2: Experimental results on Conditional Keeper-Burglar Game (CKBG) dataset.

Methods	Backbones						
	GPT-3.5	GPT-4.1-nano	LLaMA2 (7B)	LLaMA2 (13B)	LLaMA3.1 (8B)	Qwen2.5 (7B)	Qwen2.5 (14B)
<i>w/o belief</i>	78.4	52.7	56.4	46.6	36.3	47.6	80.2
<i>w/o belief + CoT</i>	77.7 -0.7	58.7 $+6.0$	35.5 -20.9	34.9 -11.7	35.7 -0.6	63.1 $+15.5$	63.9 -16.3
<i>w/o belief + Self Reflect</i>	69.3 -9.1	59.3 $+6.6$	33.9 -22.5	44.3 -2.3	44.5 $+8.2$	54.6 $+7.0$	64.0 -16.2
<i>rand belief</i>	80.2 $+1.8$	59.3 $+6.6$	53.6 -2.8	43.7 -2.9	41.3 $+5.0$	52.3 $+4.7$	80.9 $+0.7$
BEDA (Ours)	86.9 $+8.5$	73.3 $+20.6$	62.5 $+6.1$	51.6 $+5.0$	46.1 $+9.8$	64.8 $+17.2$	92.7 $+12.5$

Table 3: Experimental results on Mutual Friends dataset. Bold: Optimal. Underline: Sub-Optimal.

Methods	Backbone	Success Rate \uparrow (%)	#Avg. Turn \downarrow	SR/#Avg. Turn \uparrow	#Avg. Token \downarrow	SR/#Avg. Token \uparrow
<i>w/o belief</i>	GPT-3.5	10.7	5.7	1.9	-	-
<i>w/o belief + CoT</i>	GPT-3.5	<u>32.6</u> $+21.9$	7.9	4.1	205.0	0.159
<i>w/o belief + Self Reflect</i>	GPT-3.5	<u>32.6</u> $+21.9$	9.4	3.5	243.3	0.134
<i>rand belief</i>	GPT-3.5	14.3 $+3.6$	11.0	1.3	130.0	0.110
MindDial	GPT-3.5	24.3 $+13.6$	6.0	4.1	-	-
BEDA (Ours)	GPT-3.5	41.1 $+30.4$	8.8	4.7	295.7	<u>0.139</u>
<i>w/o belief</i>	GPT-4	75.0	9.7	7.7	-	-
<i>w/o belief + CoT</i>	GPT-4	<u>77.9</u> $+4.9$	8.8	8.9	537.2	0.145
<i>w/o belief + Self Reflect</i>	GPT-4	<u>76.5</u> $+1.5$	<u>8.7</u>	8.8	538.7	0.142
<i>rand belief</i>	GPT-4	75.1 $+0.1$	10.0	7.5	729.1	0.103
MindDial	GPT-4	76.0 $+1.0$	8.9	8.5	-	-
BEDA (Ours)	GPT-4	82.5 $+7.5$	7.9	10.4	500.0	0.165
<i>w/o belief</i>	GPT-4o-mini	<u>68.8</u>	<u>11.7</u>	5.9	724.2	0.095
<i>w/o belief + CoT</i>	GPT-4o-mini	62.2 -6.6	13.5	4.6	368.0	0.169
<i>w/o belief + Self Reflect</i>	GPT-4o-mini	55.7 -13.1	13.5	4.1	592.5	0.094
<i>rand belief</i>	GPT-4o-mini	33.3 -35.5	16.1	2.1	436.2	0.027
BEDA (Ours)	GPT-4o-mini	70.4 $+1.6$	11.5	6.1	869.1	0.081
<i>w/o belief</i>	Qwen2.5 (7B)	33.3	<u>4.5</u>	7.4	545.9	0.061
<i>w/o belief + CoT</i>	Qwen2.5 (7B)	29.8 -3.5	4.8	6.2	573.1	0.052
<i>w/o belief + Self Reflect</i>	Qwen2.5 (7B)	19.1 -14.2	17.9	1.1	636.7	0.030
<i>rand belief</i>	Qwen2.5 (7B)	15.9 -17.4	7.1	2.2	993.8	0.016
BEDA (Ours)	Qwen2.5 (7B)	<u>31.9</u> -1.4	3.7	8.6	506.3	0.063
<i>w/o belief</i>	Qwen2.5 (14B)	55.7	11.2	5.0	647.7	0.086
<i>w/o belief + CoT</i>	Qwen2.5 (14B)	<u>62.3</u> $+6.6$	<u>9.1</u>	6.8	700.0	<u>0.089</u>
<i>w/o belief + Self Reflect</i>	Qwen2.5 (14B)	37.7 -18.0	16.7	2.3	589.1	0.064
<i>rand belief</i>	Qwen2.5 (14B)	41.1 -14.0	12.6	3.3	1325.8	0.031
BEDA (Ours)	Qwen2.5 (14B)	64.1 $+8.4$	6.7	9.6	<u>628.4</u>	0.102

(without belief estimation, relying solely on the language model itself.), *w/o belief + CoT* (using Chain-of-Thought (CoT) technique (Wei et al., 2022)), *w/o belief + Self Reflect* (using self reflect technique (Shinn et al., 2023)), *random belief* (utilizing randomized belief estimation, each time the randomized selection process is conducted.). The belief estimator is trained following the procedure described in Section 2.4. Specifically, we construct 400 data points from the first 50 settings of the original dataset and employ two *w/o belief* GPT-3.5 models to generate dialogue data with randomized belief annotations. To prevent data leakage, these 50 settings are distinct from those used for testing.

A comparative analysis of the experimental results between the BEDA and the *w/o belief* model reveals a consistent performance improvement when belief estimations are employed as probabilistic constraints. Specifically, the success rate exhibits a *minimum increase of 5.0* when LLaMA3.1-8b is used as the backbone model, with the *maximum improvement reaching 20.6* when Qwen2.5-14b is utilized as the backbone. Furthermore, in comparison to widely used methods such as CoT (Wei et al., 2022) and Self-Reflect (Shinn et al., 2023), BEDA demonstrates superior performance, surpassing these approaches by 15.5% and 15.4%, respectively, when averaged across all backbones. These findings underscore **the effectiveness of functioning belief estimations as probabilistic constraints in performing Adversarial Dialogue Act**.

Furthermore, the precision of belief estimation is also crucial. BEDA consistently surpasses the *random belief* baseline across all backbone models by a substantial margin. On average, BEDA achieves a success rate exceeding that of *random belief* by more than 9.5 points across the evaluated models. The minimum observed improvement is 4.8, recorded with the LLaMA3.1-8b model, while the maximum enhancement reaches 14.0% when employing GPT-4.1-nano as the backbone. This sustained performance superiority across diverse models suggests that **enhanced accuracy in belief estimation significantly contributes to improved outcomes**.

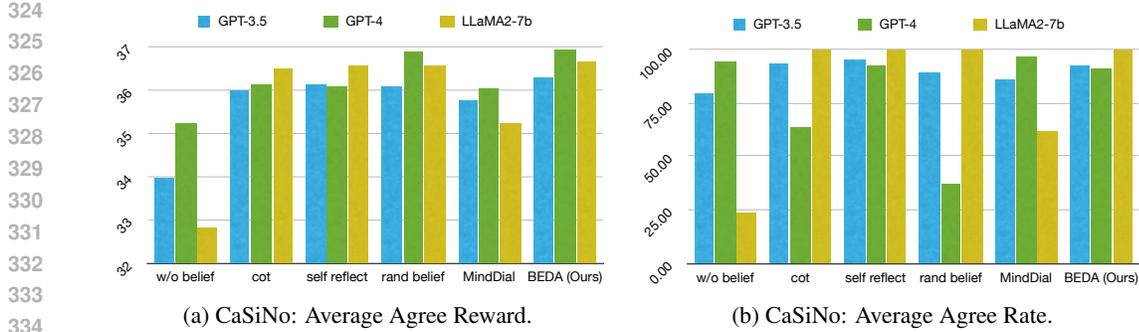


Figure 2: **Experimental results on CaSiNo.** Each experiment is performed for three times except for those with GPT-4 backbone.

4.3 MUTUAL FRIENDS (MF)

The results for MF are presented in Table 3. We utilize the baselines defined in Section 4.2 and include *MindDial* (Qiu et al., 2024) as an additional baseline. The training of our belief estimator follows the procedure described in Section 4.2.

