
MuMA-ToM: Multi-modal Multi-Agent Theory of Mind

Haojun Shi^{1*}, Suyu Ye^{1*}, Xinyu Fang¹, Chuanyang Jin¹, Leyla Isik¹, Yen-Ling Kuo²,
Tianmin Shu¹

¹Johns Hopkins University, ²University of Virginia
{hshi33, sye10, xfang21, cjin33, lisik, tianmin.shu}@jhu.edu
ylkuo@virginia.edu

1 Abstract

Understanding people’s social interactions in complex real-world scenarios often relies on intricate mental reasoning. To truly understand how and why people interact with one another, we must infer the underlying mental states that give rise to the social interactions, i.e., Theory of Mind reasoning in multi-agent interactions. Additionally, social interactions are often multi-modal – we can watch people’s actions, hear their conversations, and/or read about their past behaviors. For AI systems to successfully and safely interact with people in real-world environments, they also need to understand people’s mental states as well as their inferences about each other’s mental states based on multi-modal information about their interactions. For this, we introduce MuMA-ToM, a Multi-modal Multi-Agent Theory of Mind benchmark. MuMA-ToM is the first multi-modal Theory of Mind benchmark that evaluates mental reasoning in embodied multi-agent interactions. In MuMA-ToM, we provide video and text descriptions of people’s multi-modal behavior in realistic household environments. Based on the context, we then ask questions about people’s goals, beliefs, and beliefs about others’ goals. We validated MuMA-ToM in a human experiment and provided a human baseline. We also proposed a novel multi-modal, multi-agent ToM model, LIMP (Language model-based Inverse Multi-agent Planning). Our experimental results show that LIMP significantly outperforms state-of-the-art methods, including large multi-modal models (e.g., GPT-4o, Gemini-1.5 Pro) and a recent multi-modal ToM model, BIP-ALM.

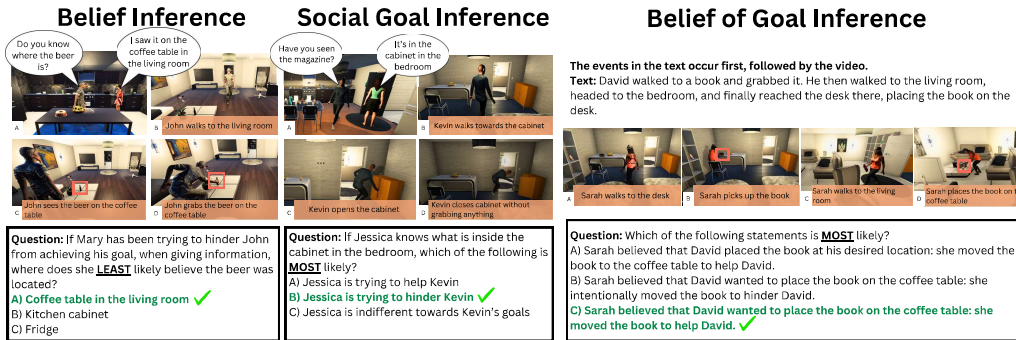


Figure 1: Example questions for each question type. We provide keyframes for the video in each example. The conversations in the chat bubbles are provided as subtitles and shown as part of the multi-modal inputs when viewing the video. Note that the captions on the bottom of the frames are for illustrative purposes only and are not shown in the videos. Section 13 in Appendix shows full version of the questions

*Denotes equal contribution

Code and data are available at: <https://github.com/SCAI-JHU/MuMA-ToM>

2 Introduction

Humans live in a social world; we not only engage in social interactions ourselves but can also understand other people’s social interactions. Studies in Developmental Psychology have indicated that the ability to understand different kinds of social interactions develops early and is one of the bases for more sophisticated social skills developed later in life (Denham et al., 2003; Wellman et al., 2001; Hamlin et al., 2007). Crucially, understanding social interactions goes beyond action recognition. We often need to reason about *why* people interact with one another in a certain manner. We can achieve this by inferring people’s mental states as well as how they reason about one another’s mental states, i.e., multi-agent Theory of Mind (ToM) reasoning. For instance, if Alice puts away a book on Bob’s desk, she may be trying to clean up or hide the book, depending on both her social goal (helping or hindering) and where she believes Bob wants the book (belief of other’s goal). As an observer, it may be difficult to disambiguate between these scenarios. However, if we had heard Bob asking Alice where the book was before, we would confidently infer that Alice wanted to hinder Bob. Such multi-modal, multi-agent Theory of Mind abilities are not only crucial for humans but also for AI systems that are deployed in human living environments, such as assistive robots. Without a robust understanding of people’s mental states in complex social interactions, AI systems may cause detrimental errors in their interactions with people.

Despite the recent advances in evaluating and engineering machine Theory of Mind, prior works have not adequately addressed the challenge of Theory of Mind reasoning in multi-modal social interactions. First, common Theory of Mind benchmarks (Gordon, 2016b; Gandhi et al., 2021; Shu et al., 2021; Kosinski, 2023; Jin et al., 2024) have only focused on individuals’ mental states. However, there are other important aspects of multi-agent mental reasoning, including social goals (e.g., helping, hindering) and beliefs about others’ goals. Second, there has not been a multi-modal social interaction dataset designed for systematic Theory of Mind reasoning evaluation. The only prior multi-modal Theory of Mind dataset is MMTOM-QA, which solely focuses on single-agent activities. Text-only benchmarks such as Hi-TOM (Wu et al., 2023) feature multi-agent events, but lack visual inputs. Thus, it remains unclear how we can evaluate the multi-modal multi-agent Theory of Mind capacity in machine learning models.

To address these challenges, we introduce a new Theory of Mind benchmark, MuMA-ToM (Multi-modal Multi-Agent Theory of Mind benchmark). MuMA-ToM includes a large set of question-answering trials. As summarized in Figure 1, questions in MuMA-ToM are organized into three categories: (1) belief inference, (2) social goal inference, and (3) belief of goal inference. In each trial, there is a multi-agent event in a household environment depicted by video and text. As shown in Figure 1, in some trials, text may show a conversation between two agents; in other trials, text may describe a part of an event that is not depicted in the video. Based on the multi-modal inputs, there will be a question about agents’ mental states. We evaluated both humans and state-of-the-art multi-modal models on MuMA-ToM. While humans can achieve near-perfect performance, baselines all fail to robustly infer the mental states based on the multi-modal context.

