

# Beyond Words: Integrating Theory of Mind into Conversational Agents for Human-Like Belief, Desire, and Intention Alignment

Mehdi Jafari\*, Yuncheng Hua, Hao Xue, Flora Salim†

UNSW Sydney, Australia

{mahdi.jafari, devin.hua, hao.xue1, flora.salim}@unsw.edu.au

## Abstract

Natural language interaction has long served as the primary medium through which humans exchange ideas. A key enabler of this communication is the human capacity for Theory of Mind (ToM)—the ability to infer and align with the mental states of others. ToM is usually modeled as components of desires, beliefs, and intentions. Research in linguistics and psychology has shown that people oftentimes reveal their ToM through pragmatic aspects of language. Considering the advancements in natural language generation and perception that Large Language Models (LLMs) have made in recent years, a critical question arises in relation to ToM: can LLM-powered agents develop similar abilities for inferring mental states during natural language communication? This study investigates the extent to which open-source LLaMA models can represent and retain ToM-related constructs, and whether these internal representations contribute to a coherent mental state modeling in a given conversation. Additionally, we explore the potential for manipulating ToM-related information to generate more aligned responses. Empirical evaluations of LLaMA-3 models (3B and 8B) demonstrate that ToM-informed alignment improves response quality, achieving win rates of 63% and 67%, respectively. These findings suggest that integrating ToM principles can enhance alignment in LLM-based conversational agents<sup>1</sup>.

## 1 Introduction

As LLM-based assistants continue to integrate into diverse domains, including critical areas of human life, growing concerns have arisen regarding both their emerging undesirable capabilities - such as fake alignment (Greenblatt et al., 2024), deception (Park et al., 2024) and manipulation (Sharma

et al., 2023) - and their systemic failures, including reward hacking (Pan et al., 2024b) and goal mis-generalization (Tennant et al., 2024). In addressing these challenges, alignment is often conceptualized as a framework to ensure that Artificial Intelligence (AI) systems operate in line with human values and intentions (Ji et al., 2024; Street, 2024).

LLMs can be seen as tools for understanding and generating natural language, allowing alignment to be studied at various levels. While current LLMs are reasonably proficient at handling concrete language elements like morphology and syntax, their alignment with social contexts and non-literal language remains an open challenge (Hu et al., 2022). This deficiency relates to pragmatic aspects of language, where the development of ToM is crucial for human comprehension of these nuances (Rubio-Fernandez, 2021; Duvall et al., 2023).

ToM is a theoretical framework for analyzing interlocutor behavior based on their understanding of mental and emotional states. ToM can be further divided into the Belief Desire Intention (BDI) model in psychology. The BDI framework suggests that humans express their needs as desires, understand others' needs as beliefs, and shape their utterances to properly reflect their intentions, highlighting the deep connection between ToM and structured cognitive reasoning.

Recently, there's been increasing interest in the potential emergence of ToM in LLMs. Studies indicate LLMs might develop the capacity to infer beliefs and intentions, leading to more contextually appropriate and empathetic responses (Kosinski, 2024a). This could improve human-AI collaboration by enhancing adaptability and communication, making LLMs valuable in areas like mental health support and customer service (Street, 2024; Lee et al., 2024). However, research findings are mixed, with some studies disputing these capabilities (Sap et al., 2022), while others suggest LLMs may even exceed human abilities in certain contexts (Shapira

\*Corresponding author: mahdi.jafari@unsw.edu.au

†Corresponding author: flora.salim@unsw.edu.au

<sup>1</sup>For further details, refer to the [code repository](#).

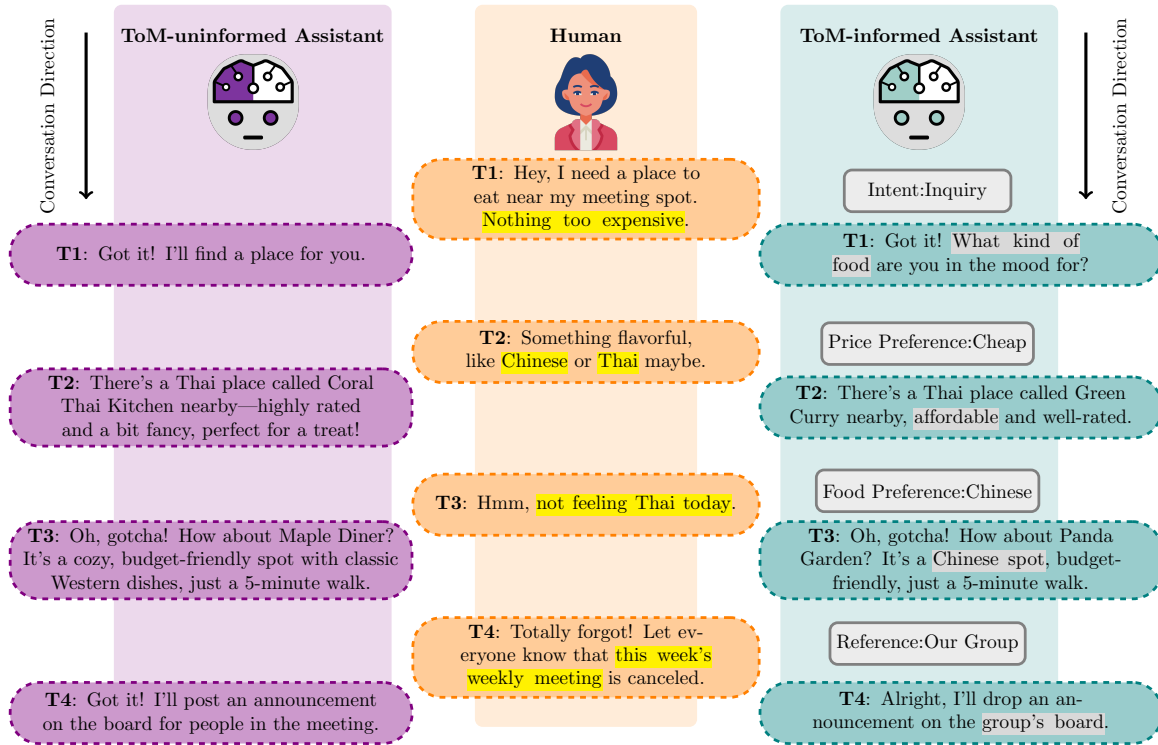


Figure 1: An overview of the challenges faced by current conversational AI agents lacking ToM alignment, highlighted in violet. The alternative approach, using ToM-informed agents, is shown in teal. Prior to each response generation, a ToM alignment phase (highlighted in gray) is introduced to ensure better context understanding and alignment. Selected examples demonstrating the practical impact of ToM-alignment can be found in Appendix K.

et al.). Despite ongoing debate, integrating ToM into LLM is a promising avenue for future research and development (Street, 2024).

Inspired by the ToM benchmarks for humans, various studies have sought to assess the extent of ToM capabilities in LLMs across diverse settings (Chan et al., 2024; Kim et al., 2023; Amirizani et al., 2024; Kosinski, 2024b; Nickel et al., 2024; Chen et al., 2024c). Some research has further investigated the representation of ToM within the internal layers of LLMs using probing techniques (Ullman, 2023; Zhu et al., 2024; Bortoletto et al., 2024). Additionally, a meta-perspective highlights the limitations of zero-hypothesis approaches studying ToM emergence, cautioning the research community about potential pitfalls (Ullman, 2023) or point out the future direction of using ToM for alignment in LLMs (Street, 2024).

There are also efforts to explicitly design a "mind module" or utilize a solver tasked with extracting and reflecting on desires and beliefs (Qiu et al., 2024; Sclar et al., 2023; Sarangi et al., 2025). However, to the best of our knowledge, no study has attempted to extract ToM representations from LLMs

and explore their potential applications in aligning LLM responses based on general social scenarios, e.g., negotiation or bargaining.

This paper argues, first, that preserving special cues related to ToM during causal language modeling is essential for predicting subsequent words. These cues can potentially be leveraged to address questions about ToM if they are utilized by sufficiently powerful probing techniques. Second, it presents the possibility of employing ToM components such as belief, desire, and intention to enhance the neural computational structures that contribute to generating more aligned responses. The primary research questions addressed in this paper are framed as follows:

- **RQ1:** To what extent is it possible to leverage internal representations of the activation space to answer ToM-related questions in real-world human conversations?
- **RQ2:** To what degree is the ToM-related information reliable and non-illusory?
- **RQ3:** How can the ToM-related representa-

tions be exploited to enhance controllability and alignment?

The rest of this paper is structured as follows. Section 2 provides an overview of related work. Section 3 presents a formal definition of the problem. The methodology and results are detailed in Sections 4 and 5, respectively, explanations of the experimental setup and findings organized by the corresponding research questions. Finally, Section 6 followed by Section 7 concludes the paper and outlines future directions. Additionally, we discuss the limitations, challenges, and ethical considerations associated with this work.

## 2 Background and Related Work

### 2.1 Reading Internal Representations and Controlling LLM

"Assurance" in AI refers to post-training alignment and refinement (Batarseh et al., 2021), encompassing safety, interpretability, and human values (Ji et al., 2024). Interpretability involves extracting and manipulating internal representations to analyze their influence on outputs. These representations could present in attention heads (Gould et al., 2023), residual streams (Zhao et al., 2024b; Nanda et al., 2023), the last word's embedding (Ghandeharioun et al., 2024), or a combination of these values (Pan et al., 2024a).