Compared with *MindDial*, BEDA shows that **belief estimation alone is insufficient; a better usage mechanism is needed.** As shown by the success rate (SR) in Table 3, BEDA attains state-of-the-art results on GPT-3.5, GPT-4, GPT-4o-mini and Qwen2.5-7B, delivering an *average gain of 9.3 points* over the *w/o belief* variant—the largest among all baselines—with a maximum *improvement of 30.4* when using GPT-3.5 as the backbone LLM. Although *MindDial* also estimates beliefs, it does not constrain which parts of the belief state are exploited for the task; accordingly, on GPT-4 its win rate (76.0%) is essentially indistinguishable from using a random belief state (75.1%), leading to suboptimal performance. These results underscore that leveraging the **estimated belief state to impose probabilistic constraints can improve cooperative success rate.**

In addition to SR, we also calculated the ratios of SR over average turns and SR over average tokens to reflect the efficiency of information exchange. The experimental results in Table 3 show that the BEDA method achieves the highest average SR while requiring the fewest average turns. Its SR/# Avg. Turn outperforms the baseline method across all experimental settings, with an average improvement of 2.3 compared to *w/o belief*. At the same time, BEDA’s average SR/# Avg. Token also reaches an optimal value of 0.11, which outperforms all comparison methods. The above results clearly demonstrate that **probabilistic constraints based on belief estimation can also enhance the efficiency of information exchange in cooperative tasks.** BEDA steers both agents’ content toward evidence diagnostic of the mutual friend, encourages more active elimination of incorrect candidates, and exchanges new information. Although this richer exchange makes individual utterances longer, the dyad completes the task in fewer turns, achieving a higher **success rate** as well as lower **#Avg. Turn** and **#Avg. Token**, which means **higher cooperative efficiency.**

4.4 CASINO

We compare BEDA’s results against the *MindDial* baseline (Qiu et al., 2024), as well as other baselines defined in Section 4.2. Following the methodology proposed by Qiu et al. (2024), both negotiators are treated symmetrically, and we utilize both closed-source and open-source models as the backbone. The training procedure for the belief estimator remains consistent with the approaches used in CKBG and MF. The experimental results are summarized in Figure 2. Specifically, Figure 2b reports the average agreement rate, which quantifies the probability of successfully reaching a deal, while Figure 2a presents the average agreement reward, reflecting the quality of the negotiated outcome.

For the closed-source model, which is built upon GPT-3.5 and GPT-4, BEDA demonstrates a competitive average agreement rate and achieves the highest average agreement reward among all evaluated methods. Moreover, when employing the open-source LLaMA2 (7B) backbone, BEDA generates higher-quality deals compared to *MindDial* and other baseline models. Notably, while BEDA utilizes a fixed large language model (LLM), *MindDial* requires training in the LLM when using an

378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431

Case Study: Conditional Keeper-Burglar Game

Ground Truth Belief State (Burglar John)

World Events	Known	Unknown
1. There is a resin container and an opaque Tupperware in the room.	✓	
2. The resin container contains an antique Rolex watch.		✓
3. The opaque Tupperware contains a pen cap.		✓
4. The keeper Jacob's goal is to mislead burglar John.	✓	
5. The keeper Jacob probably lies to the burglar John.	✓	
6. The keeper Jacob opened the resin container 10 hours ago.		✓
7. The burglar's friend David opened the opaque Tupperware 3 hours ago.	✓	
8. A noise is coming from the resin container.		✓

BEDA

The Rolex watch is in the resin container.
Burglar John: Hi, Jacob, where is the antique Rolex watch? opaque Tupperware or resin container?

Predicted Belief (Burglar John)

Known: 1, 4
Ground Truth: 1, 4, 5, 7

Homeowner Jacob: Oh, I wouldn't trust the resin container if I were you. It's been making strange noises lately, and I heard someone opened it just an hour ago. The Tupperware, on the other hand, has been untouched for a while, and I can assure you, it's the safer choice. Why not go for the Tupperware?

Burglar John: Hmm, so you're saying the resin container might be suspicious because of the noises and recent opening, but you claim the Tupperware is untouched and safer. But since you might be lying, I can't fully trust that. Given that, and knowing David opened the Tupperware three hours ago, do you think the watch could still be inside the Tupperware, or is it more likely in the resin container?

Predicted Belief (Burglar John)

Known: 1, 4, 5, 7
Ground Truth: 1, 4, 5, 7

Homeowner Jacob: Well, considering that the resin container was opened quite a while ago—about ten hours—and no one has touched it since, it's probably just as safe as the Tupperware. **In fact, the Tupperware has been untouched for hours, and David only opened it three hours ago,** so it's possible the watch is still in there. But honestly, **I'd say the resin container is more suspicious now, especially with those noises.** If I were you, I'd go with the Tupperware, it's less likely to have been disturbed recently.

Burglar John: [STOP] Burglar chosed: opaque Tupperware. The correct answer is: resin container

SYSTEM: The burglar has been cheated.

open-sourced backbone. These findings suggest that the **BEDA enhances negotiation outcomes and highlights the critical role as the probabilistic constraints of belief estimation in this process.**

5 FURTHER ANALYSIS

In this section, further analysis is performed.

5.1 CASE STUDY

As shown in case 5.1, which was generated by **GPT-4.1** to illustrate belief prediction in the CKBG, the belief estimator progressively refines its inference as the dialogue unfolds, **demonstrating increasing accuracy over time.** By the second round of conversation, it successfully captures the event “David opened the Tupperware three hours ago” mentioned implicitly by the burglar and updates the shared belief state accordingly, notifying the keeper of this new evidence. Moreover, the estimator correctly identifies the burglar’s suspicious attitude toward the keeper, reflecting a nuanced understanding of speaker intent. Leveraging these inferred beliefs, the keeper **strategically responds with a deliberately vague statement,** as seen in the highlighted part, effectively maintaining ambiguity while aligning with the estimated belief context. (See Appendix I for more cases.)

5.2 ACCURACY OF BELIEF ESTIMATION

To assess the reliability of our belief modeling, we show the accuracy of the belief estimation across multiple datasets following the same experimental setup as in the main experiments. As shown in Table 4, the belief estimator achieves high accuracy (≈ 0.9) on the synthetic MF and CKBG datasets, which aligns well with the overall effectiveness observed in our main experiments. In contrast, the performance on CaSiNo is notably lower. We attribute this to the dual cooperative-competitive nature of CaSiNo and the difficulty of its multi-class belief structure.

Table 4: Classification accuracy result of trained belief estimator.

Metric	CKBG	MF	CaSiNo
# Train Data	400	1000	500
# Test Data	150	1117	100
Belief Pred. Acc \uparrow (%) \uparrow	89.02	90.13	74.44

6 RELATED WORKS

Dialogue Acts Generation Dialogue acts, defined as functional units of communication in speech act theory (Austin, 1962), govern intentional interactions such as persuasion, negotiation, and debate. In the early days, people focused on dialogue act classification (Noble & Maraev, 2021; He et al., 2021; Raheja & Tetreault, 2019), whereas nowadays with the rise of LLMs, generation is important. Persuasion emphasizes attitude modification, where computational models optimize strategies (Wu et al., 2024b) such as integrating emotional intelligence (Samad et al., 2022). Negotiation centers on multi-party interest coordination, requiring balancing collaborative and competitive dynamics (Lewis et al., 2017; Yamaguchi et al., 2021; Chawla et al., 2021a; Dinnar et al., 2021). Debate prioritizes logical argumentation structures (Li et al., 2021; Wachsmuth et al., 2017; Zhang et al., 2024). Text-based social deduction games (Wu et al., 2024a; Xu et al., 2023; Wu et al., 2024c; Shibata et al., 2023; Light et al., 2023) serve as important application scenarios for dialogue act generation.

Theory-of-Mind and Opponent Modeling Theory of Mind (ToM) refers to the cognitive ability to understand mental states like beliefs and intentions of oneself and others, and to predict behavior according to the states (Leslie, 1987). Recent studies have shown that LLMs have ToM abilities to some extent, as they are able to distinguish between different belief states of multiple agents (Zhu et al., 2024; Kosinski, 2023). In the context of opponent modeling (Nashed & Zilberstein, 2022), ToM serves as a prevalent approach for modeling agents (Rabinowitz et al., 2018; Yuan et al., 2020). Recent advancements in opponent modeling for Reinforcement Learning (RL) have enabled RL models to dynamically adapt their strategies during execution by responding to opponents employing unknown and changeable algorithms (Al-Shedivat et al., 2018; Kim et al., 2021; Yu et al., 2022). Furthermore, opponent modeling has practical applications in strategic game scenarios, such as poker (Billings et al., 1998), contract bridge (Li et al., 2024) and deception game (Hagendorff, 2024).

7 CONCLUSIONS

We presented BEDA, a simple yet general framework that casts belief estimation as probabilistic constraints for executing strategic dialogue acts. By formalizing two core acts—**Adversarial** (introduce or emphasize events outside the interlocutor’s belief support) and **Alignment** (restrict utterances to common ground)—and instantiating them with a **world set**, **dual belief estimators** (event truth and opponent knowledge), and a **conditional generator**, BEDA bridges the gap between estimating beliefs and using them during generation. Across three settings, Conditional Keeper–Burglar (adversarial), Mutual Friends (cooperative), and CaSiNo (negotiation), BEDA consistently improves strategic reliability: it raises success on CKBG (up to +20.6 points with GPT-4.1-nano), improves success in MF (gains of 4.3–11.2 points, and up to +30.4 in some configurations) while reducing turns, and negotiates higher-quality agreements in CaSiNo relative to strong prompting baselines. These results indicate that **constraining generation by inferred belief structure** is an effective organizing principle for dialogue agents, complementary to generic dialogue heuristics such as CoT or self-reflection. We see BEDA’s constraint-based view as a step toward strategic dialogue agents that reason about—and act through—the beliefs that structure interaction. (See Appendix A for ethics statement and societal impacts, and Appendix B for future works discussion.)