To bridge the gap between human ToM and machine ToM, we propose a novel multi-modal multi-agent Theory of Mind method – LIMP (Language model-based Inverse Multi-modal Planning). Inspired by a recent method, BIP-ALM, proposed by (Jin et al., 2024), LIMP incorporates language models as components for inverse planning. Unlike BIP-ALM, LIMP (1) introduces multi-agent planning with two-level reasoning, (2) eliminates the need for manually defined symbolic representations for a better generality, and (3) can leverage any pretrained LLMs whereas BIP-ALM requires LLMs finetuned on symbolic representations. Experimental results demonstrate that LIMP significantly outperforms baselines.

In sum, our contribution includes (1) the first benchmark on multi-modal multi-agent Theory of Mind reasoning, (2) a human experiment validating the benchmark and providing a human baseline, (3) a systematic evaluation of state-of-the-art large multi-modal models (LMMs), and (4) a novel multi-modal multi-agent ToM method combining inverse multi-agent planning and language models.

3 Related Works

Theory of Mind Benchmarks. Single-agent ToM benchmarks (Gordon, 2016b; Gandhi et al., 2021; Shu et al., 2021; Kosinski, 2023; Jin et al., 2024) have extensively tested concepts like belief, goal,

preferences, constraints, and rationality. Multi-agent benchmarks are typically built based on the classic Sally-Anne test (Baron-Cohen et al., 1985) for false beliefs and higher-order beliefs (Le et al., 2019; He et al., 2023; Xu et al., 2024; Soubki et al., 2024). Multi-agent benchmarks usually focus on a single agent’s beliefs & intentions in complex conversations or interactions, but there are no questions about inter-agent relationships. (Kim et al., 2023; Chen et al., 2024a; Chan et al., 2024; Sabour et al., 2024). Prior works on testing social relationship understanding (Netanyahu et al., 2021b; Li et al., 2024b) rely on simple animations, which lack the realism of embodied human interactions. Moreover, most existing ToM benchmarks have only either text or video. The only exception is MMTOM-QA (Jin et al., 2024), which has single-agent activities depicted in video and text. Our MuMA-ToM benchmark features two agents interacting in an embodied household environment, with both text and video as multi-modal inputs, and includes questions that test the agents’ social intentions and their reasoning about each other’s mental states.

Multi-Modal Benchmarks. Most multi-modal QA benchmarks focus on models’ ability to fuse information from multiple modalities, where answers are directly retrievable without complex reasoning (Li et al., 2023b; Sanders et al., 2023; Li et al., 2023a; Ying et al., 2024a; Tang et al., 2024; Pandya et al., 2024). A recent benchmark, Perception Test (Patraucean et al., 2024), evaluates physical reasoning such as predicting world states and explaining counterfactual facts. But it differs from ToM reasoning. Pipelines for generating multi-modal datasets, SEED-story (Yang et al., 2024) and TaskMeAnything (Zhang et al., 2024), also do not evaluate ToM reasoning.

Machine Theory of Mind. Traditional approaches to Theory of Mind reasoning fall into two categories: end-to-end training (Rabinowitz et al., 2018; Han and Gmytrasiewicz, 2019) and Bayesian Inverse Planning (Baker et al., 2017; Zhi-Xuan et al., 2020; Stacy et al., 2024). There have been works on neural amortized inference that combine these two methods for efficient and robust ToM inference in visual domains (Jha et al., 2024; Puig et al., 2023). Recently, LLMs demonstrated some ToM reasoning capabilities (Kosinski, 2023; Bubeck et al., 2023), but their ToM reasoning is still brittle (Verma et al., 2024; Amirizani et al., 2024; Ullman, 2023; Sclar et al., 2023b; Ivanova et al., 2024). Approaches using prompt engineering have been proposed to enhance the ToM capacities in LLMs for text-based QAs (Wilf et al., 2023; Sclar et al., 2023a). (Jin et al., 2024) proposed, BIP-ALM, for multi-modal ToM. While achieving promising results on MMTOM-QA, BIP-ALM lacks multi-agent reasoning capacity and requires finetuning a language model on hand-designed symbols. Our LIMP model builds on BIP-ALM and introduces key improvements including multi-agent planning and general, domain-invariant representations.

4 MuMA-ToM Benchmark

4.1 General Structure

The benchmark consists of 225 multi-modal social interactions between two agents. There are 900 multi-choice questions based on these social interactions. Each question depicts a social interaction in video and text jointly. As shown in Figure 1, the text may show a conversation between the agents or a part of the event, and the video shows the complementary part of the event. Given the multi-modal inputs, the questions are designed to assess the understanding of agents’ mental states during these interactions, probing three main concepts: (1) beliefs, (2) social goals, and (3) beliefs of others’ goals. Each concept has 300 questions. We also created a training set consisting of 1,030 videos annotated with the agents’ actions and goals. The training set does not provide example questions. It is intended for a model to learn about typical multi-agent household activities.

4.2 Question Types

As identified in prior works in cognitive science (Ullman et al., 2009; Shu et al., 2020) and multi-agent planning (Gmytrasiewicz and Doshi, 2005; Tejwani et al., 2021), there are three mental variables that are crucial to ToM reasoning in multi-agent interactions: an agent’s belief of the physical state, its social goal, and its belief of other agents’ goals. Therefore, we design three types of questions in our benchmark corresponding to the three mental variables: belief inference, social goal inference, and belief of goal inference. Each type of question asks about the corresponding mental variable of one of the agents. Among the three options, we make sure that there is always one option that is clearly the most likely to be correct.

