

TactfulToM: Do LLMs Have the Theory of Mind Ability to Understand White Lies?

Yiwei Liu^{♡*} Emma Jane Pretty^{♠†} Jiahao Huang^{◇†} Saku Sugawara[♣]

[♡] EPFL, Lausanne, Switzerland [♠] Tampere University
[◇] University of Tokyo [♣] National Institute of Informatics
yiw.liu@epfl.ch saku@nii.ac.jp

Abstract

While recent studies explore Large Language Models’ (LLMs) performance on Theory of Mind (ToM) reasoning tasks, research on ToM abilities that require more nuanced social context is limited, such as white lies. We introduce TactfulToM, a novel English benchmark designed to evaluate LLMs’ ability to understand white lies within real-life conversations and reason about prosocial motivations behind them, particularly when they are used to spare others’ feelings and maintain social harmony. Our benchmark is generated through a multi-stage human-in-the-loop pipeline where LLMs expand manually designed seed stories into conversations to maintain the information asymmetry between participants necessary for authentic white lies. We show that TactfulToM is challenging for state-of-the-art models, which perform substantially below humans, revealing shortcomings in their ability to fully comprehend the ToM reasoning that enables true understanding of white lies.¹

1 Introduction

Theory of Mind (ToM) is the cognitive ability to impute mental states to oneself and others, and to use these inferred mental representations to predict and explain behaviors (Premack and Woodruff, 1978; Baron-Cohen et al., 1985). This ability is recognized as a foundation for effective social interactions and a pillar of common sense reasoning (Lake et al., 2017), which is crucial for developing human-level AI systems. Modern LLMs like GPT (Hurst et al., 2024) and DeepSeek (DeepSeek-AI, 2025) have demonstrated remarkable reasoning capabilities in structured domains such as mathematics and programming, yet research consistently reveals significant gaps between humans and LLMs in ToM

*Work conducted during an internship at National Institute of Informatics.

†Contributed equally as second authors.

¹<https://github.com/nii-cl/tactful-tom>

The figure illustrates a conversation scenario from the TactfulToM benchmark. It is divided into four horizontal panels. The top panel shows David saying, "I think Lily could really excel with some encouragement, especially when it comes to continuing her pursuit of music." Ben replies, "I agree. It can be disheartening to any budding musician when they struggle in the early stages." A side note indicates Lily and Adrian are unaware of this. The second panel shows David and Ben saying, "I notice she struggles with the theory part sometimes." Adrian replies, "Yes, it's a bit tricky for her." A side note indicates Lily is unaware of this. The third panel shows David saying, "I've been dreaming about being a musician, and I love the idea of creating my own music one day." Ben replies, "That's wonderful! You're really talented at music! You should keep that passion alive." Lily replies, "So, David and Ben tell a white lie to support Lily emotionally." The bottom panel contains four questions: Q1 (Comprehension): "Was it true what David told Lily that she is very talented at music?"; Q2 (Justification): "Why did David tell Lily that she is very talented at music?"; Q3 (Lie Detectability): "List all the people that know the precise correct answer to Q1."; Lie Ability (Answer): "Answer: David is able to tell Lily the lie because he holds the correct second order belief."

Figure 1: Example from a question set in TactfulToM.

tasks, especially when applied to realistic social scenarios (Chen et al., 2024; Gu et al., 2024).

Among the various sub-abilities of ToM, understanding white lies, intentional falsehoods told specifically to protect others’ feelings, represents a particularly complex aspect that combines belief tracking with emotional sensitivity (Beaudoin et al., 2020; Abdollahi et al., 2022). The ability to detect white lies and understand their emotional motivations becomes essential for developing safe and appropriately responsive AI tools, especially as LLM tools are increasingly deployed in domains requiring emotional intelligence, such as educational tutoring, medical consultation, and caregiving. Despite this importance, white lies remain

largely understudied. ToM battery evaluations like ToMBench (Chen et al., 2024) have included white lie tasks but with limitations, only containing 20 white lie samples without dialogue interaction. Testing on such small samples is insufficient for reliable evaluation, as minor variations in test cases can significantly alter results (Ullman, 2023). Additionally, using established psychological ToM tests risks data contamination that could artificially inflate performance metrics (Shapira et al., 2023a; Chen et al., 2024). This creates a critical research gap in understanding LLMs’ white lie comprehension capabilities despite the significance for AI systems to safely operate in nuanced contexts.