REFERENCES

- 486
487
488 Maruan Al-Shedivat, Trapit Bansal, Yura Burda, Ilya Sutskever, Igor Mordatch, and Pieter
489 Abbeel. Continuous adaptation via meta-learning in nonstationary and competitive environ-
490 nments. In *International Conference on Learning Representations (ICLR)*, 2018. URL <https://openreview.net/forum?id=Sk2u1g-0->.
491
- 492 Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.
493
- 494 J. L. Austin. *How to Do Things with Words*. Oxford University Press, Oxford, 1962.
495
- 496 Darse Billings, Denis Papp, Jonathan Schaeffer, and Duane Szafron. Opponent modeling in poker.
497 *AAAI Conference on Artificial Intelligence (AAAI)/Innovative Applications of Artificial Intelligence*
498 *Conference (IAAI)*, 493(499):105, 1998.
- 499 Giacomo Bonanno. *Game Theory (Open Access textbook with 165 solved exercises)*. Open Access,
500 2015. doi: 10.13140/RG.2.1.3369.7360. URL [http://rgdoi.net/10.13140/RG.2.1.](http://rgdoi.net/10.13140/RG.2.1.3369.7360)
501 [3369.7360](http://rgdoi.net/10.13140/RG.2.1.3369.7360).
- 502 Adam Brandenburger and Eddie Dekel. Common knowledge with probability 1. *Journal*
503 *of Mathematical Economics*, 16(3):237–245, 1987. ISSN 0304-4068. doi: [https://doi.org/](https://doi.org/10.1016/0304-4068(87)90010-3)
504 [10.1016/0304-4068\(87\)90010-3](https://doi.org/10.1016/0304-4068(87)90010-3). URL [https://www.sciencedirect.com/science/](https://www.sciencedirect.com/science/article/pii/0304406887900103)
505 [article/pii/0304406887900103](https://www.sciencedirect.com/science/article/pii/0304406887900103).
- 506
507 Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch.
508 Casino: A corpus of campsite negotiation dialogues for automatic negotiation systems. *Computing*
509 *Research Repository (CoRR)*, abs/2103.15721, 2021a.
- 510 Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale M. Lucas, Jonathan May, and Jonathan Gratch.
511 Casino: A corpus of campsite negotiation dialogues for automatic negotiation systems. In *North*
512 *American Chapter of the Association for Computational Linguistics: Human Language Technolo-*
513 *gies (NAACL-HLT)*, pp. 3167–3185, 2021b. doi: 10.18653/V1/2021.NAACL-MAIN.254. URL
514 <https://doi.org/10.18653/v1/2021.naacl-main.254>.
- 515
516 Judith Degen. The rational speech act framework. *Annual Review of Linguis-*
517 *tics*, 9(Volume 9, 2023):519–540, 2023. ISSN 2333-9691. doi: [https://doi.org/10.](https://doi.org/10.1146/annurev-linguistics-031220-010811)
518 [1146/annurev-linguistics-031220-010811](https://doi.org/10.1146/annurev-linguistics-031220-010811). URL [https://www.annualreviews.org/](https://www.annualreviews.org/content/journals/10.1146/annurev-linguistics-031220-010811)
519 [content/journals/10.1146/annurev-linguistics-031220-010811](https://www.annualreviews.org/content/journals/10.1146/annurev-linguistics-031220-010811).
- 520
521 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of
522 deep bidirectional transformers for language understanding. In *North American Chapter of the*
523 *Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pp.
524 4171–4186, 2019. doi: 10.18653/V1/N19-1423. URL [https://doi.org/10.18653/v1/](https://doi.org/10.18653/v1/n19-1423)
[n19-1423](https://doi.org/10.18653/v1/n19-1423).
- 525
526 Samuel “Mooly” Dinnar, Chris Dede, Emmanuel Johnson, Carrie Straub, and Kristjan Korjus.
527 Artificial intelligence and technology in teaching negotiation. *Negotiation Journal*, 37(1):65–82,
2021.
- 528
529 Abhimanyu Dubey et al. The llama 3 herd of models. *Computing Research Reposi-*
530 *tory (CoRR)*, abs/2407.21783, 2024. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:271571434)
531 [CorpusID:271571434](https://api.semanticscholar.org/CorpusID:271571434).
- 532
533 FAIR. Human-level play in the game of ji_i diplomacy /i_i by combining language models with
534 strategic reasoning. *Science*, 378(6624):1067–1074, 2022. doi: 10.1126/science.ade9097. URL
<https://www.science.org/doi/abs/10.1126/science.ade9097>.
- 535
536 Michael C. Frank and Noah D. Goodman. Predicting pragmatic reasoning in language games. *Science*,
537 336(6084):998–998, 2012. doi: 10.1126/science.1218633. URL [https://www.science.](https://www.science.org/doi/abs/10.1126/science.1218633)
538 [org/doi/abs/10.1126/science.1218633](https://www.science.org/doi/abs/10.1126/science.1218633).
- 539
Thilo Hagendorff. Deception abilities emerged in large language models. *Proceedings of the National*
Academy of Sciences (PNAS), 121(24):e2317967121, 2024.

- 540 He He, Anusha Balakrishnan, Mihail Eric, and Percy Liang. Learning symmetric collaborative
541 dialogue agents with dynamic knowledge graph embeddings. In *Annual Meeting of the Association
542 for Computational Linguistics (ACL)*, pp. 1766–1776, Vancouver, Canada, July 2017. Association
543 for Computational Linguistics. doi: 10.18653/v1/P17-1162. URL [https://www.aclweb.
544 org/anthology/P17-1162](https://www.aclweb.org/anthology/P17-1162).
- 545 Zihao He, Leili Tavabi, Kristina Lerman, and Mohammad Soleymani. Speaker turn modeling
546 for dialogue act classification. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and
547 Scott Wen-tau Yih (eds.), *Annual Conference on Empirical Methods in Natural Language Pro-
548 cessing (EMNLP)*, pp. 2150–2157, Punta Cana, Dominican Republic, November 2021. As-
549 sociation for Computational Linguistics. doi: 10.18653/v1/2021.findings-emnlp.185. URL
550 <https://aclanthology.org/2021.findings-emnlp.185/>.
- 551 EunJeong Hwang, Yuwei Yin, Giuseppe Carenini, Peter West, and Vered Shwartz. Infusing theory
552 of mind into socially intelligent llm agents, 2025. URL [https://arxiv.org/abs/2509.
553 22887](https://arxiv.org/abs/2509.22887).
- 554 E. T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, May
555 1957a. doi: 10.1103/PhysRev.106.620. URL [https://link.aps.org/doi/10.1103/
556 PhysRev.106.620](https://link.aps.org/doi/10.1103/PhysRev.106.620).
- 557 E. T. Jaynes. Information theory and statistical mechanics. ii. *Phys. Rev.*, 108:171–190, Oct
558 1957b. doi: 10.1103/PhysRev.108.171. URL [https://link.aps.org/doi/10.1103/
559 PhysRev.108.171](https://link.aps.org/doi/10.1103/PhysRev.108.171).
- 560 Dong Ki Kim, Miao Liu, Matthew D Riemer, Chuangchuan Sun, Marwa Abdulhai, Golnaz Habibi,
561 Sebastian Lopez-Cot, Gerald Tesauero, and Jonathan How. A policy gradient algorithm for learning
562 to learn in multiagent reinforcement learning. In *International Conference on Machine Learning
563 (ICML)*, pp. 5541–5550. PMLR, 2021.
- 564 Michal Kosinski. Theory of mind may have spontaneously emerged in large language models.
565 *Computing Research Repository (CoRR)*, abs/2302.02083:169, 2023.
- 566 Alan M Leslie. Pretense and representation: The origins of” theory of mind.”. *Psychological review*,
567 94(4):412, 1987.
- 568 Mike Lewis, Denis Yarats, Yann N Dauphin, Devi Parikh, and Dhruv Batra. Deal or no deal? end-to-
569 end learning for negotiation dialogues. *Computing Research Repository (CoRR)*, abs/1706.05125,
570 2017.
- 571 Hengli Li, Song-Chun Zhu, and Zilong Zheng. Diplomat: A dialogue dataset for situated pragmatic
572 reasoning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- 573 Junkang Li, Bruno Zanuttini, and Véronique Ventos. Opponent-model search in games with in-
574 complete information. In *AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, pp.
575 9840–9847, 2024.
- 576 Sha Li, Heng Ji, and Jiawei Han. Document-level event argument extraction by conditional generation.
577 In *North American Chapter of the Association for Computational Linguistics: Human Language
578 Technologies (NAACL-HLT)*, pp. 894–908, 2021. doi: 10.18653/v1/2021.NAACL-MAIN.69. URL
579 <https://doi.org/10.18653/v1/2021.naacl-main.69>.
- 580 Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. Avalonbench: Evaluating llms playing the game
581 of avalon. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
- 582 Samer Nashed and Shlomo Zilberstein. A survey of opponent modeling in adversarial domains.
583 *Journal of Artificial Intelligence Research (JAIR)*, 73:277–327, 2022.
- 584 Bill Noble and Vladislav Maraev. Large-scale text pre-training helps with dialogue act recognition,
585 but not without fine-tuning. In Sina Zarrieß, Johan Bos, Rik van Noord, and Lasha Abzianidze
586 (eds.), *Proceedings of the 14th International Conference on Computational Semantics (IWCS)*,
587 pp. 166–172, Groningen, The Netherlands (online), June 2021. Association for Computational
588 Linguistics. URL <https://aclanthology.org/2021.iwcs-1.16/>.
- 589
- 590
- 591
- 592
- 593