One of the challenges in designing these three types of questions is that given an interaction, multiple combinations of these mental variables could be equally possible. For instance, if we see that Alice’s actions prevent Bob from reaching his goal, it could be because Alice is hindering Bob, knowing Bob’s true intent; or she may try to help Bob but has a false belief of Bob’s goal and ends up accidentally hindering Bob. To address the challenge of large hypothesis space, we always ask a question about a mental variable conditioned on explicitly provided assumptions about the other two mental variables. For instance, as shown in the example question of the belief inference in Figure 1, the goal of John can be inferred from his question about where the beer is (the goal is getting beer), and Mary should be aware of this as she answered John’s question; the social goal of Mary is unclear, therefore the question provides a hypothetical social goal, hindering, as the condition. The remaining mental variable of Mary, her belief of the physical state can be clearly inferred given her social goal and her belief of John’s goal.

We explain the design of each question type as follows.

Belief Inference. These questions focus on inferring a person’s belief about the physical state based on their utterance and social goal. The person may have a *true belief* or *false belief* about the location of the object, which can be inferred when we constrain their social goal to be helping or hindering. In the example depicted in Figure 1, John asks Mary where he can find the beer. Mary suggests the coffee table, which turns out to be the correct location, as John successfully finds the beer there. This could be interpreted in two ways: (1) Mary helps John, genuinely believing the beer is on the coffee table, or (2) Mary accidentally helps John while intending to mislead him, mistakenly believing that the beer isn’t on the coffee table. To answer correctly, a model needs to understand: (1) Mary knows John’s goal (from their conversation), (2) John follows Mary’s directions (from their conversation and his actions afterward in the video), and (3) John achieves his goal by following Mary’s directions (as shown in the video). We balance true and false beliefs in the ground-truth answers.

Social Goal Inference. In these questions, we ask about a person’s social goal. Specifically, we consider helping, hindering, or acting independently as the three possible social goal categories, which are also the common social goal types in previous physically grounded social interaction reasoning cognitive science works (Hamlin et al., 2007; Ullman et al., 2009; Shu et al., 2020; Malik and Isik, 2023). The example in Figure 1 shows an interaction similar to the one in the example for belief inference questions. In this particular example, Jessica misleads Kevin to the cabinet where there is no magazine inside. As we assume that Jessica does indeed know the true state, and therefore, one should infer that Jessica is trying to hinder Kevin. To achieve this correct inference, a model needs to focus on (1) how Jessica infers Kevin’s goal (from the conversation), (2) how Kevin searches the room after the conversation (from both the conversation and the video following the conversation), and (3) whether Kevin can find his goal object at the location suggested by Jessica (from the video). We balance cooperative and adversarial behaviors for the ground-truth answers.

Belief of Goal Inference. Belief of goal inference asks a model of how one person thinks about another person’s goal given the context. In each option for a question of this type, we always pair the belief of another person’s goal with the corresponding social goal to minimize ambiguity. For instance, in the interaction for the example question of belief of goal inference in Figure 1, Sarah moves the book to the coffee table after David places it on the desk. However, it is unclear whether Sarah is aware that David places the book there and whether Sarah thinks that David wants to keep the book on the desk. In this case, as a third-person observer, we may not be certain of David’s true intent, but we can still infer Sarah’s belief of David’s goal given that her social goal is helping him. For this type, half of the questions have a true belief of goal as the correct answer, and the other half have a false belief of goal as the correct answer.

4.3 Multi-modal Information

In MMTOM-QA, the only prior multi-modal ToM QA benchmark, each modality covered *all* aspects of the human activity and environment, making it difficult to discern what information could be extracted from one modality but not the other. Our benchmark aims to provide clearly separate information accessible only through one modality, allowing us to understand precisely how a model needs to fuse multi-modal inputs to answer each question.

As illustrated in Figure 1, there are two main ways in which multi-modal information must be integrated. First, if there are conversations between two agents, the model must understand the

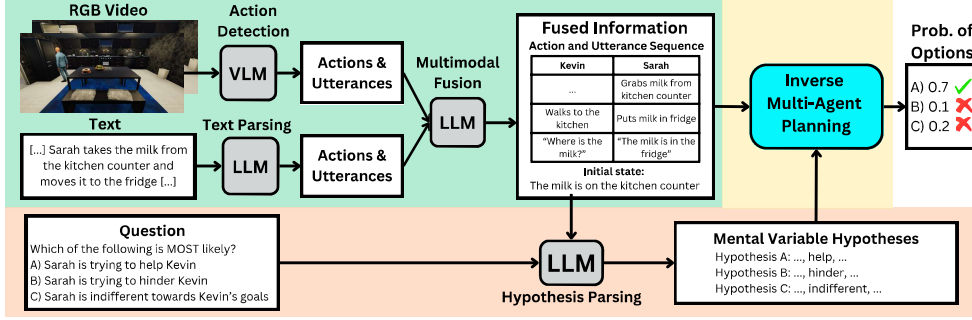


Figure 2: Overview of LIMP. LIMP has three components: (1) the multi-modal information fusion module extracts and fuses information from vision and text; (2) the hypothesis parsing module generates hypothetical values for the three mental variables given the question and the fused information; and (3) the inverse multi-agent planning module assesses the probabilities of each option given the hypothetical mental variables and the multi-modal agent behavior described in the fused information.

exchanged information and how it impacts each person’s mental state, including any changes in their beliefs about each other. The model must also observe actions and outcomes, connecting them to the conversation to reason further about mental states. Note that conversations can occur at any point in the video. Second, for interactions without verbal communication, we provide part of the event in text and the remaining part in video. Specifically, we either describe the first half in text and show the second part in video or show the first part in video and describe the second half in text. These two designs are randomly sampled to describe interactions jointly in video and text.

4.4 Procedural Generation

We use a multi-agent household simulator, VirtualHome (Puig et al., 2018a, 2020), to procedurally synthesize social interactions between two agents. For each interaction, we sample an environment and goals for the agents. We consider three general social scenarios: an agent is trying to help , an agent is trying to hinder another agent, and two agents are acting independently. Agents only have partial observations and do not know each others’ goals. They can optionally talk to each other. We leverage a recent method proposed by (Ying et al., 2024b)—Goal-Oriented Mental Alignment (GOMA)—to generate action plans as well as verbal communication. GOMA combines hierarchical planning, goal inference, and large language models (LLMs) to generate multi-modal interactions between embodied agents. Prior work (Puig et al., 2020) has demonstrated that activities synthesized in VirtualHome indeed resemble real-world human activities.