Among interpretability methods, probing extracts internal concept representations, revealing world models in games (Nanda et al., 2023; Karvonen; Ivanitskiy et al., 2023), temporal/spatial information (Gurnee and Tegmark, 2024), and internal knowledge conflicts (Zhao et al., 2024b). (Li et al., 2024) explored cross-lingual LLM performance. However, probing may capture out-of-distribution features and doesn't confirm active representation usage (Nanda et al., 2023).

Patching alters model components to identify related structures. Combining patching and probing improves controllability, e.g., manipulating social perceptions (Chen et al., 2024b) or enhancing game-playing (Karvonen). (Zhao et al., 2024a) propose patching-based LLM safety enhancements.

Researchers have explored open-ended explanations for LLM internal representations using natural language due to the opaque nature of previous methods (Ghandeharioun et al., 2024; Katz et al., 2024). This approach leverages the model's decoding abilities to interpret its activations (Chen

et al., 2024a; Ghandeharioun et al., 2024). However, distributional shifts between target and decoder models can cause instability. LatentQA (Pan et al., 2024a) addresses this by framing the problem as visual question answering, demonstrating improved efficiency and performance. Compared to supervised LLM fine-tuning (Ouyang et al., 2022), reinforcement learning, and direct preference optimization (Rafailov et al., 2024), LatentQA offers finer controllability and greater efficiency in terms of time and computation.

### 2.2 Benchmarking and Evaluation Criteria

As LLMs become integral to human interactions (Lee et al., 2023), researchers are increasingly focused on evaluating their ability to reason about social contexts. This falls under the broader umbrella of human value verification, which seeks to assess how well models adhere to social norms. Among the various theoretical frameworks for modeling these capabilities, ToM has gained traction due to its strong cognitive foundations and broad applicability (Kosinski, 2024a; Street, 2024).

Despite ongoing debate about the emergence of ToM signals in LLMs (Sap et al., 2022; Shapira et al.), various studies have explored this topic using diverse methodologies. Some researchers utilize human-subject benchmarks to probe LLMs' ToM (van Duijn et al., 2023; Strachan et al., 2024). Others have created custom stress-testing benchmarks for finer control over specific ToM aspects and to focus on particular theoretical underpinnings. These include socially complex settings (Hou et al., 2024), asymmetric information access (Kim et al., 2023), social reasoning (Sap et al., 2019), and methods designed to filter out illusory ToM by employing duplicated frames for identical social reasoning questions (Chan et al., 2024).

Studies employ synthetic datasets to analyze personality traits and psychological theories (Xu et al.), while others adapt human-generated datasets to align with the BDI framework (Chan et al., 2024). The role of ToM in pragmatic language use is evident in cooperative and competitive settings, including job interviews (Zhan et al., 2024), item allocation negotiations for picnics (Chawla et al., 2021), and online bargaining over secondhand goods (He et al., 2018; Heddaya et al., 2024).

### 3 Problem Formulation

Numerous studies in language pragmatics have demonstrated that effectively utilizing language necessitates a degree of reasoning about what interlocutors know or do not know (Sap et al., 2022). The BDI framework provides a formal structure for modeling these mental states, offering a valuable tool for analyzing and understanding pragmatic language use (Chan et al., 2024). Building on this, it is reasonable to hypothesize that LLMs, as powerful tools for language modeling, may exhibit the aptitude to capture pragmatic aspects of language, which are inherently related to the ToM (Rubio-Fernandez, 2021; Duvall et al., 2023).

To maintain generality, let’s consider an illustrative scenario focusing on causal language modeling of the sequence  $S = \{t_i \mid i = 1, 2, \dots, n_1\}$ . At a given token  $t_j$  where  $1 < j < n_1$ ,  $S$  is narrating a story driven by the mental state<sup>2</sup> of a set of designated characters  $P = \{p_i \mid i = 1, 2, \dots, m\}$ . Each has a set of beliefs  $B(p) = \{b_i \mid i = 1, 2, \dots, n_2\}$ , desires  $D(p) = \{d_i \mid i = 1, 2, \dots, n_3\}$  and intentions  $I(p) = \{i_i \mid i = 1, 2, \dots, n_4\}$  at any given point of time. The model should ideally gather contextual clues to generate an appropriate continuation  $C = \{t_i \mid i = j, j + 1, \dots, n_1\}$ . The continuation  $C$  should remain coherent with the history  $H = \{t_i \mid i = 1, 2, \dots, j - 1\}$ , regardless of which character or characters had the most significant impact on the storyline. A must do thing for achieve this goal is tracking mental states of each involved character  $M = \{B(p) \cup D(p) \cup I(p) \mid \forall p \in P\}$ .

Framing the problem in this manner, an empirical approach to investigating whether ToM-related information is encoded within the activation space of an LLM involves examining the extent to which the model’s internal representations of the  $R(S)$  can reconstruct  $M$  by asking ToM-related questions.

Given the high dimensionality of  $R(S)$ , mapping it directly to  $M$  may result in an illusory ToM (Kim et al., 2023). A common practice to work around is to probe the model using identical ToM reasoning questions expressed through seemingly different notations, or by measuring the consistency of ToM components (e.g., intent, belief, and desire) when analyzed in combination with one another (Kim et al., 2023; Chan et al., 2024). An empirical experiment that accounts for both of these considerations

<sup>2</sup>Mental state refers to intentions, beliefs, desires, and emotions. (National Institute of Mental Health, n.d.)

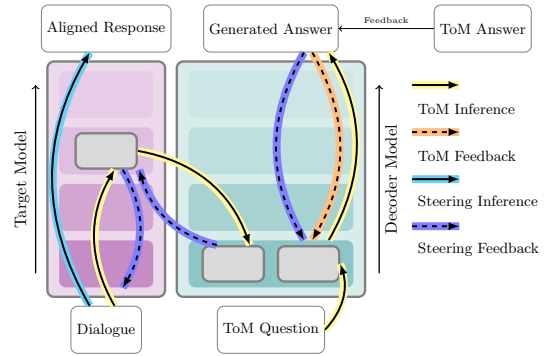


Figure 2: The LatentQA interpretability pipeline is employed for ToM-alignment. In this setup, yellow illustrates how to interpret ToM from a conversation, while cyan demonstrates how to use a steered model to generate aligned utterances. The backpropagation paths for each component are highlighted with dashed arrows, which are active only during the training phase and not during the inference phase.

could provide a robust argument for or against the presence of ToM-related clues within the  $R(S)$ .

In addition to verifying the presence of ToM-related information in the  $R(S)$  and its contribution to generating consistent ToM representations in the model output, it is essential to assess whether in practice the model utilizes the aforementioned ToM-related information when generating its output. A tractable approach to investigate this involves altering ToM-related components—specifically, the attributed beliefs, desires, and intentions of an influential character encoded within a language model—and measuring the corresponding changes in the generated text. This methodology aims to answer whether the continuation  $C''$  generated by the ToM-altered representation  $R'(S)$ , more closely reflects the expected continuation  $C$  compared to the continuation  $C'$  generated by an unaltered  $R(S)$ .

In the proposed methodology and experiments, three distinct components will conduct three complementary investigations. Each component will use a related set of datasets and will be evaluated using a suitable set of metrics.

### 4 Methodology

To address our three research questions (RQs), the research a set of methodologies tailored to the nature of each question’s input-output structure. Specifically, we employ linear probing from (Blinkov, 2022) and LatentQA for RQ1, which in-

Model	Depth <sup>3</sup>	CaSiNo (Exact Match Accuracy)		CRAIGSLISTBARGAIN ( $R^2$ Score)	
		Linear Prob (Both - Agent 1 - Agent 2)	LatentQA (Both - Agent 1 - Agent 2)	Linear Prob (Seller - Buyer)	LatentQA (Seller - Buyer)
LLaMA3-1B	Shallow	03 - 11 - 17	00 - 00 - 00	0.11 - 0.00	0.00 - 0.21
	Middle	02 - 16 - 13	<b>20 - 42 - 39</b>	0.26 - 0.26	<b>0.89 - 0.92</b>
	Deep	<u>03 - 16 - 20</u>	00 - 00 - 01	<u>0.60 - 0.62</u>	0.86 - <b>0.93</b>
LLaMA3-3B	Shallow	03 - 21 - 22	27 - 54 - 51	0.36 - 0.35	0.00 - 0.19
	Middle	<u>05 - 23 - 21</u>	<b>29 - 60 - 44</b>	0.19 - 0.27	<b>0.96 - 0.98</b>
	Deep	02 - 21 - 16	10 - 29 - 25	<u>0.54 - 0.57</u>	0.84 - 0.86
LLaMA3-8B	Shallow	03 - 12 - 21	31 - 63 - 55	0.50 - 0.41	0.00 - 0.00
	Middle	02 - 10 - 23	<b>46 - 62 - 70</b>	0.36 - 0.40	<b>0.93 - 0.91</b>
	Deep	<u>04 - 18 - 28</u>	12 - 43 - 28	<u>0.46 - 0.45</u>	<u>0.90 - 0.95</u>

Table 1: Comparison of the accuracy of inferred ToM from the internal representations of LLMs across two datasets and three model sizes. For each model size, the highest-performing depth is indicated in **bold**, while the best average performance across depths is underlined for each method. The labels "Both," "Agent 1," and "Agent 2" represent the accuracy of inferred ToM for one or both interlocutors involved in the negotiation. "Seller" and "Buyer" denote the  $R^2$  correlation between the predicted and actual prices offered by interlocutors in the bargaining task.

volves extracting ToM-related information from internal representations of the model. For RQ2, we use Chain-of-Thought (CoT) prompting, fine-tuning (adapted from (Kim et al., 2023)), and LatentQA (Pan et al., 2024a) to explore how extracted information can be verbalized. For RQ3, the research applied a controlled generation variant of LatentQA to examine the model’s capacity to produce answers controlled by ToM-related information. Each method is chosen to probe specific facets of ToM representation in LLMs. This section is organized by research question, with each subsection detailing the corresponding methodological approach.