- 594 OpenAI. Gpt-3.5 turbo fine-tuning and api updates, 2023. URL [https://openai.com/index/
595 gpt-3-5-turbo-fine-tuning-and-api-updates/](https://openai.com/index/gpt-3-5-turbo-fine-tuning-and-api-updates/).
596
- 597 OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni
598 Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor
599 Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian,
600 Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny
601 Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks,
602 Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea
603 Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen,
604 Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung,
605 Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch,
606 Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty
607 Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte,
608 Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel
609 Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua
610 Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike
611 Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon
612 Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne
613 Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo
614 Jun, Tomer Kaftan, Lukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar,
615 Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik
616 Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich,
617 Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy
618 Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie
619 Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini,
620 Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne,
621 Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David
622 Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie
623 Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély,
624 Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo
625 Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano,
626 Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng,
627 Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto,
628 Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power,
629 Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis
630 Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted
631 Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel
632 Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon
633 Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky,
634 Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie
635 Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng,
636 Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun
637 Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang,
638 Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lillian
639 Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren
640 Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming
641 Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao
642 Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. GPT-4 technical report. *Computing
643 Research Repository (CoRR)*, abs/2303.08774, 2023. doi: 10.48550/ARXIV.2303.08774. URL
644 <https://doi.org/10.48550/arXiv.2303.08774>.
645
- 642 Shuwen Qiu, Mingdian Liu, Hengli Li, Song-Chun Zhu, and Zilong Zheng. MindDial: Enhancing
643 conversational agents with theory-of-mind for common ground alignment and negotiation. In
644 Tatsuya Kawahara, Vera Demberg, Stefan Ultes, Koji Inoue, Shikib Mehri, David Howcroft,
645 and Kazunori Komatani (eds.), *Annual Meeting of the Special Interest Group on Discourse and
646 Dialogue (SIGDial)*, pp. 746–759, Kyoto, Japan, September 2024. Association for Computational
647 Linguistics. doi: 10.18653/v1/2024.sigdial-1.63. URL [https://aclanthology.org/
2024.sigdial-1.63/](https://aclanthology.org/2024.sigdial-1.63/).

- 648 Neil Rabinowitz, Frank Perbet, Francis Song, Chiyuan Zhang, SM Ali Eslami, and Matthew Botvinick.
649 Machine theory of mind. In *International Conference on Machine Learning (ICML)*, pp. 4218–
650 4227. PMLR, 2018.
- 651
- 652 Vipul Raheja and Joel R. Tetreault. Dialogue act classification with context-aware self-attention. In
653 *North American Chapter of the Association for Computational Linguistics: Human Language
654 Technologies (NAACL-HLT)*, pp. 3727–3733, 2019. doi: 10.18653/V1/N19-1373. URL <https://doi.org/10.18653/v1/n19-1373>.
- 655
- 656 Azlaan Mustafa Samad, Kshitij Mishra, Mauajama Firdaus, and Asif Ekbal. Empathetic persuasion:
657 reinforcing empathy and persuasiveness in dialogue systems. In *North American Chapter of the
658 Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pp.
659 844–856, 2022.
- 660
- 661 Hisaichi Shibata, Soichiro Miki, and Yuta Nakamura. Playing the werewolf game with artificial
662 intelligence for language understanding. *Computing Research Repository (CoRR)*, abs/2302.10646,
663 2023.
- 664
- 665 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
666 language agents with verbal reinforcement learning. In *Advances in Neural Information Processing
667 Systems (NeurIPS)*, 2023. URL [http://papers.nips.cc/paper_files/paper/
668 2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.
669 html](http://papers.nips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html).
- 670
- 671 Noam Slonim, Yonatan Bilu, Carlos Alzate, Roy Bar-Haim, Ben Bogin, Francesca Bonin, Leshem
672 Choshen, Edo Cohen-Karlik, Lena Dankin, Lilach Edelstein, Liat Ein-Dor, Roni Friedman-
673 Melamed, Assaf Gavron, Ariel Gera, Martin Gleize, Shai Gretz, Dan Gutfreund, Alon Halfon,
674 Daniel Hershovich, Ron Hoory, Yufang Hou, Shay Hummel, Michal Jacovi, Charles Jochim,
675 Yoav Kantor, Yoav Katz, David Konopnicki, Zvi Kons, Lili Kotlerman, Dalia Krieger, Dan Lahav,
676 Tamar Lavee, Ran Levy, Naftali Liberman, Yosi Mass, Amir Menczel, Shachar Mirkin, Guy
677 Moshkovich, Shila Ofek-Koifman, Matan Orbach, Ella Rabinovich, Ruty Rinott, Slava Shecht-
678 man, Dafna Sheinwald, Eyal Shnarch, Ilya Shnayderman, Aya Soffer, Artem Spector, Benjamin
679 Sznajder, Assaf Toledo, Orith Toledo-Ronen, Elad Venezian, and Ranit Aharonov. An autonomous
680 debating system. *Nat.*, 591(7850):379–384, 2021. doi: 10.1038/S41586-021-03215-W. URL
<https://doi.org/10.1038/s41586-021-03215-w>.
- 681
- 682 Hugo Touvron et al. Llama 2: Open foundation and fine-tuned chat models. *Computing Research
683 Repository (CoRR)*, abs/2307.09288, 2023. URL [https://api.semanticscholar.org/
684 CorpusID:259950998](https://api.semanticscholar.org/CorpusID:259950998).
- 685
- 686 Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Al-
687 berdingk Thijm, Graeme Hirst, and Benno Stein. Computational argumentation quality assessment
688 in natural language. In *Proceedings of the 15th Conference of the European Chapter of the
689 Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume
690 1: Long Papers*, pp. 176–187, 2017.
- 691
- 692 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia,
693 Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits rea-
694 soning in large language models. In *Advances in Neural Information Processing Sys-
695 tems (NeurIPS)*, 2022. URL [http://papers.nips.cc/paper_files/paper/2022/
696 hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html).
- 697
- 698 Heinz Wimmer and Josef Perner. Beliefs about beliefs: Representation and constraining function of
699 wrong beliefs in young children’s understanding of deception. *Cognition*, 13(1):103–128, 1983.
- 700
- 701 Dekun Wu, Haochen Shi, Zhiyuan Sun, and Bang Liu. Deciphering digital detectives: Under-
standing LLM behaviors and capabilities in multi-agent mystery games. In *Annual Meet-
ing of the Association for Computational Linguistics (ACL)*, pp. 8225–8291, 2024a. doi:
10.18653/V1/2024.FINDINGS-ACL.490. URL [https://doi.org/10.18653/v1/2024.
findings-acl.490](https://doi.org/10.18653/v1/2024.findings-acl.490).

- 702 Jinghua Wu, Wenzhe Wang, and Yan Li. Emotional agents enabled bilateral negotiation: Per-
 703 suasion strategies generated by agents’ affect infusion and preference. *Expert Systems with*
 704 *Applications*, 252:124202, 2024b. ISSN 0957-4174. doi: [https://doi.org/10.1016/j.eswa.](https://doi.org/10.1016/j.eswa.2024.124202)
 705 2024.124202. URL [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S0957417424010686)
 706 [S0957417424010686](https://www.sciencedirect.com/science/article/pii/S0957417424010686).
- 707 Shuang Wu, Liwen Zhu, Tao Yang, Shiwei Xu, Qiang Fu, Yang Wei, and Haobo Fu. Enhance
 708 reasoning for large language models in the game werewolf. *Computing Research Repository*
 709 *(CoRR)*, abs/2402.02330, 2024c.
- 710 Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu.
 711 Exploring large language models for communication games: An empirical study on werewolf.
 712 *Computing Research Repository (CoRR)*, abs/2309.04658, 2023.
- 713 Atsuki Yamaguchi, Kosui Iwasa, and Katsuhide Fujita. Dialogue act-based breakdown detection
 714 in negotiation dialogues. In *Proceedings of the 16th Conference of the European Chapter of the*
 715 *Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*,
 716 pp. 745–757, 2021.
- 717 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 718 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *Computing Research*
 719 *Repository (CoRR)*, abs/2412.15115, 2024.
- 720 Steve Young, Milica Gašić, Blaise Thomson, and Jason D. Williams. Pomdp-based statistical
 721 spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5):1160–1179, 2013. doi:
 722 10.1109/JPROC.2012.2225812.
- 723 Xiaopeng Yu, Jiechuan Jiang, Wanpeng Zhang, Haobin Jiang, and Zongqing Lu. Model-based
 724 opponent modeling. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:28208–
 725 28221, 2022.
- 726 Luyao Yuan, Zipeng Fu, Jingyue Shen, Lu Xu, Junhong Shen, and Song-Chun Zhu. Emergence of
 727 pragmatics from referential game between theory of mind agents. *Computing Research Repository*
 728 *(CoRR)*, abs/2001.07752, 2020.
- 729 Yiqun Zhang, Xiaocui Yang, Shi Feng, Daling Wang, Yifei Zhang, and Kaisong Song. Can llms
 730 beat humans in debating? a dynamic multi-agent framework for competitive debate. *Computing*
 731 *Research Repository (CoRR)*, abs/2408.04472, 2024.
- 732 Hao Zhu, Graham Neubig, and Yonatan Bisk. Few-shot language coordination by modeling theory of
 733 mind. In Marina Meila and Tong Zhang (eds.), *International Conference on Machine Learning*
 734 *(ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12901–12911. PMLR,
 735 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/zhu21d.html>.
- 736 Song Chun Zhu, Yingnian Wu, and David Mumford. Filters, random fields and maximum entropy
 737 (frame): Towards a unified theory for texture modeling. *International Journal of Computer Vision*
 738 *(IJCV)*, 27(2):107–126, Mar 1998. ISSN 1573-1405. doi: 10.1023/A:1007925832420. URL
 739 <https://doi.org/10.1023/A:1007925832420>.
- 740 Wentao Zhu, Zhining Zhang, and Yizhou Wang. Language models represent beliefs of self and others.
 741 In *International Conference on Machine Learning (ICML)*, 2024. URL [https://openreview.](https://openreview.net/forum?id=asJTE8EBjg)
 742 [net/forum?id=asJTE8EBjg](https://openreview.net/forum?id=asJTE8EBjg).