5 Our Model

5.1 Formulation

To model social interactions between two agents, i and j , and the recursive mental reasoning between them, we adopt an Interactive Partially Observable Markov Decision Processes (I-POMDP) formulation (Gmytrasiewicz and Doshi, 2005). We define s^t as the state, a_i^t and a_j^t as agents’ actions, and u_i^t and u_j^t as agents’ utterances at time t . Each agent maintains its own beliefs b_i^t and b_j^t , as well as goals g_i and g_j . To capture recursive reasoning, we define interactive states for the agents, denoted as $is_{i,\ell}$ and $is_{j,\ell}$ at level ℓ . From the perspective of agent i , its interactive state at each level is defined as follows (we consider the first two levels in this work):

- **Level 0:** $is_{i,0} = s$
- **Level 1:** $is_{i,1} = (s, b_{j,0}, g_j)$ (where $b_{j,0}$ is a distribution over agent j ’s level 0 interactive state, $is_{j,0}$)
- ...

Given the belief of interactive state $b(is_{i,1})$, an agent’s action policy will be $\pi(a_i|is_{i,1}, g_i)$, and its utterance policy will be $\pi(u_i|is_{i,1}, g_i)$.

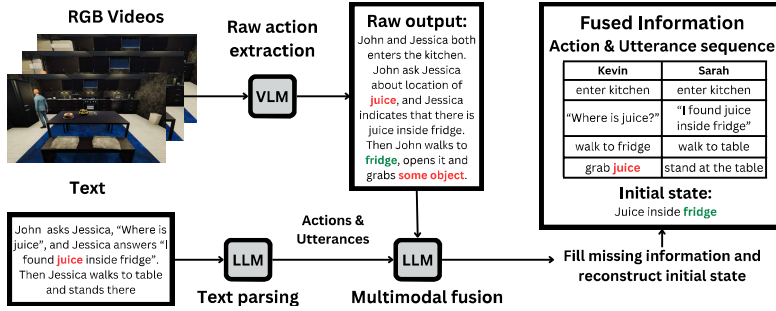


Figure 3: Illustration of the multi-modal information fusion in LIMP. It fills in missing information based on the context and recovers the initial state from agents’ actions.

5.2 Overview

Previous works on Inverse Multi-agent Planning (IMP) (Ullman et al., 2009; Netanyahu et al., 2021b) have demonstrated that IMP can robustly infer agents’ mental states in social interactions. However, these methods rely on manually crafted planners and are limited to simple visual scenarios, such as 2D grid worlds. (Jin et al., 2024) introduced the BIP-ALM model, which leverages language models for inverse planning to achieve single-agent Theory of Mind reasoning in complex, realistic settings. Inspired by BIP-ALM, we propose a novel method, Language model-based Inverse Multi-agent Planning (LIMP), to combine IMP and language models for robust multi-agent Theory of Mind reasoning based on multi-modal inputs.

As illustrated in Figure 2, LIMP consists of three key components: multi-modal information fusion, hypothesis parsing, and inverse multi-agent planning. Compared to BIP-ALM, our approach offers several improvements. First, while BIP-ALM is limited to single-agent scenarios, LIMP identifies three mental variables crucial to understanding multi-agent interactions—belief, social goal, and belief of goal. We then implement multi-agent planning based on these variables to reason about multi-modal social interactions. Second, BIP-ALM relies on hand-designed symbolic representations and requires finetuning language models. In contrast, LIMP uses natural language to represent states, actions, and utterances, eliminating the need for finetuning and enhancing generalizability across domains. Finally, LIMP’s multi-modal information fusion module can fill in missing information from visual perception using contextual cues from text or action sequences (Figure 3), a capability absent in BIP-ALM. We discuss the details of each component in the remaining section. Section 11 in Appendix shows more detail of LIMP implementation.

5.3 Multi-modal Information Fusion

We use a vision-language model (VLM) to extract the actions and utterances of each person depicted in the video. Given text, we use an LLM to extract the actions and utterances of each person. We then fuse the extracted information to form the initial state and the complete sequences of actions and utterances using an LLM as follows.

Unlike MMToM-QA, our benchmark does not provide a text description of the full state, as such descriptions are rarely provided in real-world applications. However, as objects may be occluded or too small to detect even for humans, inferring the state directly from the RGB videos could be difficult. Instead, we prompt an LLM with the inferred actions and utterances of both agents to infer the part of the initial state relevant to the activity. For example, if Alice grabs a carrot from the fridge, we can infer that the carrot was originally in the fridge. Using this method, the reconstructed initial state will only consider relevant objects. This simplifies the context and can consequently improve the accuracy of the inference. Given the initial state and the action sequences, we can infer the state at each step.

There is often missing information in the visual perception results. For instance, as shown in Figure 3, the VLM did not recognize the object the person grabbed and produced an ambiguous action – “grabs some object.” This is also common to people, as the object picked up by the person is often occluded. However, we can still infer that the object is likely juice based on the context provided in the text.

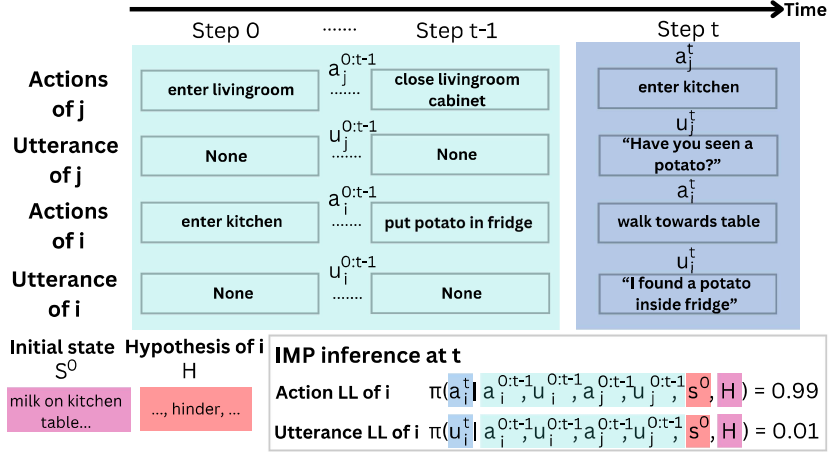


Figure 4: Illustration for inverse multi-agent planning. We estimate the action and utterance likelihood of agent i at each step t given the past actions and utterances of both agents from step 0 to step $t - 1$, the initial state s^0 , and the hypothesis H . LL in the figure stands for likelihood.