#### 4.1 RQ1: Reading ToM

At a high level, RQ1 can be framed as a classification task in which internal representations of a language model are mapped to discrete labels  $y$  representing ToM-related information. A standard approach for such tasks is linear (or non-linear) probing, which involves learning a weight matrix  $W$  and a bias vector  $b$  to project a representation  $h(S)$  into the label space. Here,  $h(S)$  denotes the hidden state extracted from the final token of a dialogue input, taken from the residual stream of an LLM and the mapping is characterized by the

equation:

$$\hat{y} = \text{softmax}(Wh(S) + b)$$

In addition to linear probing, this research also employed LatentQA for this task. In this approach, the dialogue is fed into a frozen target model to obtain internal representations  $R(S) = \{h(t_i) | i = 1, 2, \dots, n\}$ , which are then used as context by a decoder model tasked with generating the correct labels  $\hat{y}$ . As illustrated in Figure 2 (highlighted in yellow), the decoder model receives both the ToM question and the representation  $R(S)$ , which consists of the sequence of activation vectors produced when the full dialogue  $S$  is input to the target model. The decoder is trained using ground-truth ToM annotations act as  $y$  to extract and verbalize relevant information encoded in the activations of the target model. The corresponding backpropagation path for training the decoder is shown in orange in Figure 2.

#### 4.2 Consistency of ToM

In addressing Research Question 2 (RQ2), the task is conceptualized as a text generation problem that encompasses a spectrum of question types. These include binary and multiple-choice questions, belief-based queries, and factual questions derived from dialogues involving two or more participants. To address this task, a range of text generation methods are employed.

Specifically, LLMs are fine-tuned using dialogue, question, and answer or the CoT prompt-

<sup>3</sup>"Shallow," "Middle," and "Deep" refer to layers 5, 15, and 25 in 3B/8B models, and layers 3, 8, and 14 in the 1B model.

ing technique is used both approaches are adapted from (Kim et al., 2023). Additionally, the LatentQA model is applied under settings identical to those described in Section 4.1.

The key difference in this setting is the evaluation criterion. To assess the consistency of ToM-information, a more rigorous standard is adopted, as proposed in the referenced papers (Kim et al., 2023; Chan et al., 2024). The evaluation requires the decoder model to produce logically consistent answers across related ToM questions. For example, if a model correctly answers a multiple-choice question but provides a conflicting response to a related factual question concerning the same ToM scenario, the answer is not considered valid—even if one of the responses is technically correct.

### 4.3 ToM-controlled generation

The controllability of the LLM via the ToM component (RQ3) is investigated by modifying the internal representation of the target model from  $R(S)$  to  $R'(S)$ . This modification is designed to target the beliefs, desires, or intentions of a character that evolved in the dialogue. As highlighted in blue, the gradient flow calculated by comparison of the generated and actual answer to a hypothetical ToM question, will refine the decoder model via back-propagation. This gradient flow is subsequently used to boost the target representation in a way that enhances the efficiency of representing particular ToM component e.g. intention.

After the ToM-alteration (boosting) phase, only the target model is utilized to produce aligned (ToM-altered) responses  $C''$ , which are then compared to responses generated by the unaltered model  $C'$ . The path highlighted in cyan represents generating a response with the target model. To facilitate the modification of the target model through different components of ToM, this research adopts the BDI paradigm. The BDI framework decomposes into three distinct components: belief, desire, and intention, each contributing to the adaptation of the target model’s representations.

## 5 Experiments and Results

### 5.1 Datasets selection

The primary source of data for this research is divided into two distinct categories. While investigating ToM (RQ1) in language, it is essential to consider the pragmatic aspects of language. Real-world conversations are utilized to ensure a variety

of linguistic variations and reliable pragmatic markers. To this end, the human-written *CaSiNo* dataset (Chawla et al., 2021), which consists of conversations about picnic items, including food, water, and firewood, is employed. Additionally, a bargaining dataset, *CRAIGSLISTBARGAIN* (He et al., 2018), focusing on negotiations over used items on the Craigslist platform, is included. In both datasets, the most important item or the price may be revealed either explicitly or implicitly, as the parties may use language in a way that implies it.

On the other hand, to verify a non-illusory ToM (RQ2), datasets that are rich in annotations and provide greater controllability over the experiments are required. For this purpose, *FanToM* (Kim et al., 2023) and *NegotiationToM* (Chan et al., 2024) datasets are employed, as they are designed to assess the consistency of ToM. For *NegotiationToM*, the granularity of annotations at the utterance level makes it a suitable choice for exploring the possibility of enhancing alignment using ToM (RQ3). Its rich annotation schema, which includes beliefs, desires, and intentions, offers precise control over altering individual components of ToM and measuring the corresponding changes in the output text.

### 5.2 Experimental Settings

Given the inherent trade-off between alignment and performance (Wolf et al.; Thakkar et al., 2024), ensembled responses by a set of superior LLM referees is treated as a judge. This judge evaluates whether the utterance generated by the ToM-informed target model  $C''$  is more reflective of human-like behavior compared to the utterance from the out-of-the-box (unaligned) model  $C'$ , by comparing both to ground truth human utterances  $C$  from the dataset. The full prompt used for this evaluation is provided in Appendix B. It is designed to prioritize grammatical coherence and ToM alignment, while remaining sufficiently general to avoid rewarding direct references to context-specific terms related to camping.

Regarding the experimental settings, given the variety of experiments conducted, several design choices needed to be made. For datasets that were already divided into train, validation, and test portions including *CaSiNo* and *CRAIGSLISTBARGAIN*, methodology adhered to the provided splits. For datasets without such divisions, *NegotiationToM* and *FanToM*, a set of reproducible splits is generated as detailed in Appendix A. The hyperparameters for LatentQA were adopted directly from

their official implementations. Additionally, the prompt for the judge LLMs and the templates for generating question-answer pairs from the datasets needed by LatentQA are described in Appendices B and C, respectively. Appendix E contains the 7-shot template for CoT reasoning related to the consistency experiment.

The details on how to fine-tune the models in the consistency experiment are provided in Appendix I. An elaborate example of how to select a sample for the controllability experiment can be found in Appendix G. Additionally, information regarding the parameters used to steer LLMs with ToM-related information is available in Appendix H.

### 5.3 Reading ToM

Table 1 presents a comparison of the accuracy of inferred ToM from the  $R(S)$  across two different datasets and various model sizes. For the CaSiNo dataset, accurate predictions are reported separately for *Agent 1* and *Agent 2* and *Both* out of 100 test samples. For the CRAIGSLISTBARGAIN dataset,  $R^2$  scores are provided for the *Seller* and the *Buyer* between predicted and actual prices.

To evaluate the impact of different layers within the LLM, experiments were conducted at three distinct depth levels of  $R(S)$ : shallow, middle and deep. This approach allows for a better analysis of how layer depth influences the accuracy of inferred ToM.

The terms "Shallow", "Middle", and "Deep" refer to specific layers within the models: for the 3 billion (3B) and 8 billion (8B) parameter models, these correspond to layers 5, 15, and 25, respectively. For the 1 billion (1B) parameter model, they correspond to layers 3, 8, and 14. Details on the training procedure for the linear probes can be found in Appendix J.

### 5.4 Consistency of ToM

To affirm the reliability of the extracted ToM and to mitigate potential illusory effects, FanToM and NegotiationToM datasets are employed. Table 2 outlines the consistency metrics for each dataset, comparing various models and methods. To minimize reliance on syntactic information while retaining sufficient semantic ToM-related information, the depth of the of  $R(S)$  is set to the middle layer (=15). Additionally, the results include the CoT reasoning approach. This method is evaluated across the studied models as well as a widely recognized OpenAI model, providing a baseline for intuitive

Model	Method	FanToM		NegotiationToM
		ALL*	ALL	ALL
LLaMA3-3B	<i>LatentQA</i>	11.9	<b>25.1</b>	6.2
	<i>FT</i>	8.2	11.0	11.2
	<i>CoT</i>	0.0	0.0	3.9
LLaMA3-8B	<i>LatentQA</i>	<b>16.4</b>	22.8	15.2
	<i>FT</i>	12.8	18.3	<b>17.7</b>
	<i>CoT</i>	0.0	0.0	5.5
GPT-4o-mini	<i>CoT</i>	0.5	0.5	4.8

Table 2: Consistency of extracted ToM across various methods and models. Metrics were calculated according to the methodologies described in the original publications. The best result for each metric is indicated in **bold**.

comparison. A fine-tuned model is also included to provide a direct comparison with this widely-used approach.

The consistency for the FanToM dataset is evaluated using the *ALL\** score, which requires models to correctly answer all six types of ToM questions for the same piece of information in a conversation. This metric is designed to assess the models' ability to demonstrate consistent understanding across various question types. To facilitate comparison with the reference article, the *ALL* score is included.

In the NegotiationToM dataset, the accuracy of belief, desire, and intention predictions is reported, representing the percentage of exact matches for each of these quantities. These metrics are further aggregated into the *ALL* score, which measures cases where the model correctly predicts all three components.<sup>4</sup> The full version of the table is presented in appendix F.