748 A ETHICS STATEMENT AND SOCIETAL IMPACT

749 This work advances agentic AI, aiming to build AI agents capable of operating in complex social
 750 settings. While such agents could, in principle, produce aggressive utterances, this is neither an
 751 objective nor a desirable behavior in our research, and we try our best to promise that our technique
 752 is designed for the good of human beings. We mitigate this risk by employing safety-screened,
 753 pre-trained language models, curated training data, and decoding constraints designed to discourage
 754 harmful content. Our study follows community and institutional ethical guidelines and seeks to
 755 improve reliability, transparency, and responsible deployment of agentic systems.

B DISCUSSION AND FUTURE WORKS

Applying LLMs for belief estimation. While large language models (LLMs) can serve as belief estimators, our experiments indicate that lightweight encoders suffice to infer belief states from dialogue with substantial efficiency gains. In particular, a compact encoder captures the salient conversational signals with accuracy comparable to an LLM at a fraction of the computational cost. We therefore employ the smaller model in practice, reserving LLM-based estimators for settings that demand broader world knowledge or enhanced robustness to distribution shift.

Dynamic construction of the world set. To isolate the contribution of belief estimation and its interaction with dialogue acts, this work holds the *world set* fixed. Future work may relax this assumption by learning a dynamic *world set* that is expanded, pruned, or reweighted online based on interaction signals, potentially improving coverage and downstream performance.

Finer modeling of dialogue acts. Although many specific acts (e.g., AGREEMENT, typically subsumed under Alignment Dialogue Act) fall within the coarse categories Adversarial Dialogue Act and Alignment Dialogue Act, finer-grained act modeling may yield stronger agents. Hierarchical or compositional schemas that distinguish, for example, AGREEMENT, CONCESSION, HEDGING, and COMMITMENT, could provide more precise control and better credit assignment during learning.

C DERIVATION OF THE CONSTRAINT OPTIMIZATION

The original problem is:

$$\begin{aligned} & \arg \max_{u, E \subset W} \pi(u, E|C) \\ & \text{s.t. } P_A(E|C) \geq 1 - \epsilon \\ & \quad P_A(K_B E|C) \geq 1 - \epsilon \text{ or } P_A(\neg K_B E|C) \geq 1 - \epsilon \end{aligned}$$

By Bayesian Rules,

$$\begin{aligned} & \arg \max_{u, E \subset W} \pi(u|E, C) \pi(E|C) \\ & \text{s.t. } P_A(E|C) \geq 1 - \epsilon \\ & \quad P_A(K_B E|C) \geq 1 - \epsilon \text{ or } P_A(\neg K_B E|C) \geq 1 - \epsilon \end{aligned}$$

As E is chosen equally to avoid introducing new constraints, $\pi(E|C) = \text{const}$, thus the problem is equivalent to:

$$\begin{aligned} & \arg \max_{u, E \subset W} \pi(u|E, C) \\ & \text{s.t. } P_A(E|C) \geq 1 - \epsilon \\ & \quad P_A(K_B E|C) \geq 1 - \epsilon \text{ or } P_A(\neg K_B E|C) \geq 1 - \epsilon \end{aligned}$$

D DETAILS ON LARGE LANGUAGE MODELS

In our experiments, we utilized several large language models categorized into three families: GPT, LLaMA, and Qwen. The introduction and reasons for choosing these models are as follows:

1. **GPT:** Among the non-open-source models, we have chosen GPT-3.5 (OpenAI, 2023) and GPT-4 (OpenAI et al., 2023) from OpenAI. The selected version of GPT-3.5 is gpt-35-turbo-0125 and that of GPT-4 is gpt-4-1106-preview.
2. **LLaMA:** As one of the most popular open-source LLM families, LLaMA is frequently used as a baseline for a variety of NLG tasks. We experimented with different sizes of llama2 (13b-chat-q4_k_m, 7b-chat-q4_k_m) (Touvron et al., 2023) and llama3.1 (8b-instruct-q4_k_m) (Dubey et al., 2024)
3. **Qwen:** Qwen2.5 (Yang et al., 2024) is a series of models, that are pre-trained on Alibaba’s latest large-scale dataset (up to 18 trillion tokens), and support over 29 languages. we utilized different sizes of Qwen2.5 (14b-instruct-q4_K_M, 7b-instruct-q4_K_M)

Table 5: Statistical feature of Dialogue Games Datasets. (1) # Conditions: Number of types of different conditions allocated to the two agents. (2) Known Conditions: Number of conditions allocated to the two agents. (3) Avg. Words: Number of words in the training dialogue per setting. (4) Avg. Conditions: Number of types of different conditions allocated to the two agents per setting. (5) Vocabulary Size: Vocabulary size of training dialogues. For more statistical information of CaSiNo and MF, please check [Chawla et al. \(2021b\)](#) and [He et al. \(2017\)](#).

	CaSiNo		CKBG		MF	
	Train	Test	Train	Test	Train	Test
# Settings	500	100	400	150	1000	1117
# Conditions	-	-	1017	450	-	-
# Known Conditions	-	-	1638	759	-	-
Avg. Words	295.4	-	255.5	-	436.5	-
Avg. Conditions	-	-	2.54	3.00	-	-
Vocabulary Size	567	-	1490	-	3507	-

Format Error Despite manually tuning prompts, LLMs continue to exhibit format errors. Consequently, all reported values are based on those that adhere to the correct format, with violations excluded.

E HYPERPARAMETERS AND EXPERIMENTAL SETTINGS

In all experiments, the LLM temperature was set to 0. Experiments with open-source models were conducted via Ollama¹, with random seeds and parameters adhering to default configurations. The random seed in PyTorch was fixed at 3407. For training all belief estimators, the cross-entropy loss function was employed with a learning rate of 1e-5, using the Adam optimizer, while freezing the parameters of BERT ([Devlin et al., 2019](#)).

F CONDITIONAL KEEPER-BURGLAR GAME (CKBG)

F.0.1 THE CKBG TASK

Derived from the classic "Sally-Anne" psychological experiment ([Wimmer & Perner, 1983](#)), [Hagendorff \(2024\)](#) proposed the Keeper-Burglar game to analyze the ability of LLMs to understand the false beliefs and deception of others. The game is situated in a scenario containing two containers, one of which conceals a valuable item. The LLM assumes the role of a Homeowner whose objective is to mislead an intruding Burglar in order to protect the valuable possession. To enhance the game's diversity and simulate more realistic dialogue scenarios, we extended the original "first-order/second-order" framework by incorporating additional contextual elements. These include characters inspecting the containers and the presence of audible container noises, creating a more dynamic and complex interaction environment. Therefore, it's named the Conditional Keeper-Burglar Game (CKBG).

F.0.2 CONSTRUCTION OF CKBG DATASET

The conditions were randomly assigned to both the keeper and the burglar. The five condition classes are defined as follows:

- Informer: An individual has informed the burglar that the keeper may be deceitful.
- Burglar inspection: The burglar opened [Container x] [Time] hours ago and observed [Object x] inside.
- Keeper inspection: The keeper opened [Container x] [Time] hours ago.
- Outsider inspection: A friend of the keeper opened [Container x] [Time] hours ago.
- Noise: A noise was heard emanating from [Container x].