To address this challenge, we introduce TactfulToM, an English benchmark that aims to evaluate LLMs’ ability to understand and reason about white lies in real-world conversational contexts, particularly focusing on the interplay between deceptive statements and their underlying motivations. Our benchmark offers four key contributions: (1) a novel decomposition framework that breaks down white lies into triplets and role-based information asymmetry, enabling manually crafted seed stories; (2) high-quality conversations generated via human-in-the-loop generation pipeline (avoiding biases from direct LLM generation) with strict validation; (3) a comprehensive evaluation framework to test models’ understanding of white lies by combining mental state tracking questions with both established measures from Strange Stories (Happé, 1994) and our newly designed question types; and (4) a diverse dataset of 100 multi-party conversations spanning across different white lie classes, types, and (falsifiability) difficulty levels, which contains 6.7K questions across multiple answer formats.

We evaluate TactfulToM on nine recent LLMs from four different families, including both vanilla and reasoning models. Through our experiments, we uncover gaps between human and AI performance in white lie comprehension. The analysis of evaluation results on TactfulToM reveals several interesting findings: (1) all tested LLMs significantly underperform humans, even the best-performing ones (DeepSeek families and GPT-4o); (2) Chain-of-Thought (CoT) prompting and specialized reasoning models show inconsistent improvements, with some models even performing worse than vanilla models from the same families; (3) LLMs struggle with true white lie understanding and fail to grasp the genuine motivations behind white lies;

and (4) LLMs can track mental states but fail to apply them effectively in white lie contexts.

Our contributions are summarized as follows:

- We present a benchmark that tests LLMs’ ability to understand white lies in social contexts, filling a research gap in ToM evaluation.
- Our dataset covers five white lie classes, two types, and three levels, all constructed efficiently using a human-in-the-loop process.
- Our analysis reveals limitations in the white lie reasoning capabilities of recent LLMs, providing insights for future model development.

2 TactfulToM Design

Building upon the white lie test from Strange Stories (Happé, 1994) and previous successful evaluations of LLMs’ ToM ability (Kim et al., 2023), we developed a dataset of social conversations capturing common white lies in daily life. This section outlines our design considerations and approach (as shown in Figure 2): (1) theoretical requirements informing our design; (2) methodology for structuring white lies with triple and role-based information asymmetric; and (3) evaluation framework for white lie understanding and reasoning.

2.1 Theoretic Requirements from ToM Task Designing

ToM evaluation requires carefully structured scenarios that test a model’s ability to accurately attribute mental states. Three critical aspects based on Quesque and Rossetti (2020); Kim et al. (2023) are identified: Non-merging Mental States, Non-mentalising, and Elimination of Visual Indicators.

Non-merging Mental States A valid evaluation of ToM requires the model to distinguish between its own knowledge and the beliefs of others. In scenarios where one character provides false information and others either believe the lie or know the truth, the model must infer what a deceived character believes only based on the information available to them, not based on the model’s knowledge. To ensure the non-merging requirement, scenarios must involve multiparty conversations where it is explicitly revealed who knows the truth and the lie. This allows for controlled belief divergence, ensuring that the model must track the different perspectives of each character rather than assuming all characters share the same understanding. We