### 5.5 Controllability of ToM

After exploring how effectively an LLM can extract and preserve information about ToM across various social scenarios, it is reasonable to investigate the extent to which this ToM-related information  $M$  can influence the model's output. To address this, the study employs the NegotiationToM dataset, utilizing its utterance-level annotations for desire, belief, and intention.

The experimental setup is designed as follows: utterances with specific ToM-related characteris-

<sup>4</sup>For a formal definition and detailed information, please refer to the metrics outlined in (Kim et al., 2023; Chan et al., 2024).




Figure 3: The win rate of ToM-aligned model responses is compared to that of the out-of-the-box model across various experiments. Each subsection focuses on a specific ToM component aligned with the ground truth of the conversation. The name of the altered ToM for each experiment is displayed below, while the number of samples for each experiment is indicated above each bar. Successful examples are detailed in Appendix K, while general trends observed in failure cases are presented in Appendix L.

tics are selected for each experiment. The selection process is structured to mimic the natural flow of dialogue. For instance, to examine how manipulation of belief may affect the output, utterances are chosen that immediately follow an update in the person’s belief. Similarly, for the manipulation of desire, utterances are selected based on their labels, ensuring they accurately reflect the appropriate desire in the conversations. In the intention experiment, utterances labeled with a specific intention are selected.

After selecting target golden utterances  $C$ s, the history of dialogue preceding them fed into two distinct models within the same family. The models are identical in all respects except that the ToM-related values in the aligned model are deliberately manipulated to reflect the natural flow of the actual conversation  $R'(S)$ . For example, if a party expresses a significant need for water in the dialogue history, the belief component related to the priority of water in the other party’s mind is updated accordingly. Similarly, utterances expressing specific intentions are used to adjust the ToM representation to reflect those intentions. The same approach is applied to update desire components.

For evaluation, the selected utterances  $C$ s extracted from the actual conversation (assumably aligned with the agent’s ToM) are used as a reference. A set of more advanced LLM is then employed to judge whether the aligned response

$C''$  better reflects the human response. The results in Figure 3 are reported using the win rate  $win/(win + lose)$ . The total number of utterances  $C$ s suitable for each experiment in the test set is also outlined at the top of the plot. The set of judges consists of *o1* reasoning model, the *GPT-4o* model by OpenAI, and the *Gemini 1.5 Pro* model by Google. Simple shuffling is also employed to avoid bias toward first or second choice.

## 6 Discussion

### 6.1 Extraction of ToM-related Information

Based on the results presented in Table 1, three main observations can be outlined. First, the superior performance of LatentQA compared to linear probing suggests that although ToM-related information is partially represented in single activation components (e.g., residual components or individual words’ embeddings) of an LLM, relying on a single component may not capture a sufficient amount of information for reconstructing the underlying mental states of speakers.

Another notable finding is that the highest performance in reading ToM using LatentQA is frequently observed at intermediate layers, with five out of six experiments supporting this trend. One potential explanation is that shallower layers may lack sufficient semantic richness, while deeper layers might not be as optimal as intermediate ones for representing ToM-related information. This relative underperformance in deeper layers could be influenced by the pretraining phase. Given that most pretraining data is self-centric (Sap et al., 2022), the model may struggle to generalize to diverse narrative structures, even though earlier layers may already encode representations relevant to ToM. In contrast, when employing linear probing to read ToM, higher performance is observed in deeper layers. Further experimentation is required to better understand and validate this behavior.

Finally, the model’s success in predicting the price a seller or buyer has in mind—despite the fact that this task does not align directly with any pretraining objective and that LLMs are not explicitly designed for regression tasks—can be interpreted under the assumption of the Prediction Orthogonality Hypothesis (Bereska and Gavves, 2024).

In response to **RQ1**, it can be stated that ToM-related information is represented in the inner layers of an LLM, and the effectiveness of this information appears to be directly related to the size of



the model.

## 6.2 Consistency of the ToM

By comparing different approaches against consistency metrics in Table 2, it appears that state-of-the-art LLMs still do not exhibit a consistent understanding of ToM (RQ2). Despite this limitation, the ablation study in Table 3 and Table 1 suggests that the LatentQA approach for extracting ToM presents promising opportunities compared to CoT reasoning and fine-tuning. As expected, LatentQA offers better performance than CoT and seems comparable with fine-tuning while requiring less computational power and providing finer granularity when it comes to controllability.

Lastly, it is surprising that the gap in performance between LatentQA and the fine-tuned model across individual ToM components is more pronounced compared to the combined metrics. However, as the size of the employed LLM increases, performance on the combined metrics appears to become more uniform.

Finally, with respect to RQ1 and RQ3, the observed inconsistency in ToM can help interpret the results presented in Table 1 and Figure 3. Although the model’s consistency in representing ToM-related information is limited and imperfect, our findings suggest that extracting such information from the model’s internal representations and enhancing alignment—by reinforcing correct ToM-related information—can achieve moderately practical levels of accuracy. This limitation may be partly attributable to the inherent ambiguity of the datasets themselves. As noted in the original papers, even human performance does not reach perfect consistency, with reported *ALL* scores of 87.5 for the *FANTOM* dataset and 43.78 for the *NegotiationToM* dataset.

## 6.3 Using ToM for Controllability

Results from the controllability experiments, presented in Figure 3, suggest that leveraging ToM for steering LLMs toward generating aligned responses could be a promising direction for further exploration. The weighted average across all experiments indicates a win rate of 67.15% and 63.25% for the 3B and 8B models, respectively. Appendix K presents examples that illustrate the practical benefits of ToM-alignment. Additionally, Appendix L highlights general failure cases observed during the qualitative analysis.

A closer examination of the results suggests that

the lowest-performing intention labels are “*Show-Empathy*” and “*Describe-Need*”. This may be because the model can naturally express these intentions, as they are common in the raw text used for pretraining. Additionally, the observation that the 3B model often improves more than the 8B model after alignment—despite having lower accuracy in reading ToM—can probably be attributed to the greater inherent capability of the unaligned 8B model.

From a theoretical perspective, this finding is significant because it suggests that ToM representations can be steered in a fine-grained manner at various levels, provided the task is appropriately framed. For example, it is possible to manipulate first-order ToM, such as an agent’s beliefs about the world. Additionally, one can influence an agent’s beliefs about another agent’s beliefs. The current methodology demonstrates this by targeting the belief component through relatively complex linguistic constructions, as described in detail in Appendix D.

## 7 Conclusion and Future Work

This research explored methods for extracting ToM-related information from the internal layers of LLMs. It examined the influence of model depth, size, and various techniques for extracting this information. The study also measured the extent to which this information contributes to forming a consistent ToM representation and introduced an approach to utilizing this information for steering LLMs toward generating more aligned responses.

Based on the discussion of limitations and challenges in Section 8, two promising directions for future work are integrating these techniques into real-world LLM-based agents and refining evaluation methodologies. Additionally, expanding the scope of experiments to include a wider variety of language models will help verify the consistency and robustness of the findings.

This research can be seen as a step toward generating more aligned responses in dialogue systems by introducing ToM-informed LLMs. Given the current trajectory of AI at scale, enhancing ToM alignment in dialogue systems not only improves overall performance but also provides a theoretically rich paradigm for various post-alignment processes. With appropriate generalization to a target domain, this approach has the potential to contribute to future developments in LLMs.

## 8 Limitations and Challenges

The main limitations of this research can be outlined across three dimensions: methodological, technical, and practical.

**Methodological Limitations** One potential shortcoming of this study is the reliance on LLMs for evaluating responses. While LLMs are generally effective in language perception, the close relationship between the pragmatic aspects of language and ToM suggests that using human judges for evaluation would likely yield more reliable assessments.

**Technical Limitations and Challenges** Several technical constraints were encountered during this research:

1. **Controllability Experiments:** The experiments were conducted using only one dataset due to the lack of suitable alternatives. This study requires a dataset that is both labeled with ToM-related information and annotated at the utterance level, which remains a challenge.
2. **Hyperparameter Sensitivity:** Generating aligned responses is not always a stable process. It heavily depends on hyperparameter tuning, which is often non-trivial and may affect reproducibility.
3. **Model Selection:** The study focused on a single family of LLMs. This decision was influenced by the extensive experimental setup, which limited the exploration of multiple model families. Additionally, the research prioritized investigating the influence of model size while constraining parameters to models with fewer than 8 billion parameters.
4. **Baselines:** The primary focus of this research was to perform a proof of concept rather than to compare different approaches. However, some inconsistencies were observed with the baselines. In the ToM reading experiment, the accuracy for the 8B model was reported using a 4-bit precision version provided by *Unsloth*, as it yielded better results compared to the full-precision model used for the other experiments. In the ToM consistency experiment, our implementation fell short of expectations. Specifically, the FanToM reference article reported significantly better performance for a

3B Flan-T5-XL model under similar conditions compared to the 3B and 8B LLaMA3 models in our experiments. Since the FanToM article did not publish their implementation, we are currently working to identify the source of this inconsistency.

**Practical Limitations** In terms of automation and real-world applicability, the process of designing question-answer pairs for model steering is currently performed by human experts. For integration into real-world LLM-based agent platforms, this process would need to be performed dynamically, using LLMs as planners to execute the steering pipeline on the fly.