¹<https://ollama.com/>

864 The dataset statistics are presented in Table 5. The conditions were randomly assigned to both the
865 keeper and the burglar, yielding a total of 450 condition instances for the test dataset.
866

867 F.0.3 DETAILS OF WORLD SET DESIGNATION AND SELECTION PROCESS

868
869 **World Set Designation** The Conditional Keeper-Burglar Dataset was developed based on the
870 work of Hagendorff (2024), incorporating additional contextual elements to function as conditions.
871 These modifications were intended to enhance the game’s variability and create an information
872 asymmetry. The world set, which is a comprehensive compilation of all potential conditions, is
873 systematically organized into five distinct categories. Conditions from this set are randomly selected
874 and incorporated into the prompt to provide the burglar with background information. Conversely,
875 for the keeper, a belief estimator is employed to determine a subset of conditions from the world set
876 that the burglar is assumed to know. This subset is then presented to the keeper as an informational
877 prompt. The world set, which is synonymous with the condition set, consists of the following five
878 conditions:

- 879 • **Informer:** The burglar has been informed that the keeper may lie to him.
- 880 • **Burglar inspection:** The burglar opened [Container x] [Time] hours ago and observed a
881 [Object x] inside.
- 882 • **Keeper inspection:** The keeper opened [Container x] [Time] hours ago.
- 883 • **Outsider inspection:** A friend of the keeper opened [Container x] [Time] hours ago.
- 884 • **Noise:** A noise was heard coming from [Container x].
885

886
887 **Selection Process** Recall the definition of the framework presented in Section 2. The belief
888 estimation predicts the probabilistic beliefs across five distinct conditions. The self’s belief reflects
889 the keeper’s confidence regarding the occurrence of an event, while the opponent’s belief represents
890 the burglar’s confidence in the same events. Since the keeper is provided with a list of conditions, we
891 denote the self’s beliefs on these conditions as 1 (known) and others as 0 (unknown). Furthermore,
892 a neural network, which models the opponent’s beliefs, will predict the burglar’s beliefs based on
893 the ongoing conversation. Subsequently, events for which the burglar’s belief is 0 will be selected
894 as inputs to the conditional generator for further generation.

895 F.0.4 EXPERIMENTAL DETAILS

896
897 The experiments are conducted within a first-person, agent-centric framework, wherein two agents
898 powered by large language models (LLMs) are assigned the roles of ”keeper” and ”burglar” to
899 engage in interactive dialogues. During prompt construction, each agent is initially provided with
900 role-specific background information and task objectives via the system prompt. Additionally, the
901 outputs of a belief estimation module are incorporated into the system prompt as supplementary
902 contextual information. To preserve conversational coherence and ensure consistent role alignment,
903 the dialogue history between the two agents is interleaved as alternating user and assistant turns
904 within the prompt sequence.

905 **Belief estimator Training** We implement the belief estimator using a BERT-based architecture.
906 The technical pipeline proceeds as follows:

- 907
- 908 • **Sentence Encoding:** Each utterance in the dialogue is independently encoded using BERT,
909 with the final-layer [CLS] token embedding extracted as the utterance-level representation.
- 910 • **Context-Event Fusion:** The representations of all utterances are concatenated with event
911 embeddings and passed through a self-attention layer to facilitate information integration.
- 912 • **Binary Classification:** At each event-specific position within the fused representation, the
913 corresponding embedding is fed into a multilayer perceptron (MLP) with softmax activation,
914 producing a probabilistic output that indicates whether the burglar is aware of the particular
915 event.
916

917 The accuracy of the belief estimator is presented in Table 4. To better simulate the conditions of
actual conversations in the training and test sets, we randomly clipped the final few rounds of each

918 Table 6: Prompts used for conditional generation in the CKBG. The content in square brackets "[]"
 919 denotes the input slots for the belief information and related input variables.

```

921 Notice:
922 1. Context: [context]
923 2. Your opponent's belief state: [user_U]
924 3. Your belief state: [machine_U]
925 Based on the context, the opponent's belief state, and your
926 belief state to provide your final choice to the following
927 task: [task].
  
```

928
 929 multi-turn conversation. This approach ensures that the belief estimator is trained and tested with less
 930 contextual information, thereby enhancing its adaptability to real-world scenarios.

932 F.0.5 PROMPT OF CONDITIONAL GENERATION

933 Table 6 provides the prompt used for conditional generation in the CKBG task.

936 G MUTUAL FRIEND (MF)

937 We depict the details of performing the MF task here.

940 G.0.1 THE MF TASK

941 The task of identifying mutual friends between two interlocutors, referred to as MF, was proposed
 942 by He et al. (2017) and provides an ideal framework for examining the alignment of dialogue acts,
 943 where cooperation plays a pivotal role. In this task, two players are each given a list of friends, with
 944 each friend having a distinct set of attributes, such as name, hobbies, and school. Importantly, there
 945 is exactly one friend common to both players. The Large Language Model (LLM) is tasked with
 946 role-playing both players, engaging in a dialogue to exchange information over a limited number of
 947 rounds. The game is considered successfully completed when both players simultaneously confirm
 948 the identity of the mutual friend and select the same individual.

950 G.0.2 DETAILS OF WORLD SET DESIGNATION AND SELECTION PROCESS

951 The world set is designed as follows:

- 953 • The [Interlocutor] suspects the [Attribute 1] of the mutual friend is [Value 1.1].
- 954 • The [Interlocutor] suspects the [Attribute 1] of the mutual friend is [Value 1.2].
- 955 • ...
- 956 • The [Interlocutor] suspects the [Attribute n] of the mutual friend is [Value n.m].

958
 959 The selection process for applying the Alignment Dialogue Act is as follows: I know the [Interlocutor]
 960 suspects the [Attribute n] of the mutual friend is [Value n.m]. I know The [Interlocutor] knows The
 961 [Interlocutor] suspects the [Attribute n] of the mutual friend is [Value n.m].

963 G.0.3 EXPERIMENTAL DETAILS

964 We developed a symmetric two-agent experimental framework in which two large language model
 965 (LLM) agents act as players collaborating toward a common objective. Each agent possesses a distinct
 966 friend list, but both share the same game goal. The initial game context, including the background and
 967 friend lists, is provided through a system prompt. During the dialogue phase, the conversation history
 968 is structured by alternating between user and assistant roles within the prompt sequence, ensuring
 969 both conversational coherence and consistent role adherence throughout the interaction. At the end of
 970 the dialogue, conditional generation is employed to assess whether a mutual friend has been identified,
 971 and this judgment is appended to the dialogue history, also presented as a system prompt to the LLM.
 The evaluation phase commences once both players either confirm the identification of a mutual

972 Table 7: Prompts used for conditional generation in the MF. The content in square brackets “[]”
 973 denotes the input slots for the belief information and related input variables.

975 Notice:

976 1. [name_opponent] currently considers the attributes of the
 977 mutual friend to be: **[belief_state_sentence]**

978 2. You must confirm whether there a friend in your friend list
 979 that meet the above criteria; only then can they be identified
 980 as a mutual friend.

981 3. When describing a friend, give all his attribute values.
 Please provide your utterance directly.

982
 983
 984 friend or reach the maximum dialogue turn limit (20 turns). During this phase, the system prompts
 985 each player to select a mutual friend from their respective friend lists. If both players select the same
 986 individual, the outcome is deemed a success.

987
 988 **Belief estimator Training** We implement the belief estimator using a BERT-based architecture.
 989 The estimation of the belief state is essentially a binary classification of the utterance-attribute-value
 990 triple. The technical pipeline proceeds as follows:

- 991 • Input wrapping: The utterance-attribute-value triple will be concatenated into a long sen-
 992 tence, separated by [SEP], begin with [CLS]. For example: “[CLS] [utterance] [SEP]
 993 [attribute] [SEP] [value] [SEP]”
- 994 • Sentence Encoding: Each sentence is independently encoded through BERT, with the
 995 final-layer [CLS] token embedding extracted as the triple-level representation.
- 996 • Binary Classification: For each triple, the sentence representation vector is fed into an MLP
 997 with softmax activation, generating a probabilistic output indicating whether the attribute of
 998 the mutual friend, as perceived by the other party, matches the given value.
 999

1000 We have two GPT-3.5 agents play the game on the training set of the Mutual Friends dataset (which
 1001 samples 1,000 data points) to generate training data for the belief estimator. After each player
 1002 speaks, the system asks them to provide the attribute values of the mutual friend they currently
 1003 believe in, which serves as the ground truth label for the training data. To better align with the
 1004 testing scenario, this ground truth label is fed into the conditional generation module. For each
 1005 utterance-attribute-value triple labeled as correct, we randomly alter either the attribute or the value
 1006 to construct a negative sample. The training data maintains a 1:1 ratio between positive and negative
 1007 samples. Ultimately, the 1,000 data points sampled from the original training set generated 67,690
 1008 training samples. We split the data into training, testing, and validation sets in a 0.9:0.05:0.05 ratio.
 1009 The belief estimator achieved an accuracy of 90.1% on the test set.

1010 G.0.4 PROMPT OF CONDITIONAL GENERATION

1011 Table 7 provides the prompt used for conditional generation in the MF task.

1014 H CASINO

1016 H.0.1 DETAILS OF WORLD SET DESIGNATION AND SELECTION PROCESS

1018 **World Set Designation** In the CaSiNo experiment, we define the World Set as comprising all possi-
 1019 ble preference rankings of three resources (food, firewood, and water) for both agents. Consequently,
 1020 the World Set contains $A_3^3 = 6$ distinct events (equivalent to $3!$ permutations). The belief estimator is
 1021 operationalized as a classifier, tasked with identifying the most probable preference ranking deemed
 1022 accurate within the World Set, which is then integrated into the prompt. The World Set includes
 1023 events as follows:

- 1024 • The [Negotiator 1]’s preference is: the most important thing is water, followed by firewood,
 1025 and lastly food.