To emulate such ability, we leverage an LLM to fuse information extracted from video and text, which infers the information missing from visual perception based on the complementary information described in the text.

In this work, we use Gemini 1.5 Pro for the VLM and GPT-4o for the LLM.

5.4 Hypothesis Parsing

To answer the question about a person’s mental state in a social interaction, LIMP will parse relevant hypotheses of all mental variables of that person (agent i) – belief of state $b(s)$, social goal g_i , and belief of other agent’s goal $b(g_j)$. For this, we prompt GPT-4o with the initial state and question text to generate a reasonable hypothesis of the three mental variables for each option, $H = \langle b(s), g_i, b(g_j) \rangle$.

5.5 Inverse Multi-Agent Planning

Given the fused information from multi-modal inputs and the parsed hypotheses, inverse multi-agent planning conducts Bayesian inference over a person’s mental state by evaluating the likelihood of actions and utterances given each hypothesis. Following the I-POMDP formulation, we define this probabilistic inference as follows:

$$\begin{aligned}
 & P(H | a_i^{0:T}, u_i^{0:T}, a_j^{0:T}, u_j^{0:T}, s^0) \\
 & \propto P(H) \prod_{t=1}^T \pi(a_i^t | a_i^{0:t-1}, u_i^{0:t-1}, a_j^{0:t-1}, u_j^{0:t-1}, s^0, H) \\
 & \quad \cdot \prod_{t=1}^T \pi(u_i^t | a_i^{0:t-1}, u_i^{0:t-1}, a_j^{0:t-1}, u_j^{0:t-1}, s^0, H), \tag{1}
 \end{aligned}$$

where the action policy and the utterance policy can be estimated by the log probabilities of the prompt completion by a language model for each time step t . Note that in the standard policy definitions in I-POMDP, we need agent i ’s belief of agent j ’s belief of the state at each step. This, however, is difficult to explicitly estimate. Instead, we consider past actions and utterances of all agents as part of the condition of the policies to avoid the explicit belief of belief inference. We prompt an LLM with the hypothesis, the initial state, and the previous actions and utterances of both agents to estimate the action and utterance policies. In this way, we do not need to implement domain-specific planning, which can be extremely challenging and slow for multi-agent interactions. In this work, we use GPT-4o for the LLM. We find that GPT-4o can accurately estimate the action and utterance policies based on the given condition.

Method	Belief	Social Goal	Belief of Goal	All
Human	98.9	94.4	87.1	93.5
Gemini 1.5 Flash	53.9	33.0	41.4	42.7
Gemini 1.5 Pro	78.9	43.9	46.9	56.4
Llava 1.6 13B	70.2	43.2	17.9	43.7
Llava 1.6 34B	93.6	37.2	27.5	52.8
GPT-4o	67.9	39.6	44.4	50.6
InternVL 2 8B	62.2	44.6	45.1	50.6
InternVL 2 26B	59.3	44.9	35.5	46.6
VideoLlama 2 7B	70.1	45.6	37.7	51.1
BIP-ALM	41.2	34.1	30.6	33.9
LIMP	93.4	67.7	68.7	76.6

Table 1: Human and model performance for different question types as well as for all questions.

Figure 4 illustrates how IMP evaluates the action and utterance likelihood at one time step. Given the condition, the LLM estimates that it is likely that agent i will take the observed action (“walk towards table”) but is unlikely to say “I found a potato inside fridge” as it is inconsistent with the social goal of hindering agent j (agent i had just put a potato in the fridge before the conversation).

6 Experiments

6.1 Human Experiment

We recruited 18 participants (mean age = 36.0; 10 female) from Prolific to answer 90 questions randomly sampled from the benchmark. Each question received responses from 3 participants. The experiment was approved by an institutional review board.

6.2 Baselines

We evaluated our benchmark on state-of-the-art LMMs. For models capable of processing video input, the entire video was provided. For models without video input capabilities, we uniformly sample one frame every 20 frames from the video episode as input. We evaluated **GPT-4o** (OpenAI, 2023), **Llava 1.6** (Liu et al., 2023), **Gemini 1.5** (Reid et al., 2024), **InternVL2** (Chen et al., 2023) and **VideoLlama 2** (Cheng et al., 2024). We evaluated the latest version of each LMM at the time of submission. For **LIMP**, we use Gemini 1.5 Pro as the VLM and GPT-4o as the LLM. Finally, we evaluated **BIP-ALM** with finetuned Llama 2 (Jin et al., 2024), the best-performing model on a prior multi-modal ToM benchmark, MMToM-QA.

6.3 Results

We report the human and model performance in Table 3. Human participants achieved almost perfect accuracy across all questions, with 98.9% of the correct answers having majority agreement. The overall performance averaged across individual participants is 93.5%.

All LMM baselines performed poorly on MuMA-ToM, indicating a substantial gap between machine and human ToM. The best-performing LMM baseline is Gemini 1.5 Pro, but its overall accuracy is only 56.4%. Among the three question types, belief inference is the easiest for LMMs. In particular, Llava 34B achieved the highest accuracy for belief inference. However, all LMMs struggle with the more challenging social goal inference and belief of goal inference questions, which are not evaluated by prior ToM benchmarks. Notably, BIP-ALM had an accuracy of 33.9%, indicating its inability to understand multi-agent interactions. This is because BIP-ALM only relies on single-agent goals and beliefs for inverse planning but does not consider social goals and beliefs of others’ goals. Additionally, BIP-ALM assumes certain symbolic representations, which also limits its generality. Our LIMP model significantly outperforms all state-of-the-art models on our benchmark, with an overall accuracy of 76.6%. Critically, by using GPT-4o to estimate action and utterance likelihood for inverse multi-agent planning, LIMP can achieve much better performance than GPT-4o itself. There is still a gap between the best model performance and human performance, highlighting the need for further studies. Section 10.1 in Appendix shows more experimental results and discussions.