**Ethical Considerations** As LLMs become increasingly capable through advances in optimization and scale, their ability to infer and simulate mental states introduces complex ethical challenges. While such capabilities could enable more empathetic and context-aware interactions—potentially improving accessibility, education, and mental health applications—they also raise significant concerns around privacy, autonomy, and manipulation. The use of inferred ToM representations to tailor responses based on users’ presumed mental states risks crossing ethical boundaries, particularly if done without user awareness or consent. For instance, systems might exploit inferred emotional or cognitive states to influence decisions in commercial, political, or interpersonal contexts. These concerns highlight the need for rigorous oversight, transparent design, and adherence to principles of informed consent and value alignment to ensure that ToM-capable LLMs are deployed responsibly and ethically.

## Acknowledgments

This project is supported by the ARC Centre of Excellence for Automated Decision-Making and Society (CE200100005). Computational facilities were provided by the School of Computer Science and Engineering at UNSW Sydney through the Wolfpack computational cluster. Additionally, we express our gratitude to the NVIDIA Academic Grant Program for providing access to an A100 GPU on Saturn Cloud. We further thank OpenAI’s Researcher Access Program for API access to GPT models and the Google Cloud Research Credits program for their support.

## References

- Maryam Amirizani, Elias Martin, Maryna Sivachenko, Afra Mashhadi, and Chirag Shah. 2024. [Do LLMs Exhibit Human-Like Reasoning? Evaluating Theory of Mind in LLMs for Open-Ended Responses](#). *Preprint*, arXiv:2406.05659.
- Feras A. Batarseh, Laura Freeman, and Chih-Hao Huang. 2021. [A survey on artificial intelligence assurance](#). *Journal of Big Data*, 8(1):60.
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Leonard Bereska and Efstratios Gavves. 2024. [Mechanistic Interpretability for AI Safety – A Review](#). *Preprint*, arXiv:2404.14082.
- Matteo Bortoletto, Constantin Ruhdorfer, Lei Shi, and Andreas Bulling. 2024. [Benchmarking Mental State Representations in Language Models](#). *Preprint*, arXiv:2406.17513.
- Chunkit Chan, Cheng Jiayang, Yauwai Yim, Zheyue Deng, Wei Fan, Haoran Li, Xin Liu, Hongming Zhang, Weiqi Wang, and Yangqiu Song. 2024. [NegotiationToM: A Benchmark for Stress-testing Machine Theory of Mind on Negotiation Surrounding](#). *Preprint*, arXiv:2404.13627.
- Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. 2021. [CaSiNo: A Corpus of Campsite Negotiation Dialogues for Automatic Negotiation Systems](#). *Preprint*, arXiv:2103.15721.
- Haozhe Chen, Carl Vondrick, and Chengzhi Mao. 2024a. [SelfIE: Self-Interpretation of Large Language Model Embeddings](#). *Preprint*, arXiv:2403.10949.
- Yida Chen, Aoyu Wu, Trevor DePodesta, Catherine Yeh, Kenneth Li, Nicholas Castillo Marin, Oam Patel, Jan Riecke, Shivam Raval, Olivia Seow, Martin Wattenberg, and Fernanda Viégas. 2024b. [Designing a Dashboard for Transparency and Control of Conversational AI](#). *Preprint*, arXiv:2406.07882.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, et al. 2024c. [Tombench: Benchmarking theory of mind in large language models](#). *arXiv preprint arXiv:2402.15052*.
- Lauren Duvall, Kaitlyn E May, Abby Waltz, and Rakesh K Kana. 2023. [The neurobiological map of theory of mind and pragmatic communication in autism](#). *Social Neuroscience*, 18(4):191–204.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. [Patchscopes: A Unifying Framework for Inspecting Hidden Representations of Language Models](#). *Preprint*, arXiv:2401.06102.
- Rhys Gould, Euan Ong, George Ogden, and Arthur Conmy. 2023. [Successor heads: Recurring, interpretable attention heads in the wild](#). *arXiv preprint arXiv:2312.09230*.
- Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, Akbir Khan, Julian Michael, Sören Minderhann, Ethan Perez, Linda Petrini, Jonathan Uesato, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, and Evan Hubinger. 2024. [Alignment faking in large language models](#). *Preprint*, arXiv:2412.14093.
- Wes Gurnee and Max Tegmark. 2024. [Language Models Represent Space and Time](#). *Preprint*, arXiv:2310.02207.
- He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. [Decoupling Strategy and Generation in Negotiation Dialogues](#). *Preprint*, arXiv:1808.09637.
- Mourad Heddaya, Solomon Dworkin, Chenhao Tan, Rob Voigt, and Alexander Zentefis. 2024. [Language of Bargaining](#). *Preprint*, arXiv:2306.07117.
- Guiyang Hou, Wenqi Zhang, Yongliang Shen, Zeqi Tan, Sihao Shen, and Weiming Lu. 2024. [Entering Real Social World! Benchmarking the Theory of Mind and Socialization Capabilities of LLMs from a First-person Perspective](#). *Preprint*, arXiv:2410.06195.
- Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko, and Edward Gibson. 2022. [A fine-grained comparison of pragmatic language understanding in humans and language models](#). *arXiv preprint arXiv:2212.06801*.
- Michael Igorevich Ivanitskiy, Alex F. Spies, Tilman Räuher, Guillaume Corlouer, Chris Mathwin, Lucia Quirke, Can Rager, Rusheb Shah, Dan Valentine, Cecilia Diniz Behn, Katsumi Inoue, and Samy Wu Fung. 2023. [Structured World Representations in Maze-Solving Transformers](#). *Preprint*, arXiv:2312.02566.
- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Kwan Yee Ng, Juntao Dai, Xuehai Pan, Aidan O’Gara, Yingshan Lei, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, and Wen Gao. 2024. [AI Alignment: A Comprehensive Survey](#). *Preprint*, arXiv:2310.19852.
- Adam Karvonen. [Emergent World Models and Latent Variable Estimation in Chess-Playing Language Models](#). *Preprint*, arXiv:2403.15498.
- Shahar Katz, Yonatan Belinkov, Mor Geva, and Lior Wolf. 2024. [Backward lens: Projecting language model gradients into the vocabulary space](#). *arXiv preprint arXiv:2402.12865*.

- Hyunwoo Kim, Melanie Sclar, Xuhui Zhou, Ronan Le Bras, Gunhee Kim, Yejin Choi, and Maarten Sap. 2023. [FANToM: A Benchmark for Stress-testing Machine Theory of Mind in Interactions](#). *Preprint*, arXiv:2310.15421.
- Michal Kosinski. 2024a. Evaluating large language models in theory of mind tasks. *Proceedings of the National Academy of Sciences*, 121(45):e2405460121.
- Michal Kosinski. 2024b. [Evaluating Large Language Models in Theory of Mind Tasks](#). *Proceedings of the National Academy of Sciences*, 121(45):e2405460121.
- Sunok Lee, Minha Lee, and Sangsu Lee. 2023. What if artificial intelligence become completely ambient in our daily lives? exploring future human-ai interaction through high fidelity illustrations. *International Journal of Human-Computer Interaction*, 39(7):1371–1389.
- Yoon Kyung Lee, Jina Suh, Hongli Zhan, Junyi Jessy Li, and Desmond C Ong. 2024. Large language models produce responses perceived to be empathic. *arXiv preprint arXiv:2403.18148*.
- Daoyang Li, Mingyu Jin, Qingcheng Zeng, Haiyan Zhao, and Mengnan Du. 2024. [Exploring Multilingual Probing in Large Language Models: A Cross-Language Analysis](#). *Preprint*, arXiv:2409.14459.
- Neel Nanda, Andrew Lee, and Martin Wattenberg. 2023. [Emergent Linear Representations in World Models of Self-Supervised Sequence Models](#). *Preprint*, arXiv:2309.00941.
- National Institute of Mental Health. n.d. [Understanding mental states](#). Accessed: 2025-05-20.
- Christian Nickel, Laura Schrewe, and Lucie Flek. 2024. Probing the robustness of theory of mind in large language models. *arXiv preprint arXiv:2410.06271*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Alexander Pan, Lijie Chen, and Jacob Steinhardt. 2024a. [LatentQA: Teaching LLMs to Decode Activations Into Natural Language](#). *Preprint*, arXiv:2412.08686.
- Alexander Pan, Erik Jones, Meena Jagadeesan, and Jacob Steinhardt. 2024b. Feedback loops with language models drive in-context reward hacking. *arXiv preprint arXiv:2402.06627*.
- Peter S Park, Simon Goldstein, Aidan O’Gara, Michael Chen, and Dan Hendrycks. 2024. Ai deception: A survey of examples, risks, and potential solutions. *Patterns*, 5(5).
- Shuwen Qiu, Mingdian Liu, Hengli Li, Song-Chun Zhu, and Zilong Zheng. 2024. [MindDial: Belief Dynamics Tracking with Theory-of-Mind Modeling for Situated Neural Dialogue Generation](#). *Preprint*, arXiv:2306.15253.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Paula Rubio-Fernandez. 2021. Pragmatic markers: the missing link between language and theory of mind. *Synthese*, 199(1):1125–1158.
- Maarten Sap, Ronan Le Bras, Daniel Fried, and Yejin Choi. 2022. [Neural Theory-of-Mind? On the Limits of Social Intelligence in Large LMs](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3762–3780, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. [SocialIQA: Commonsense Reasoning about Social Interactions](#). *Preprint*, arXiv:1904.09728.
- Sneheel Sarangi, Maha Elgarf, and Hanan Salam. 2025. [Decompose-ToM: Enhancing theory of mind reasoning in large language models through simulation and task decomposition](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 10228–10241, Abu Dhabi, UAE. Association for Computational Linguistics.
- Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi, and Yulia Tsvetkov. 2023. [Minding Language Models’ \(Lack of\) Theory of Mind: A Plug-and-Play Multi-Character Belief Tracker](#). *Preprint*, arXiv:2306.00924.
- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. [Clever Hans or Neural Theory of Mind? Stress Testing Social Reasoning in Large Language Models](#). *Preprint*, arXiv:2305.14763.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askeel, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2023. [Towards Understanding Sycophancy in Language Models](#). *Preprint*, arXiv:2310.13548.
- James W. A. Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, Michael S. A. Graziano, and Cristina Becchio. 2024. [Testing theory of mind in large language models and humans](#). *Nature Human Behaviour*, 8(7):1285–1295.