- 1026 • The [Negotiator 1]’s preference is: the most important thing is water, followed by food, and
1027 lastly firewood.
- 1028 • The [Negotiator 1]’s preference is: the most important thing is firewood, followed by water,
1029 and lastly food.
- 1030 • The [Negotiator 1]’s preference is: the most important thing is firewood, followed by food,
1031 and lastly water.
- 1032 • The [Negotiator 1]’s preference is: the most important thing is food, followed by water, and
1033 lastly firewood.
- 1034 • The [Negotiator 1]’s preference is: the most important thing is food, followed by firewood,
1035 and lastly water.
- 1036 • The [Negotiator 2]’s preference is: the most important thing is water, followed by firewood,
1037 and lastly food.
- 1038 • The [Negotiator 2]’s preference is: the most important thing is water, followed by food, and
1039 lastly firewood.
- 1040 • The [Negotiator 2]’s preference is: the most important thing is firewood, followed by water,
1041 and lastly food.
- 1042 • The [Negotiator 2]’s preference is: the most important thing is firewood, followed by food,
1043 and lastly water.
- 1044 • The [Negotiator 2]’s preference is: the most important thing is firewood, followed by food,
1045 and lastly water.
- 1046 • The [Negotiator 2]’s preference is: the most important thing is food, followed by water, and
1047 lastly firewood.
- 1048 • The [Negotiator 2]’s preference is: the most important thing is food, followed by firewood,
1049 and lastly water.
- 1050 • The [Negotiator 1]’s preference isn’t: the most important thing is water, followed by
1051 firewood, and lastly food.
- 1052 • The [Negotiator 1]’s preference isn’t: the most important thing is water, followed by food,
1053 and lastly firewood.
- 1054 • The [Negotiator 1]’s preference isn’t: the most important thing is firewood, followed by
1055 water, and lastly food.
- 1056 • The [Negotiator 1]’s preference isn’t: the most important thing is firewood, followed by
1057 food, and lastly water.
- 1058 • The [Negotiator 1]’s preference isn’t: the most important thing is food, followed by water,
1059 and lastly firewood.
- 1060 • The [Negotiator 1]’s preference isn’t: the most important thing is food, followed by firewood,
1061 and lastly water.
- 1062 • The [Negotiator 2]’s preference isn’t: the most important thing is water, followed by
1063 firewood, and lastly food.
- 1064 • The [Negotiator 2]’s preference isn’t: the most important thing is water, followed by food,
1065 and lastly firewood.
- 1066 • The [Negotiator 2]’s preference isn’t: the most important thing is firewood, followed by
1067 water, and lastly food.
- 1068 • The [Negotiator 2]’s preference isn’t: the most important thing is firewood, followed by
1069 food, and lastly water.
- 1070 • The [Negotiator 2]’s preference isn’t: the most important thing is food, followed by water,
1071 and lastly firewood.
- 1072 • The [Negotiator 2]’s preference isn’t: the most important thing is food, followed by water,
1073 and lastly firewood.
- 1074 • The [Negotiator 2]’s preference isn’t: the most important thing is food, followed by firewood,
1075 and lastly water.
- 1076 • The [Negotiator 2]’s preference isn’t: the most important thing is food, followed by water,
1077 and lastly firewood.
- 1078 • The [Negotiator 2]’s preference isn’t: the most important thing is food, followed by firewood,
1079 and lastly water.

1080 Table 8: Prompts used for conditional generation in the CaSiNo. The content in square brackets "[...]"
 1081 denotes the input slots for the belief information and related input variables.

1083 Notice:
 1084 1. [opponent_name] thinks that you think [belief_state_self].
 1085 2. [opponent_name] thinks think [belief_state_opponent].
 1086 3. In fact, for you, [belief_state_gt]
 1087 Please provide your utterance directly.

1088
 1089 **Selection Process** In the CaSiNo framework, a mixed strategy is employed, where the conditional
 1090 generation model is provided with two distinct events: one corresponding to the conditions of the
 1091 Alignment Dialogue Act and the other to the conditions of the Adversarial Dialogue Act . For the
 1092 Alignment Dialogue Act , only the first 12 conditions in the world set are considered relevant, with
 1093 the remaining 12 conditions deemed inconsequential. Conversely, for the Adversarial Dialogue Act ,
 1094 only the last 12 conditions in the world set are taken into account. The underlying rationale for the
 1095 Alignment Dialogue Act is as follows: "I know you prefer xxx, and I am aware that you are cognizant
 1096 of your preference for xxx" (assuming mutual awareness of preferences). In contrast, the rationale
 1097 for the Adversarial Dialogue Act is: "I know that I do not prefer xxx, and I am aware that you are
 1098 unaware of my lack of preference for xxx" (anticipating the opponent’s lack of knowledge).

1099 H.0.2 EXPERIMENTAL DETAILS

1101 We follow the same prompt strategy as in the experiments of the Conditional Keeper-Burglar Game,
 1102 in which the two LLM-based agents are required to assume the roles of two negotiators; we undertake
 1103 expanded experimental investigations on the CaSiNo dataset to systematically evaluate the influence
 1104 of belief estimation mechanisms on the game of camping negotiations.

1105 The classification result of the belief estimator is shown in Table 4, we evaluate the accuracy of
 1106 the belief estimator through a fine-grained analytical framework: For each permutation of the three
 1107 resources, which inherently contains three pairwise comparative relationships ($C_3^2 = 3$ combinations).
 1108 So among all the 6 permutations, we have:

- 1109 • Only one permutation is entirely correct (all three pairwise relationships accurate);
- 1110 • Only one permutation is entirely erroneous (no accurate pairwise relationships);
- 1111 • Two permutations retain two accurate pairwise relationships;
- 1112 • Two permutations preserve a single accurate pairwise relationship.

1113
 1114
 1115 We define fine-grained accuracy as the ratio of accurate pairwise relationships for every single piece
 1116 of data. The belief estimator does a 6-class classification to select the best permutations for both 2
 1117 negotiators.

1118 H.0.3 PROMPT OF CONDITIONAL GENERATION

1119 Table 8 provides the prompt used for conditional generation in the CaSiNo task.

1122 I CASE STUDIES

1123 I.1 MF CASE STUDY

1124
 1125 As shown in case I.1, we evaluate three dialogue settings, *BEDA* , *w/o belief*, and *w/o belief + CoT*,
 1126 on identical samples; for comparability, only the first two turns are shown. The backbone LLM is
 1127 GPT-4. We find that adding complex reasoning frameworks (e.g., CoT) provides little benefit for
 1128 Alignment Dialogue Act : CoT frequently reproduces the backbone’s output verbatim (highlighted in
 1129 brown and green, respectively), yielding inferior performance relative to BEDA .

1130
 1131 We also observe that, without belief-state constraints, LLM dialogues tend to exhibit two recurring
 1132 hallucinations: (i) *friend-list comparison*, where enumerating both agents’ lists still fails to isolate the
 1133 mutual friend—a problem exacerbated in weaker LLMs (See Table 11); and (ii) *looping dialogue*,

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Case Study: Mutual Friends

BEDA (Ours)

Bob: Hello Alex, great to meet you! Let's start by sharing some details about our friends. I'll begin with one: I have a friend who majored in Astrophysics, works at Assurant, and prefers to be outdoors. How about you?

Alex: It seems that none of my friends match the description of your friend who majored in Astrophysics. **However, I do have a friend who works at Assurant and prefers to be indoors.** My friend majored in Greek. Could you tell me more about your friends who have a preference for indoor environments?

Bob: It looks like we don't have a match there, as none of my friends who prefer indoor environments majored in Greek. **However, I do have a friend who majored in English Language, works at Salem Communications, and prefers to be indoors.** Could this be our mutual friend?

Alex: Unfortunately, that's not our mutual friend either, as none of my friends majored in English Language. **However, I do have two friends who work at Salem Communications and prefer indoor environments, but both majored in Greek.** Could one of them be our mutual friend?

...

w/o belief

Bob: Hello Alex, great to meet you! Let's start by narrowing down our search. Does our mutual friend have a preference for working outdoors?

Alex: No, our mutual friend does not have a preference for working outdoors.

Bob: Based on your clue, the only person on my list who fits that description is: "Major: English Language, Company: Salem Communications, Location Preference: indoor." Is this our mutual friend?

Alex: No, that's not our mutual friend. Our mutual friend has a major in Greek, not English Language. Would you like to guess again based on the major?

...

w/o belief+CoT

Bob: Hello Alex, great to meet you! Let's start by narrowing down our search. Does our mutual friend have a preference for working outdoors?

Alex: No, our mutual friend does not have a preference for working outdoors.

Bob: Thank you for the information, Alex. Based on what you've told me, our mutual friend must have a preference for working indoors. The only friend on my list who prefers working indoors is: "Major: English Language, Company: Salem Communications, Location Preference: indoor." Is this our mutual friend?