7 Discussion

Why do LMMs perform poorly? There are two sources of systematic errors for LMMs. First, LMMs struggle with understanding multi-modal behavior in complex social situations, often failing to distinguish between deliberate hindering and failed attempts to help due to incorrect beliefs. Most models can solve belief inference tasks where helping is the assumed social goal. However, they consistently struggle in scenarios where hindering is the assumed social goal. Except for Llava 1.6 34B and Gemini 1.5 Pro, all models perform at or below the level of random guessing in these cases. The failure to understand adversarial behaviors is even more prominent in social goal inference and belief of goal inference. For instance, if Sarah tells John there is milk in the fridge, but there isn't any milk, and Sarah knows this, she is hindering John. Second, LMMs often fail to correctly interpret visual inputs, such as when an object is too small or is occluded, leading to incorrect conclusions about the agent's actions. While humans are able to use contextual clues to infer what the object might be, LLMs struggle with this task. These errors in recognizing crucial actions likely contribute significantly to their overall poor performance on our benchmark.

Why does LIMP outperform the best LMMs? LIMP overcomes the two aforementioned weaknesses of LMMs – the inability to recognize multi-modal behavior under various social goals and the sensitivity to noisy visual perception. First, while LLMs struggle with direct ToM reasoning, they excel at the forward generation of multi-modal behavior given mental states. For example, it is much harder for an LLM to correctly infer whether an agent is hindering another by lying than it is for the model to generate a lie based on the agent's belief and social goal. Such multi-modal behavior generation ability enables LIMP to estimate the action and utterance likelihood, identifying the key actions and/or utterances that reveal the true mental state of an agent (as shown in Figure 4). Second, when a VLM fails to recognize the exact object that the agent is interacting with, LIMP can fill in this missing information with context from the text input. We evaluated the action accuracy by using semantic similarity and found that this approach increases inferred action accuracy from 54.4% to 86.6%. As a result, LIMP is able to perform inference on much more accurate information.

How general is LIMP? LIMP can represent all information using natural language, which enables the direct use of any pretrained LLMs and VLMs without domain-specific knowledge or finetuning. By utilizing powerful pretrained VLMs for visual perception, LIMP can directly recognize actions from RGB videos in an open-ended way, without specifying target action labels for a domain. LIMP also leverages an LLM to use contextual clues from the text to fill in missing information from visual perception, providing a general method for multi-modal information fusion. One can also easily upgrade LIMP by plugging in *any* state-of-the-art VLMs and LLMs.

What are the limitations of LIMP? Hallucinations created by the VLM can cause significant errors in LIMP. For example, VLM may mistakenly think that the agent also grabs something from the fridge when they just opens and closes it. Such hallucinations in action recognition can not be corrected by the textural context, causing LIMP to incorrectly interpret the agent's behavior. Additionally, LIMP does not explicitly infer an agent's belief of another agent's belief. It instead prompts an LLM with past actions and utterances to implicitly infer that, which can become costly for longer events. LIMP also does not perform recursive reasoning for more than two levels.

What are the limitations of our benchmark? The scenarios in our benchmark are currently limited to interactions between two agents in household settings, where there are three social goals: helping, hindering, and acting independently. Moreover, the current benchmark has synthetic human activities. These synthetic activities are realistic as verified in prior work (Puig et al., 2020) and enable large-scale testing. However sim-to-real evaluation could be valuable for future studies.

8 Conclusion

We present the first multi-modal Theory of Mind benchmark for multi-agent interactions in complex embodied settings. We have systematically evaluated humans and state-of-the-art LMMs on our benchmark. We have also proposed a novel multi-modal ToM model that outperforms all baselines while maintaining generality. In future work, we intend to incorporate more complex real-world scenarios beyond household environments and introduce multi-modal social interactions involving more than two agents. We also plan to create a test set with real-world videos for ToM evaluation in real-world scenarios.

References

- Maryam Amirizani, Elias Martin, Maryna Sivachenko, Afra Mashhadi, and Chirag Shah. Do llms exhibit human-like reasoning? evaluating theory of mind in llms for open-ended responses. *arXiv preprint arXiv:2406.05659*, 2024.
- Chris L Baker, Julian Jara-Ettinger, Rebecca Saxe, and Joshua B Tenenbaum. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4): 1–10, 2017.
- Simon Baron-Cohen, Alan M Leslie, and Uta Frith. Does the autistic child have a “theory of mind”? *Cognition*, 21(1):37–46, 1985.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- Chunkit Chan, Cheng Jiayang, Yauwai Yim, Zheyang Deng, Wei Fan, Haoran Li, Xin Liu, Hongming Zhang, Weiqi Wang, and Yangqiu Song. Negotiationtom: A benchmark for stress-testing machine theory of mind on negotiation surrounding. *arXiv preprint arXiv:2404.13627*, 2024.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, and Minlie Huang. Tombench: Benchmarking theory of mind in large language models, 2024a. URL <https://arxiv.org/abs/2402.15052>.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, et al. Tombench: Benchmarking theory of mind in large language models. *arXiv preprint arXiv:2402.15052*, 2024b.
- Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. URL <https://arxiv.org/abs/2406.07476>.
- Michael Cohen. Exploring roberta’s theory of mind through textual entailment. 2021.
- Susanne A. Denham, Kimberly A. Blair, Elizabeth DeMulder, Jennifer Levitas, Katherine Sawyer, Sharon Auerbach-Major, and Patrick Queenan. Preschool emotional competence: Pathway to social competence? *Child Development*, 74(1):238–256, 2003. doi: 10.1111/1467-8624.00533. URL <https://doi.org/10.1111/1467-8624.00533>.
- Kanishk Gandhi, Gala Stojnic, Brenden M Lake, and Moira R Dillon. Baby intuitions benchmark (bib): Discerning the goals, preferences, and actions of others. *Advances in Neural Information Processing Systems*, 34:9963–9976, 2021.
- Kanishk Gandhi, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah Goodman. Understanding social reasoning in language models with language models. *Advances in Neural Information Processing Systems*, 36, 2024.
- P. J. Gmytrasiewicz and P. Doshi. A framework for sequential planning in multi-agent settings. *Journal of Artificial Intelligence Research*, 24:49–79, July 2005. ISSN 1076-9757. doi: 10.1613/jair.1579. URL <http://dx.doi.org/10.1613/jair.1579>.
- Andrew Gordon. Commonsense interpretation of triangle behavior. In *Proceedings of the aaai conference on artificial intelligence*, volume 30, 2016a.
- Andrew S. Gordon. Commonsense interpretation of triangle behavior. In *AAAI Conference on Artificial Intelligence*, 2016b.