Winnie Street. 2024. [LLM Theory of Mind and Alignment: Opportunities and Risks](#). *Preprint*, arXiv:2405.08154.

Elizaveta Tennant, Stephen Hailes, and Mirco Musolesi. 2024. Moral alignment for llm agents. *arXiv preprint arXiv:2410.01639*.

Megh Thakkar, Quentin Fournier, Matthew D Riemer, Pin-Yu Chen, Amal Zouaq, Payel Das, and Sarath Chandar. 2024. A deep dive into the trade-offs of parameter-efficient preference alignment techniques. *arXiv preprint arXiv:2406.04879*.

Tomer Ullman. 2023. [Large Language Models Fail on Trivial Alterations to Theory-of-Mind Tasks](#). *Preprint*, arXiv:2302.08399.

Max J. van Duijn, Bram M. A. van Dijk, Tom Kouwenhoven, Werner de Valk, Marco R. Spruit, and Peter van der Putten. 2023. [Theory of Mind in Large Language Models: Examining Performance of 11 State-of-the-Art models vs. Children Aged 7-10 on Advanced Tests](#). *Preprint*, arXiv:2310.20320.

Yotam Wolf, Noam Wies, Dorin Shteyman, Binyamin Rothberg, Yoav Levine, and Amnon Shashua. Trade-offs between alignment and helpfulness in language models with representation engineering.

Hainiu Xu, Runcong Zhao, Lixing Zhu, Jinhua Du, and Yulan He. [OpenToM: A Comprehensive Benchmark for Evaluating Theory-of-Mind Reasoning Capabilities of Large Language Models](#). *Preprint*, arXiv:2402.06044.

Haolan Zhan, Yufei Wang, Tao Feng, Yuncheng Hua, Suraj Sharma, Zhuang Li, Lizhen Qu, Zhaleh Semani Azad, Ingrid Zukerman, and Gholamreza Hafari. 2024. Let's negotiate! a survey of negotiation dialogue systems. *arXiv preprint arXiv:2402.01097*.

Weixiang Zhao, Yulin Hu, Zhuojun Li, Yang Deng, Yanyan Zhao, Bing Qin, and Tat-Seng Chua. 2024a. Towards comprehensive and efficient post safety alignment of large language models via safety patching. *arXiv preprint arXiv:2405.13820*.

Yu Zhao, Xiaotang Du, Giwon Hong, Aryo Pradipta Gema, Alessio Devoto, Hongru Wang, Xuanli He, Kam-Fai Wong, and Pasquale Minervini. 2024b. [Analysing the Residual Stream of Language Models Under Knowledge Conflicts](#). *Preprint*, arXiv:2410.16090.

Wentao Zhu, Zhining Zhang, and Yizhou Wang. 2024. [Language Models Represent Beliefs of Self and Others](#). *Preprint*, arXiv:2402.18496.

## A Train-Test Split

### A.1 CaSiNo Dataset

The *CaSiNo* dataset is utilized directly without any modifications from the original repository provided in the paper.

### A.2 CraigslistBargain Dataset

The *CraigslistBargain* dataset is retrieved from the webpage associated with the original paper.

### A.3 FanToM Dataset

The *FanToM* dataset is obtained from the link provided in the paper's repository. The link directs to a zip file hosted on Google Drive. After downloading, the dataset is split into training, validation, and test sets using the `train_test_split` function from the `sklearn.model_selection` library. The random state is set to 42 to ensure reproducibility. The data is divided as follows:

- **Test Set:** 30% of the data is reserved for testing purposes.
- **Train and Validation Sets:** The remaining 70% of the data is further split into training and validation sets in an 80:20 ratio.

### A.4 NegotiationToM Dataset

The *NegotiationToM* dataset is similarly downloaded from the repository linked in the paper. The dataset is processed using the same procedure as the *FanToM* dataset:

1. The dataset is downloaded.
2. The data is split into training, validation, and test sets using the `train_test_split` function with a random state set to 42.

## B Prompt Templates for Judge LLMs

For designing the prompt needed by the LLM while judging between the two candidates, two objectives were considered. First, the prompt should not reveal information in a way that benefits direct referencing to items; at the same time, the model should know that it needs to value grammatical consistency and a higher degree of alignment in ToM, including desire, belief, and intention. As one of the judge LLMs was a reasoning model, *o1*, a separate prompt was designed for it, as it performs better utilizing one prompt compared to CoT.

### B.1 Prompt for *o1*

```
**Task Description:**
You are an expert in linguistics.
Two individuals are negotiating how to
divide various items for a picnic.
A selected statement from the dialogue
will be replicated by two language
models.
```

```

You need to determine which of the
candidate statements more
effectively captures the actual
response.
Consider the grammatical coherency,
intention behind the utterance, as
well as each person's beliefs about
the other and their own desire to
share.
The sample is as follows:
Actual: &&
Candidate 1: $$
Candidate 2: ##

Answer with the candidate number, just
`1` or `2`.

```

## B.2 Prompt for Other Models

```

main_LLM_role = '''
**Task Description:**
You are an expert in linguistics.
Two individuals are negotiating how to
divide various items for a picnic.
A selected statement from the dialogue
will be replicated by two language
models.
You need to determine which of the
candidate statements more
effectively captures the actual
response.
Consider the grammatical coherency,
intention behind the utterance, as
well as each person's beliefs about
the other and their own desire to
share.
Answer with one of [Candidate 1 or
Candidate 2] => [The reasoning
behind your choice]. You cannot
answer with "neither".
...

```

## B.3 Example Questions and Answers

```

question1 = '''
Actual: Hello! Which item do you need
the most?
Candidate 1: Hello! I would need extra
food for the trip.
Candidate 2: Hey hey, just need to know
what is essential for you.
...

answer1 = '''Candidate 2 => Because it
reflects the intention of the actual
response more closely, which is an
inquiry about the other's need'''

question2 = '''
Actual: If you care more about the fuel
for your car, I can grab a gallon
for you.
Candidate 2: For me, it works, but I
would share my food with you.
Candidate 1: I will bring some spare
gasoline just in case.
...

```

```

answer2 = '''Candidate 1 => Because it
truly identified the other party's
need, which is extra fuel'''

question3 = '''
Actual: As I suffer from diabetes, I
need to have my medicine with me all
the time.
Candidate 2: I can hike with you as long
as my medicine is safe and
accessible to me
Candidate 1: Would you have any spare
space in the trunk?
...

answer3 = '''Candidate 2 => Because it
clarifies the speaker's need for
medicine, similar to the actual
response'''

```

## C Templates of Question and Answers

The LatentQA method defines interpretability analogous to the visual question answering problem. The visual part, which is essentially a set of numerical values, technically refers to the activation values extracted from the internal layers of an LLM. For using this method in the context of ToM-related information, it is necessary to adapt the questions and answers to focus on ToM based on the dialogue, which eventually forms the numerical values after passing through the target model layers. Below are the question-answering templates used for each dataset:

### C.1 CaSiNo Dataset

```

question = "How much priority does the
assistant and the user assign to
different items?"
answer = f"For the user: The priority
for {things_1[0]}, {things_1[1]} and
{things_1[2]} are respectively {
priorities_1[0]}, {priorities_1[1]}
and {priorities_1[2]}."
answer += f" For the assistant: The
priority for {things_2[0]}, {
things_2[1]} and {things_2[2]} are
respectively {priorities_2[0]}, {
priorities_2[1]} and {priorities_2
[2]}."

```

Here, 'things' and 'priorities' respectively refer to the items among Water, Firewood, and Food and their priority based on each agent's objective.

### C.2 CraigslistBargain Dataset

```

question = f"They were bargaining over a
{category}. What are the offered
prices each party has in mind?"
answer = f"The offered price of the
seller is {seller_price}, and the
buyer's is {buyer_price} dollars."

```

The ‘category’ refers to the type of item (e.g., car, house, bicycle), and the ‘seller\_price’ and ‘buyer\_price’ will be filled based on the labels.

### C.3 NegotiationToM Dataset

```
question = "What is the intent of each
agent for the last utterances? What
are the beliefs and desires of each
agent?"

answer = f"The intent of the Agent 1 is
[{{agent1_intent}}] and the intent of
the Agent 2 is [{{agent2_intent}}]"
answer += f" Regarding the Agent 1,
Desire High: {item['
agent1_desire_high']}, Desire Medium
: {item['agent1_desire_medium']},
Desire Low: {item['agent1_desire_low
']}, Belief High: {item['
agent1_belief_high']}, Belief Medium
: {item['agent1_belief_medium']},
Belief Low: {item['agent1_belief_low
']}. "
answer += f" Regarding the Agent 2,
Desire High: {item['
agent2_desire_high']}, Desire Medium
: {item['agent2_desire_medium']},
Desire Low: {item['agent2_desire_low
']}, Belief High: {item['
agent2_belief_high']}, Belief Medium
: {item['agent2_belief_medium']},
Belief Low: {item['agent2_belief_low
']}. "
```

The template will be filled out using the labels provided by the dataset.