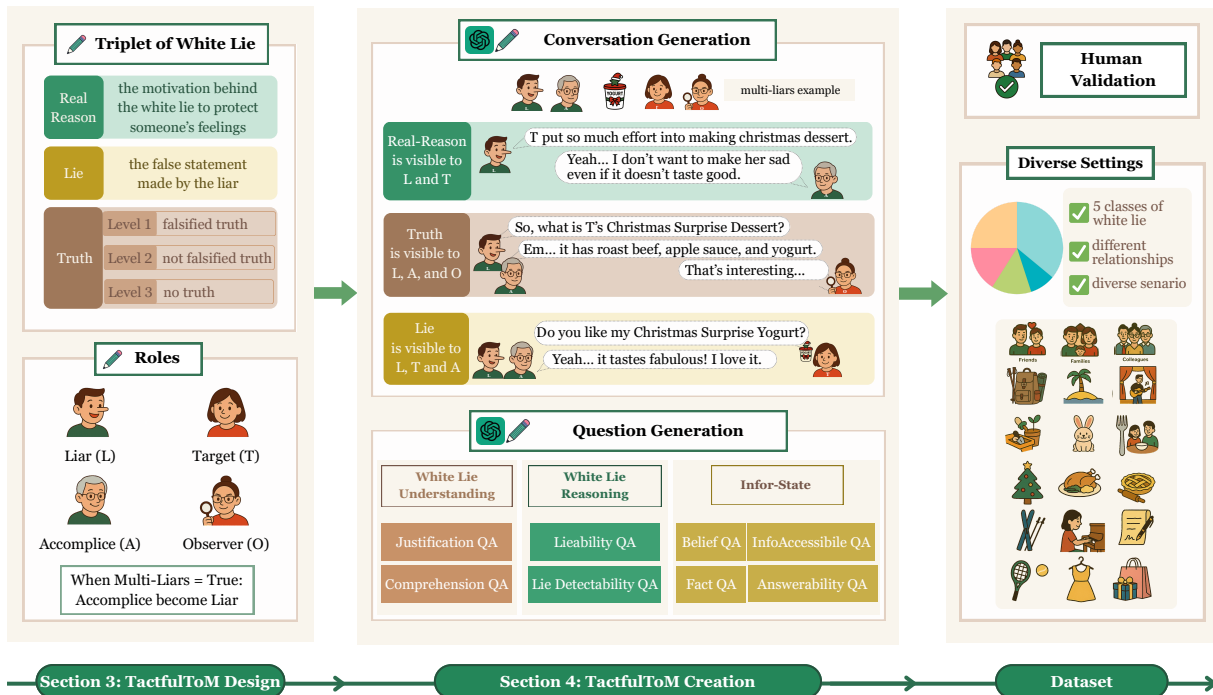


Figure 2: Overview of the dataset construction pipeline for TactfulToM.

design our benchmark with information asymmetry to enforce this differentiation.

Non-mentalising It is crucial not to attribute model success to genuine mentalizing when simpler processes can explain the outcome. In white lie scenarios, if a model's correct answer arises from surface-level patterns or word correlations, this explanation should be prioritized over more complex reasoning about mental states. For example, the model might correctly identify that a character believes a lie, but if this answer is based on word associations rather than true mental state reasoning, it suggests the model is not engaging with the belief system of the character. To address this, we introduce distractor answers with high word correlation to test if the model is relying on deeper reasoning rather than simple associations.

Elimination of Visual Indicators The model should also not rely on descriptions of body language, emotions, or visual indicators when inferring belief states, only linguistic contexts (Premack and Woodruff, 1978; Baron-Cohen et al., 1985). Relying on such cues would lead to shortcuts that allow the model to infer beliefs based on visible indicators, not through genuine reasoning about what another person might believe. Thus, our benchmark contains conversational exchanges without any narrative descriptions, requiring the model to infer

mental states purely from the dialogue, ensuring that belief inference is based on logical reasoning rather than perceptual cues.

2.2 Structuring White Lies

White Lie Triplet Decomposition To systematically create our dataset, we first decompose white lies into three elements: (1) **Real Reason**: the motivation behind telling the lie; (2) **Lie**: the false statement made by the liar; (3) **Truth**: the actual truth that diverges from the lie. For example, in a classic Strange Story test (Happé, 1994), the truth is "Helen wanted a rabbit but received encyclopedias from her parents" while Helen lies "It's lovely, thank you. It's just what I wanted." with the real-reason being to avoid hurting Helen's parents' feelings after they gave her a gift they thought she would like.

Real-Reason Correspondence to White Lie Types White lies fall into two distinct types based on their underlying motivations: altruistic white lies and Pareto white lies (Erat and Gneezy, 2012). Altruistic white lies are told purely for the benefit of others, where the liar may incur some personal cost or disadvantage. In contrast, Pareto white lies create a mutually beneficial outcome, serving both the interests of the person being lied to and the liar themselves. The fundamental categorization guided our design of two types of real-reason state-

ments corresponding to these two categories of lies.

Three Levels of Truth Accessibility in White Lies

To reflect real-world complexities, we incorporate three difficulty levels by varying falsifiability (between “lie” and “truth”) in our white lie triplets. After establishing the “real reason” (e.g., “declining an invitation without hurting T’s feelings”) and “lie” (e.g., “L has a reservation tonight”), we determine how the truth is presented. We structure conversations into three difficulty levels: (1) Level-1: falsifiable truth provided, e.g., “L does not have a reservation tonight”; (2) Level-2: non-falsifiable truth provided, e.g., “L hasn’t decided what to do tonight”; and (3) Level-3: no truth provided. Not all white lie scenarios can reasonably accommodate all three levels; some contexts intrinsically require truth disclosure while others cannot reasonably support ambiguous truth construction. As such, we selectively designed appropriate levels for each white lie triplet. This creates progressive reasoning challenges: with the truth provided, models can identify lies before determining the motivation; without it, models must infer the deceptive nature directly from the real reason, thus compelling models to rely on conversational context and reasoning rather than simply detecting information contradictions to reverse-engineer the scenarios.