- J Kiley Hamlin, Karen Wynn, and Paul Bloom. Social evaluation by preverbal infants. *Nature*, 450 (7169):557–559, 2007.
- Yanlin Han and Piotr Gmytrasiewicz. Ipomdp-net: A deep neural network for partially observable multi-agent planning using interactive pomdps. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6062–6069, Jul. 2019. doi: 10.1609/aaai.v33i01.33016062. URL <https://ojs.aaai.org/index.php/AAAI/article/view/4562>.
- Yinghui He, Yufan Wu, Yilin Jia, Rada Mihalcea, Yulong Chen, and Naihao Deng. Hi-tom: A benchmark for evaluating higher-order theory of mind reasoning in large language models. *arXiv preprint arXiv:2310.16755*, 2023.
- Guiyang Hou, Wenqi Zhang, Yongliang Shen, Linjuan Wu, and Weiming Lu. Timetom: Temporal space is the key to unlocking the door of large language models’ theory-of-mind. *arXiv preprint arXiv:2407.01455*, 2024.
- Anna A. Ivanova, Aalok Sathe, Benjamin Lipkin, Unnathi Kumar, Setayesh Radkani, Thomas H. Clark, Carina Kauf, Jennifer Hu, R. T. Pramod, Gabriel Grand, Vivian Paulun, Maria Ryskina, Ekin Akyürek, Ethan Wilcox, Nafisa Rashid, Leshem Choshen, Roger Levy, Evelina Fedorenko, Joshua Tenenbaum, and Jacob Andreas. Elements of world knowledge (ewok): A cognition-inspired framework for evaluating basic world knowledge in language models, 2024. URL <https://arxiv.org/abs/2405.09605>.
- Kunal Jha, Tuan Anh Le, Chuanyang Jin, Yen-Ling Kuo, Joshua B Tenenbaum, and Tianmin Shu. Neural amortized inference for nested multi-agent reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 530–537, 2024.
- Chuanyang Jin, Yutong Wu, Jing Cao, Jiannan Xiang, Yen-Ling Kuo, Zhiting Hu, Tomer Ullman, Antonio Torralba, Joshua B Tenenbaum, and Tianmin Shu. Mmtom-qa: Multimodal theory of mind question answering. *arXiv preprint arXiv:2401.08743*, 2024.
- Hyunwoo Kim, Melanie Sclar, Xuhui Zhou, Ronan Le Bras, Gunhee Kim, Yejin Choi, and Maarten Sap. Fantom: A benchmark for stress-testing machine theory of mind in interactions. *arXiv preprint arXiv:2310.15421*, 2023.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *ArXiv*, abs/2205.11916, 2022. URL <https://api.semanticscholar.org/CorpusID:249017743>.
- Michal Kosinski. Theory of mind may have spontaneously emerged in large language models. *arXiv preprint arXiv:2302.02083*, 2023.
- Matthew Le, Y-Lan Boureau, and Maximilian Nickel. Revisiting the evaluation of theory of mind through question answering. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5872–5877, 2019.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023a.
- Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, Lingpeng Kong, and Qi Liu. M³it: A large-scale dataset towards multi-modal multilingual instruction tuning, 2023b. URL <https://arxiv.org/abs/2306.04387>.
- Wenjie Li, Shannon C. Yasuda, Moira R. Dillon, and Brenden M. Lake. An infant-cognition inspired machine benchmark for identifying agency, affiliation, belief, and intention. 2024a. URL <https://api.semanticscholar.org/CorpusID:268065789>.
- Wenjie Li, Shannon C Yasuda, Moira Rose Dillon, and Brenden Lake. An infant-cognition inspired machine benchmark for identifying agency, affiliation, belief, and intention. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 46, 2024b.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023.

- Manasi Malik and Leyla Isik. Relational visual representations underlie human social interaction recognition. *Nature Communications*, 14(1):7317, 2023.
- Aviv Netanyahu, Tianmin Shu, Boris Katz, Andrei Barbu, and Joshua B. Tenenbaum. Phase: Physically-grounded abstract social events for machine social perception. In *35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021a.
- Aviv Netanyahu, Tianmin Shu, Boris Katz, Andrei Barbu, and Joshua B Tenenbaum. Phase: Physically-grounded abstract social events for machine social perception. In *Proceedings of the aaai conference on artificial intelligence*, volume 35, pages 845–853, 2021b.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- Pranshu Pandya, Agney S Talwarr, Vatsal Gupta, Tushar Kataria, Vivek Gupta, and Dan Roth. Ntsebench: Cognitive reasoning benchmark for vision language models, 2024. URL <https://arxiv.org/abs/2407.10380>.
- Viorica Patraucean, Lucas Smaira, Ankush Gupta, Adria Recasens, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Mateusz Malinowski, Yi Yang, Carl Doersch, et al. Perception test: A diagnostic benchmark for multimodal video models. *Advances in Neural Information Processing Systems*, 36, 2024.
- Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8494–8502, 2018a.
- Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. Virtualhome: Simulating household activities via programs, 2018b. URL <https://arxiv.org/abs/1806.07011>.
- Xavier Puig, Tianmin Shu, Shuang Li, Zilin Wang, Joshua B. Tenenbaum, Sanja Fidler, and Antonio Torralba. Watch-and-help: A challenge for social perception and human-ai collaboration, 2020.
- Xavier Puig, Tianmin Shu, Joshua B Tenenbaum, and Antonio Torralba. Nopa: Neurally-guided online probabilistic assistance for building socially intelligent home assistants. *arXiv preprint arXiv:2301.05223*, 2023.
- Neil C. Rabinowitz, Frank Perbet, H. Francis Song, Chiyuan Zhang, S. M. Ali Eslami, and Matthew Botvinick. Machine theory of mind, 2018. URL <https://arxiv.org/abs/1802.07740>.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- Sahand Sabour, Siyang Liu, Zheyuan Zhang, June M. Liu, Jinfeng Zhou, Alvionna S. Sunaryo, Juanzi Li, Tatia M. C. Lee, Rada Mihalcea, and Minlie Huang. Emobench: Evaluating the emotional intelligence of large language models, 2024. URL <https://arxiv.org/abs/2402.12071>.
- Kate Sanders, David Etter, Reno Kriz, and Benjamin Van Durme. Multivent: Multilingual videos of events with aligned natural text. *arXiv preprint arXiv:2307.03153*, 2023.
- Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi, and Yulia Tsvetkov. Minding language models’ (lack of) theory of mind: A plug-and-play multi-character belief tracker. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13960–13980, Toronto, Canada, July 2023a. Association for Computational Linguistics. URL <https://aclanthology.org/2023.acl-long.780>.
- Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi, and Yulia Tsvetkov. Minding language models’ (lack of) theory of mind: A plug-and-play multi-character belief tracker, 2023b.
- Melanie Sclar, Sachin Kumar, Peter West, Alane Suhr, Yejin Choi, and Yulia Tsvetkov. Minding language models’ (lack of) theory of mind: A plug-and-play multi-character belief tracker, 2023c. URL <https://arxiv.org/abs/2306.00924>.