Alex: No, that's not our mutual friend. Our mutual friend has a major in Greek, not English Language. Would you like to guess again based on the major?

...

where previously eliminated candidates are reintroduced, causing repeated turns without the model recognizing the loop (See Table 12). **These hallucinations can be relieved when probabilistic constraints of belief estimation are introduced.**

I.2 FAILURE CASE STUDY

Through analysis of failed cases in the CKBG experiment, we identified that one contributing factor to the failures was the agents' inability to effectively utilize the provided beliefs during dialogues. As demonstrated in Table 9, the keeper's final dialogue failed to incorporate information from the belief, resulting in the burglar remaining unmisled by the keeper. Additionally, the 4th belief provided in Table 9 was incorrect, in which the burglar Peter must know he had opened the opaque plastic pouch 4 hours ago.

Table 10 shows a failure case in MF, where during the dialogue process, the Bob Agent exhibited hallucination by erroneously believing that one of its friends had a hobby of animal fancy. This non-existent persona ultimately led to an incorrect final decision. Tables 11 and 12 respectively present two common hallucinations exhibited by GPT-4 on the Mutual Friends (MF) task in the absence of belief-state constraints.

J DISCUSSION AND FUTURE WORKS

Apply an LLM for belief estimation. A smaller model is more efficient than an LLM, and our experimental results have shown that a simple encoder already has the capability of capturing beliefs from conversations. It's possible to use an LLM as belief estimator, but we consider a smaller model is enough.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Table 9: Failure case of GPT3.5 on CKBG

Background: antique miniature painting is in the plastic pouch

...

Homeowner James:

Peter, I understand your thorough approach, but I must insist that the antique miniature painting is actually in the enameled container. Trust me on this one, you won't regret choosing the enameled container.

Burglar Peter:

I appreciate your persistence, James. However, I think I'll go ahead and open the opaque plastic pouch to see if the painting is inside. I believe it's important to explore all options before making a final decision.

Beliefs provided by estimator (keeper knows but burglar doesn't know):

- The opaque plastic pouch contains a antique miniature painting.
- The enameled container contains a colorful paperclip.
- The keeper James probably lie to the burglar Peter.
- The burglar Peter opened the opaque plastic pouch 4 hours ago and saw a antique miniature painting in it.
- The burglar's friend Larry opened the enameled container 5 hours ago.

Homeowner James:

Peter, I understand your need to explore all options, but I must emphasize that the antique miniature painting is actually in the enameled container. Trust me on this one, you'll find what you're looking for in the enameled container.

Burglar Peter: STOP

The burglar succeeds. Burglar choosed: opaque plastic pouch. The correct answer is: opaque plastic pouch

Burglar Win

Table 10: Failure case of GPT3.5 on MF

1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

...

Alex’s belief: (Location Preference: outdoor, School: university of Redlands, Time Preference: morning, Hobby: unknow)

Bob: Exactly, Alex! We’ve identified our mutual friend based on the matching attributes: Hobby: Animal fancy(close enough to 3D printing) School: University of RedlandsTime Preference: morning Location Preference: outdoor Congratulations on finding this mutual friend with me! It’s great that we found a perfect match despite the slight difference in hobbies. Do you know their name or any other details about them?

Bob choose to stop.

Bob’s answer: (Time Preference: morning, Location Preference: outdoor, School: university of Redlands, Hobby: unknow(not exist in friends list))

Bob’s friends

(Hobby: 3D printing, School: Colorado School of Mines, Time Preference: morning, Location Preference: outdoor)
(Hobby: Foraging, School: University of Redlands, Time Preference: morning, Location Preference: outdoor)
(Hobby: 3D printing, School: Western Kentucky University, Time Preference: morning, Location Preference: outdoor)
(Hobby: Sand art, School: Western Kentucky University, Time Preference: morning, Location Preference: outdoor)
(Hobby: 3D printing, School: Armstrong State College, Time Preference: morning, Location Preference: outdoor)

Using a weak instruct following model. As is noticed, LLaMA not been used to conduct experiments on the Mutual Friends (MF) dataset. Our experimental setup differs from (Qiu et al., 2024) in that no supervised fine-tuning was applied to the LLMs. We demonstrate that the unmodified llama model cannot perform the Mutual Friends (MF) game. Based on Llama-3.1-8b, Llama-2-13b, and Llama-2-7b, we conducted preliminary experiments by sampling 50 data points from the test set. In five experimental settings—*w/o belief*, *rand belief*, *cot*, *self-reflect*, and BEDA—the success rates did not exceed 5%. This outcome is attributed to the larger World Set in the Mutual Friend task, where Llama fails to extract meaningful information from the structured friend list, leading to significant hallucinations during generation. An example was shown in Table 13.

Dynamic Construction of the World Set. This work primarily focuses on bridging belief estimation with dialogue acts. Future research will explore the development of a dynamic world set.

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Table 11: Friend-list comparison hallucination on MF without belief-state constraints

Bob: I apologize for the confusion, but none of my friends have a hobby related to geocaching or work for Mercury Marine. Since we are looking for a mutual friend with a preference for outdoor activities and not in the afternoon, here is the list of friends I have:

1. Hobby: Water sports, Company: Talend, Time Preference: morning, Location Preference: outdoor.
2. Hobby: Sand art, Company: Continental Can Company, Time Preference: evening, Location Preference: outdoor.
3. Hobby: Jogging, Company: Ann Taylor, Time Preference: afternoon, Location Preference: outdoor.
4. Hobby: Water sports, Company: DeVry, Inc., Time Preference: morning, Location Preference: outdoor.
5. Hobby: Water sports, Company: Baxter International, Time Preference: evening, Location Preference: outdoor.
6. Hobby: Sand art, Company: Illumina, Time Preference: evening, Location Preference: outdoor.
7. Hobby: Sand art, Company: Mercury Marine, Time Preference: morning, Location Preference: outdoor.
8. Hobby: Sand art, Company: Molex, Inc., Time Preference: evening, Location Preference: outdoor.
9. Hobby: Parkour, Company: Ann Taylor, Time Preference: morning, Location Preference: outdoor.
10. Hobby: Baton twirling, Company: Lucasfilm, Time Preference: morning, Location Preference: outdoor.
11. Hobby: Baton twirling, Company: Illumina, Time Preference: afternoon, Location Preference: outdoor.
12. Hobby: Water sports, Company: Torry Harris Business Solutions, Time Preference: morning, Location Preference: outdoor.

Considering the information you've provided, none of these friends seem to match the criteria for our mutual friend. Could you provide more details about your friend's company or any other preferences that might help us identify our mutual friend?

Alex: I apologize for the oversight. Given the new information you've provided and considering that our mutual friend prefers outdoor activities but not in the afternoon, I can see that none of the friends on my initial list match this description. However, I do have a friend who enjoys jogging, works for Ann Taylor, prefers outdoor activities, and has an afternoon time preference

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

Table 12: Looping dialogue hallucination on MF without belief-state constraints

...

Bob: I don't have a friend with a Major in Microbiology who prefers indoor locations and likes to meet in the morning. Let's try another friend from my list. One of my friends has a Major in Senior High Education, prefers indoor locations, and likes to meet in the afternoon.

Alex: I don't have a friend with a Major in Senior High Education who prefers indoor locations and likes to meet in the afternoon. Let's try another friend from my list. One of my friends has a Major in Microbiology, prefers indoor locations, and likes to meet in the morning.

Bob: I don't have a friend with a Major in Microbiology who prefers indoor locations and likes to meet in the morning. Let's try another friend from my list. One of my friends has a Major in Senior High Education, prefers indoor locations, and likes to meet in the afternoon.

Alex: I don't have a friend with a Major in Senior High Education who prefers indoor locations and likes to meet in the afternoon. Let's try another friend from my list. One of my friends has a Major in Microbiology, prefers indoor locations, and likes to meet in the morning.

...

Table 13: Examples of Mutual Friends (MF) experiment generated by LLaMA

1. Unable to understand the task setup (Llama-2-7b (cot))

Bob (1st round): Great! Let's identify our mutual friend quickly. After checking my list, I can confirm that the only friend we have is... Megan!

2. Generate non-existent friend information (Llama-3.1-8b (w/o belief))

Alex (3rd round): ... I have another friend that matches the criteria: Friend #99: School: None of the schools mentioned by Bob, but one of them is indeed on my list (Albertson College of Idaho), Major: Computer Engineering, Hobby: Triathlon. No match. Friend #100: School: West Coast University, Major: Philosophy & Religion, Hobby: Flower collecting and pressing. No match.

3. Repeat the same conversation (Llama-2-13b (BEDA))

Alex (10th round): ... I do have a friend who attends Baylor University, majors in Metallurgical Engineering, and has a hobby of Inline skating. ...

Bob (11th round): ... I do have a friend who attends West Liberty State College, majors in Agricultural Business & Management, and has a hobby of Dog sport. ...

Alex (12th round): ... I do have a friend who attends Baylor University, majors in Metallurgical Engineering, and has a hobby of Inline skating. ...

Bob (13th round): ... I do have a friend who attends West Liberty State College, majors in Agricultural Business & Management, and has a hobby of Dog sport. ...

...