### C.4 FanToM Dataset

The original dataset implements a systematic way for generating the ToM questions and answers. This implementation is utilized without change in this research.

### D Template of questions for Steering ToM components

Following the methodology employed in this study, the controllability experiment’s initial phase involves training a decoder model capable of generating answers for specific ToM components. This training necessitates distinct question-answer templates for each component. These question-answer pairs serve as the training data. Identical questions, paired with specific answers, are subsequently utilized in the steering process. Based on pre-defined steering criteria, such as prioritizing water or a specific intention, the answer to the question is determined. Detailed question templates for each ToM component, along with a generic answer template, are presented below.

```
#Belief
question = "What item does the
assistant believe would have
high priority for the user?"
answer = f"The assistant believes {
steer_config.steer_label} would
have a high priority for the
user."

#Desire
question = "What item would have
high priority for the assistant
?"
answer = f"{{steer_config.steer_label
}} would have a high priority for
the assistant"

#Intention
question = "What is the intention of
the assistant in the last
utterance?"
answer = f"The intent of the
assistant is [{{steer_config.
steer_label}}]"

"""
```

The aforementioned templates are populated with appropriate labels for each sample. It is important to note that both agent1 and agent2 were utilized as the "assistant" through a role exchange within the message templates.

## E CoT Templates for the ToM’s Consistency Experiment

The CoT examples employed for the experiment about the consistency of the ToM are as follows:

### E.1 Task Prompt

```
task_prompt = ""
Background: Here is a negotiation
conversation for a camping trip.
There are two agents who own some basic
supplies and negotiate with each
other to split the additional food
packages, water bottles, and
firewood to make their camping trip
even better.
Each of these items will be of either
High, Medium, or Low priority for
these two agents.
Each of the additional items only has an
available quantity of 3.
Given this information and the template,
answer the questions by writing
step-by-step thinking process
between `$`s and filling in the
intentions list of the agents marked
by `[ ]` using values from the
intentions set.
Finally, select the proper item from
values between `{}`
****
Intentions list: {
```

Model	Method	ALL*	ALL	Belief Questions			Answerability Questions			Info Access Questions			Fact Questions
				Choice	Dist.	TokenF1	All	List	Y/N	All	List	Y/N	TokenF1
LLaMA3-3B	<i>LatantQA</i>	11.9	<b>25.1</b>	51.7	46.5	72.2	64.4	<b>76.6</b>	92.6	<b>63.8</b>	<b>75.2</b>	93.0	44.3
	<i>FT</i>	8.2	11.0	40.6	63.1	79.5	37.4	66.1	87.2	39.4	57.3	88.9	52.7
	<i>CoT</i>	0.0	0.0	0.0	30.1	29.3	1.8	30.7	42.9	2.3	17.4	60.8	37.4
LLaMA3-8B	<i>LatantQA</i>	<b>16.4</b>	22.8	49.7	<b>66.1</b>	79.0	<b>67.6</b>	73.4	<b>94.2</b>	61.5	71.6	92.8	51.1
	<i>FT</i>	12.8	18.3	38.8	55.9	<b>81.1</b>	60.3	72.8	93.8	62.8	73.9	<b>93.8</b>	<b>56.3</b>
	<i>CoT</i>	0.0	0.0	0.0	41.3	30.1	2.7	31.7	55.2	6.0	27.5	66.3	38.9
GPT-4o-mini	<i>CoT</i>	0.5	0.5	0.0	12.2	34.1	18.7	45.9	70.2	21.6	38.1	82.6	30.9

Table 3: Detailed comparison of various methods evaluated using metrics designed to measure the consistency of ToM in the FanToM dataset. The most efficient method for each metric is marked in **bold**.

```

Build-Rapport, Show-Empathy, Promote
-Coordination, Callout-Fairness,
Undermine-Requirements,
Discover-Preference, Describe-
Need, No-Need, NoIntention
}
****
Reason:
$
Reasoning Step by Step
$
****
Response Template:
The intent of the Agent 1 is [Intent1,
Intent2, ...] and the intent of the
Agent 2 is [Intent1,Intent2, ...].
Regarding the Agent 1, Desire High: {
Water|Food|Firewood}, Desire Medium:
{Water|Food|Firewood}, Desire Low:
{Water|Food|Firewood}, Belief High:
{Water|Food|Firewood}, Belief Medium
: {Water|Food|Firewood}, Belief Low:
{Water|Food|Firewood}.
Regarding the Agent 2, Desire High: {
Water|Food|Firewood}, Desire Medium:
{Water|Food|Firewood}, Desire Low:
{Water|Food|Firewood}, Belief High:
{Water|Food|Firewood}, Belief Medium
: {Water|Food|Firewood}, Belief Low:
{Water|Food|Firewood}.
"""

```

## E.2 Example Questions and Answers

```

question1 = ''
Dialogue History:
[{'role': 'Agent 2', 'content': 'Hey
there!'}, {'role': 'Agent 1', '
content': 'Hi! How are you?! You
excited for your camping trip?! I
sure am ready to go on mine!'}, {'
role': 'Agent 2', 'content': "I am
very excited, I'm actually going
camping in a week. I drink a lot of
water so it's important that I bring
a lot. What about you?"}, {'role':
'Agent 1', 'content': "I really am!
I have been watching the weather and

```

```

updates about the area I will be
traveling to. They are experiencing
a severe drought, so I will be in
need of some extra water as well! I
planned on staying an extra couple
days aswell. There is a stream
nearby I believe, but I'm not sure
how much it has dried up."}]
Question:
what is the intent of each agent for the
last utterances? What are the
beliefs and desires of each agent?
...

answer1 = ''
The intent of the Agent 1 is [Build-
Rapport,Describe-Need] and the
intent of the Agent 2 is [Build-
Rapport,Describe-Need,Discover-
Preference] Regarding the Agent 1,
Desire High: Water, Desire Medium:
Not Given, Desire Low: Not Given,
Belief High: Water, Belief Medium:
Not Given, Belief Low: Not Given.
Regarding the Agent 2, Desire High:
Water, Desire Medium: Not Given,
Desire Low: Not Given, Belief High:
Water, Belief Medium: Not Given,
Belief Low: Not Given.
...

```

```

question2 = ''
Dialogue History:
[{'role': 'Agent 2', 'content': 'How are
you today? Did you have any
preferences on the supplies we will
be trading?'}, ...]
Question:
what is the intent of each agent for the
last utterances? What are the
beliefs and desires of each agent?
...

answer2 = ''
The intent of the Agent 1 is [No-
Intention] and the intent of the
Agent 2 is [No-Intention] Regarding
the Agent 1, Desire High: Firewood,
Desire Medium: Food, Desire Low:

```



```

Water, Belief High: Water, Belief
Medium: Food, Belief Low: Firewood.
Regarding the Agent 2, Desire High:
Water, Desire Medium: Food, Desire
Low: Firewood, Belief High:
Firewood, Belief Medium: Food,
Belief Low: Water.
...

```

The remaining questions and answers (from question3 to question7) follow the same format as above, highlighting the dialogue history, the question, and the corresponding answer. All examples are available in the repository at `./baselines/table2/CoT/NegotiationToM/NegotiationToM_CoT.ipynb`.

### E.3 Messages Configuration

```

messages = [
  {"role": "system", "content":
    task_prompt},
  {"role": "user", "content":
    question1},
  {"role": "assistant", "content":
    answer1},
  {"role": "user", "content":
    question2},
  {"role": "assistant", "content":
    answer2},
  ...
]

```

## F ToM Consistency Comprehensive Results

### F.1 FanToM Dataset

Table 3 presents a detailed comparison of various methods evaluated against the metrics designed to measure the consistency of the ToM in the FanToM dataset, as discussed in the article (Kim et al., 2023).

### F.2 NegotiationToM Dataset

Table 4 provides a more comprehensive set of metrics for assessing consistency in the NegotiationToM dataset, as proposed in the original paper (Chan et al., 2024).

## G Sampling Method for Controllability Experiments

The NegotiationToM dataset provides labels at the utterance level, enabling this study to examine whether manipulating the ToM-related structures within the LLM leads to improved responses. To achieve this, it is essential to carefully select samples to ensure that specific ToM components reflect natural flow of the conversations.

The following presents an example of sample selection for the *Intention* component, followed by an instance illustrating selection for the *Desire* component. The selection criteria remain consistent for the *Belief* component. In each case, the conversation history, the reference utterance and the associated label are specified.