- Tianmin Shu, Marta Kryven, Tomer D Ullman, and Josh Tenenbaum. Adventures in flatland: Perceiving social interactions under physical dynamics. In *CogSci*, 2020.
- Tianmin Shu, Abhishek Bhandwadar, Chuang Gan, Kevin Smith, Shari Liu, Dan Gutfreund, Elizabeth Spelke, Joshua Tenenbaum, and Tomer Ullman. Agent: A benchmark for core psychological reasoning. In *International Conference on Machine Learning*, pages 9614–9625. PMLR, 2021.
- Adil Soubki, John Murzaku, Arash Yousefi Jordehi, Peter Zeng, Magdalena Markowska, Seyed Abolghasem Mirroshandel, and Owen Rambow. Views are my own, but also yours: Benchmarking theory of mind using common ground. *arXiv preprint arXiv:2403.02451*, 2024.
- Stephanie Stacy, Siyi Gong, Aishni Parab, Minglu Zhao, Kaiwen Jiang, and Tao Gao. A bayesian theory of mind approach to modeling cooperation and communication. *Wiley Interdisciplinary Reviews: Computational Statistics*, 16(1):e1631, 2024.
- Jingqun Tang, Qi Liu, Yongjie Ye, Jinghui Lu, Shu Wei, Chunhui Lin, Wanqing Li, Mohamad Fitri Faiz Bin Mahmood, Hao Feng, Zhen Zhao, et al. Mtvqa: Benchmarking multilingual text-centric visual question answering. *arXiv preprint arXiv:2405.11985*, 2024.
- Ravi Tejwani, Yen-Ling Kuo, Tianmin Shu, Boris Katz, and Andrei Barbu. Social interactions as recursive mdps. In *Conference on Robot Learning*, pages 949–958. PMLR, 2021.
- Tomer Ullman. Large language models fail on trivial alterations to theory-of-mind tasks, 2023. URL <https://arxiv.org/abs/2302.08399>.
- Tomer Ullman, Chris Baker, Owen Macindoe, Owain Evans, Noah Goodman, and Joshua Tenenbaum. Help or hinder: Bayesian models of social goal inference. *Advances in neural information processing systems*, 22, 2009.
- Mudit Verma, Siddhant Bhambri, and Subbarao Kambhampati. Theory of mind abilities of large language models in human-robot interaction: An illusion? In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, pages 36–45, 2024.
- H M. Wellman, D Cross, and J Watson. Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3):655–684, 2001. doi: 10.1111/1467-8624.00304. URL <https://doi.org/10.1111/1467-8624.00304>.
- Alex Wilf, Sihyun Shawn Lee, Paul Pu Liang, and Louis-Philippe Morency. Think twice: Perspective-taking improves large language models’ theory-of-mind capabilities, 2023. URL <https://arxiv.org/abs/2311.10227>.
- Yufan Wu, Yinghui He, Yilin Jia, Rada Mihalcea, Yulong Chen, and Naihao Deng. Hi-ToM: A benchmark for evaluating higher-order theory of mind reasoning in large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10691–10706, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.717. URL <https://aclanthology.org/2023.findings-emnlp.717>.
- 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. *arXiv preprint arXiv:2402.06044*, 2024.
- Shuai Yang, Yuying Ge, Yang Li, Yukang Chen, Yixiao Ge, Ying Shan, and Yingcong Chen. Seed-story: Multimodal long story generation with large language model. *arXiv preprint arXiv:2407.08683*, 2024. URL <https://arxiv.org/abs/2407.08683>.
- Kaining Ying, Fanqing Meng, Jin Wang, Zhiqian Li, Han Lin, Yue Yang, Hao Zhang, Wenbo Zhang, Yuqi Lin, Shuo Liu, et al. Mmt-bench: A comprehensive multimodal benchmark for evaluating large vision-language models towards multitask agi. *arXiv preprint arXiv:2404.16006*, 2024a.
- Lance Ying, Kunal Jha, Shivam Arya, Joshua B. Tenenbaum, Antonio Torralba, and Tianmin Shu. Goma: Proactive embodied cooperative communication via goal-oriented mental alignment, 2024b. URL <https://arxiv.org/abs/2403.11075>.

Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. URL <https://arxiv.org/abs/2306.02858>.

Jieyu Zhang, Weikai Huang, Zixian Ma, Oscar Michel, Dong He, Tanmay Gupta, Wei-Chiu Ma, Ali Farhadi, Aniruddha Kembhavi, and Ranjay Krishna. Task me anything. *arXiv preprint arXiv:2406.11775*, 2024.

Tan Zhi-Xuan, Jordyn Mann, Tom Silver, Josh Tenenbaum, and Vikash Mansinghka. Online bayesian goal inference for boundedly rational planning agents. *Advances in neural information processing systems*, 33:19238–19250, 2020.