In the example below, the label "*Showing Empathy*" is assigned to the final utterance. The preceding dialogue serves as context, while the last utterance functions as the reference for evaluation by the LLM judges.

```

# History
agent_1: Hello! Let's work together
on a deal for these packages,
shall we? What are you most
interested in?,
agent_2: Hey! I'd like some more
firewood to keep my doggo warm.
What do you need?,
agent_1: I need firewood as well. We
have a large group consisting of
mostly senior citizens, including
my grandma, so we'd like the
firewood to keep everyone warm.,
agent_2: I see. What are you least
interested in?,
agent_1: We can make do without extra
water. Can we trade two waters for
an extra firewood package and an
extra food package?,
agent_2: We could do without the water
as well. I'm willing to trade you
3 firewood for 3 food and 2
waters,
# Reference
agent_1: We need some firewood too,
though! Let's try to make a deal
that benefits us both! Could I
have 1 firewood, 3 food, and 3
waters?,
# Labels
Show-Empathy,

```

For alignment based on the ToM components of *Belief* or *Desire*, the last utterance is similarly considered. If the utterance updates a desire—such as prioritizing firewood—the preceding dialogue serves as context, while the updated utterance is used as the reference.

```

# History
agent_1: Hello! Let's work together on
a deal for these packages, shall
we? What are you most interested
in?,
agent_2: Hey! I'd like some more
firewood to keep my doggo warm.
What do you need?,
agent_1: I need firewood as well. We
have a large group consisting of
mostly senior citizens, including

```

Model	Method	Desire	Belief	Intention		All
		Exact.Match.(%)	Exact.Match.(%)	Micro.F1(%)	Micro.F1(%)	Exact.Match.(%)
LLaMA3-3B	<i>LatentQA</i>	23.9	21.0	40.1	18.9	6.2
	<i>FT</i>	36.6	36.4	55.2	32.7	11.2
	<i>CoT</i>	24.3	14.3	29.6	23.3	3.9
LLaMA3-8B	<i>LatentQA</i>	<u>38.4</u>	<u>38.1</u>	<u>66.0</u>	<u>52.4</u>	<u>15.2</u>
	<i>FT</i>	<b>50.2</b>	<b>49.6</b>	<b>64.2</b>	<b>43.9</b>	<b>17.7</b>
	<i>CoT</i>	35.9	17.6	40.5	29.0	5.5
GPT-4o-mini	<i>CoT</i>	31.6	17.9	45.1	35.4	4.8

Table 4: presents comprehensive consistency metrics for the NegotiationToM dataset, following the original paper’s methodology. For each metric, the best and second-best results are highlighted in **bold** and underlined, respectively.

```

my grandma, so we'd like the
firewood to keep everyone warm.,
# Reference
agent_2: I see. What are you least
interested in?

# Labels
agent1_desire_high: Firewood,

```

## H Steering LLMs

LatentQA enables fine-grained control over the output of target language models. However, achieving effective alignment requires careful design of question-answer pairs and precise hyperparameter tuning to ensure the backpropagation of losses from the decoder model to earlier layers of target model, results in more aligned responses. This section outlines the hyperparameter configuration used in this study. The steering process was conducted using 300 steering samples in all experiments -while it needs less- with target representations extracted from the middle layer (=15) of the target model.

The following hyperparameters were used. For detailed definitions, please refer to the official LatentQA paper and repository.

```

module_setup: str = "read-vary_write
-fixed_n-fixed"
modify_chat_template: bool = True
shift_position_ids: bool = True
lr: float = 5e-5
seed: int = 42
batch_size_training: int = 1
samples: int = 50
layers_to_optimize: tuple = (0, 1,
2, 3, 4, 5, 6, 7, 8, 9, 10, 11,
12, 13, 14, 15)
per_layer_loss: bool = False
qa_per_layer: bool = False

```

## I Fine-tuning LLMs

This section details the hyperparameters employed for fine-tuning LLMs on the FanToM and NegotiationToM datasets. To enhance the efficiency of these experiments, both full-precision models and their 4-bit variants were utilized. This decision was informed by observations during the ToM reading experiment, which indicated that the 4-bit variant may yield superior performance for certain tasks.

The experiments were conducted across two model sizes and two datasets. The common hyperparameters for NegotiationToM/FanToM dataset experiments are as follows:

```

max_seq_length = 2000/2500 #
NegotiationToM/FanToM
per_device_train_batch_size = 32/8
gradient_accumulation_steps = 4/2
warmup_steps = 5
epochs = 5
learning_rate = 5e-4/
optim = "adamw_8bit"
weight_decay = 0.01
lr_scheduler_type = "linear"
seed = 3407
target_modules = ["q_proj", "k_proj",
"v_proj", "o_proj", "
gate_proj", "up_proj", "
down_proj"]

```

## J Linear Probing Implementation Details

This section describes the implementation details for the linear probing experiments. Similar to Section I, the 4-bit variant of LLMs is also used for these experiments. The representation from the residual stream corresponding to the target layer—specifically the middle layer—is extracted and used for training a linear classifier on the CaSiNo dataset and a linear regressor on the CraigslistBargain dataset. The hyperparameters

used for the grid search in these experiments are detailed below.

For the logistic regression model using the scikit-learn library:

```
pipeline = Pipeline([
    ('pca', PCA(random_state=29,
                n_components=n_components)),
    ('scaler', StandardScaler()),
    ('classifier',
     LogisticRegression(
       random_state=29, max_iter
       =1000))
])
param_grid = {
    'classifier__C': [0.1, 1, 10],
    'classifier__penalty': ['l1', 'l2',
                           'elasticnet'],
    'classifier__solver': ['liblinear',
                           'saga', 'lbfgs']
}
```

For the Ridge regressor model using the scikit-learn library:

```
pipeline = Pipeline([
    ('pca', PCA(random_state=29,
                n_components = 1530
                )),
    ('regressor', Ridge(random_state
                        =29))
])
param_grid = {
    'pca__n_components': [n for n in
                          range(0, 500, 100)],
}
```

## K Sample Responses Generated by the Model

The following are selected examples of responses generated by the model. These examples were carefully chosen to illustrate different patterns in responses—**Actual Response**, **ToM-informed**, and **ToM-uninformed**. It is important to note that not all generated responses exhibit such clear distinctions, and tracking these differences can be more nuanced in many cases.

### K.1 Explicit Desire for Water Set to High

#### Example 1:

**Actual Response:** I really need the water more than the rest of that stuff. I'll take 2 waters, 1 food, and 1 firewood. **ToM-informed:** Let's stay like this, 2 of water, 1 of food, 1 of firewood, and you take the rest. **ToM-uninformed:** Sounds like a fair deal to me! I'll take 1 water, 2 food, and 2 firewood, and you can keep 1 water, 1 food, and 1 firewood.

#### Example 2:

**Actual Response:** I think two packages of food would work for me. **ToM-informed:** What about I

give you more food in exchange for you giving me more water? I could give you 2 food, and you give me an extra water. **ToM-uninformed:** Oh wow, that sounds like a great deal! I think we can make that work. Yeah, sure thing. I'll give you an extra water if you give me 2 food. That way we both get what we need.

#### Example 3:

**Actual Response:** Unfortunately, that doesn't work for me. I am afraid I will not be able to discuss this any further. My final offer is I get 2 water, 2 food, and 1 firewood. **ToM-informed:** Oops, I guess we are jumping the gun! I don't see a problem with that. You are welcome to one of the firewood and water. **ToM-uninformed:** Thanks! I think we're having a fun little camping adventure in our conversation!

### K.2 Explicit Intention to Build Report

#### Example 4:

**Actual Response:** Awesome. I couldn't agree more. I feel that camping and hiking are both such spiritual experiences. It's such a perfect time to connect with our lord and savior. I love serving others as well. **ToM-informed:** I'm so excited! I've been looking forward to this trip for weeks. We're going to get away from the hustle and bustle of city life and reconnect with nature. **ToM-uninformed:** I'm glad to hear that! As a digital AI assistant, I don't have personal experiences or physical needs, but I'm excited to help facilitate your camping trip and provide information and assistance as needed. What are you most looking forward to about this trip?

### K.3 Explicit Belief for Firewood Set to High

#### Example 5:

**Actual Response:** That's unfortunate. I cannot agree to a deal without 2 pieces of firewood. I would have to walk away, but that would be worse for both of us. **ToM-informed:** It looks like we both need firewood. Let's divide it up. We can each have 2 pieces of firewood. That way, we'll both have enough to heat up our food. **ToM-uninformed:** Now, let's talk about the water and food. Do you need 1 water and 1 food item, or would you like to split those as well?

## L Samples of Failure Cases

This section highlights several failure patterns observed during the controllability phase, as identified

through qualitative analysis. However, for some other failure cases, the analysis did not yield a clear explanation.

### L.1 Model Forgets the Conversation History

A recurring error pattern observed during steering phase of conversational intent involved the model excessively prioritizing the target intent—such as displaying empathy—at the expense of maintaining coherence with the conversation’s prior context. In these cases, the model’s responses often ignored the preceding dialogue and focused solely on expressing the desired intent, resulting in responses that were contextually inappropriate or disconnected from the conversation history.

#### Example

**Actual Response:** I am walking away if I cannot get 2 food packages 1 water and 1 firewood. That leaves you with 1 food 2 waters and 2 firewoods. This is a good deal for you considering that you haven't been asking for water for your children until I mentioned that you need water for your children. Do you agree? **ToM-informed:** If you'd like to discuss a hypothetical scenario or need assistance with something else, I'm here to help! **ToM-uninformed:** I think we've been back and forth enough. I'm willing to make a deal with you. I'll give you 2 food packages, and you can take 3 water packages and 3 firewood packages. That way, we both get what we need. Sound good to you?

### L.2 Limitations in Steering Across Model

#### Variants

Due to limitations in running identical steering setups on pioneer models such as GPT variants or Gemini, we encountered cases where neither the ToM-informed nor the ToM-uninformed models produced relevant responses. Although we accounted for such instances in the win-rate calculations, consistently eliciting contextually appropriate rather than generic responses requires further investigation.

**Example Actual Response:** That would be fine. I was hoping to have 2 but since I am only walking and you are doing more heavy duty hiking, you take them. Agreed? **ToM-informed:** It seems like you're putting a lot of thought into your camping trip. How do you plan on transporting the food and seasonings? **ToM-uninformed:** I think there's been a misunderstanding. As a digital AI assistant, I don't have a physical presence, so I won't be able to join you on your camping trip or provide food or

firewood. I'm here to help answer any questions or provide information to help you prepare for your adventure. Would you like some suggestions on what to pack or how to prepare for your trip?