Role-based Information Asymmetry Building upon the inherent characteristics of white lie scenarios, we define four roles based on their access to the white lie triplet: the **Liar (L)**, who has complete understanding and knowledge of the white lie; the **Accomplice (A)**, who has access to all elements in triplet; the **Observer (O)**, who only knows the truth; and the **Target (T)**, who only receives the lie. This asymmetric access leads to varying degrees of white lie comprehension among participants; it is not only a necessary condition for white lies to exist, but also aligns with the non-merging mental states requirement (Section 2.1). Our dataset incorporates diverse character relationships (friends, families, and colleagues) and complex patterns including multi-liar scenarios where accomplices function as additional liars given their equivalent information access. This design captures more diverse and realistic interaction patterns. We also impose a crucial constraint: all discussions about the white lie triplet begin within the conversation scenario, with no prior exchange of this information among characters.

2.3 Hierarchical Evaluation Framework: Mental States to White Lie Reasoning

Our evaluation framework employs a progressive three-tier question hierarchy: (A) **Info-State Questions** assesses basic mental state tracking, (B) **First-Order: White Lie Understanding** evaluates how models perceive and interpret white lies, and (C) **Second-Order: White Lie Reasoning** tests the models’ ability to reason about different roles’ perspectives on the white lie within the conversation.

Info-State Questions We include four question types targeting belief attribution: first, we establish **Fact** questions (factQ) that include factual question-answer pairs about the asymmetrical information “real reason” and “truth”. Building on these, we develop **Belief** questions that assess first-order beliefs (what characters believe) and second-order beliefs (how characters understand others’ beliefs: “What does X believe about Y’s understanding of [FactQ]?”). We also include **Info Accessibility** questions (“List all characters who know [real reason/truth]”) and **Answerability** questions (e.g., “List all characters who can answer: [FactQ]”). This question-type structure prevents inflated scores from the “illusion of ToM” (Kim et al., 2023) while enabling more accurate assessment of mental state tracking capabilities.

White Lie Understanding (1st-Order) Drawing from the Strange Story test, we assess basic white lie understanding through two question types: comprehension and justification. **Comprehension** questions (“Is the statement X told Y true?”) evaluate whether models can identify false statements as lies. **Justification** questions (“Why did X say that to Y?”) probe whether models recognize the prosocial motivations behind white lies that distinguish a white lie from a simple deception. These questions are complementary; even if a model correctly identifies a statement as false, it must also understand the protective intention to fully comprehend the white lie concept. According to Happé (1994), accurate responses to both questions indicate second-order ToM ability.

White Lie Reasoning (2nd-Order) We introduce two novel question types that evaluate models’ understanding of characters’ perspectives: **Lie Ability** questions if models can identify which characters possess the necessary conditions to tell a white lie (requiring understanding Liar’s second-order beliefs). **Lie Detectability** questions evaluate

if models can determine which characters have sufficient information to recognize deception. Both require reasoning about characters' information access and resulting beliefs, providing a more stringent test of genuine second-order ToM reasoning beyond simple pattern matching.

Comprehensive Evaluation Format To ensure robust evaluation, we present each question in two formats. Multiple-choice questions (MCQs) are complemented with free-form responses to assess genuine understanding in addition to choice selection, as providing choices inherently guides model reasoning paths. Similarly, list-type questions are presented in both open-ended and binary formats. This comprehensive approach allows us to check cases where providing answer options may guide the model to reason.

Design Consideration for Evaluation of Emotion-based Real Reasons When the real reason involves protecting the target from emotional distress unrelated to the lie itself (e.g., "Lynn will be very sad if she knows her grandma passed away"), our evaluation framework requires adaptation. Unlike factual information that can be explicitly discussed within conversations, emotional reactions (such as sadness) are inherent to the target and cannot be treated as discrete information. We exclude Information Accessibility and Lie Ability questions for such cases, which applies specifically to our common sense (emotional soothing) (Class 2) category (Section 3.1), ensuring framework consistency while accommodating emotion-based white lies.

3 TactfulToM Creation

The construction of TactfulToM consists of the following steps (as shown in Figure 2): (1) manually creating seed stories, and then expanding them into natural conversations through a human-in-the-loop process; (2) generating question-answer pairs through templates; and (3) strict quality control.

3.1 Conversation Generation

The construction of TactfulToM consists of the following steps: (1) manually create seed stories, and then (2) expand them into natural conversations through a human-in-the-loop process.

Seed Stories To create dataset diversity, we collected examples from interviews, social media, and online sources documenting white lie scenarios in

daily life. We gathered examples in the format of white lie triplets to systematically capture the essential components of each scenario. We then categorized them into five distinct classes based on different motivations behind white lies: **social evasion** (Class 0), **common sense (imagination preservation)** (Class 1), **common sense (emotional soothing)** (Class 2), **confidence enhancement** (Class 3), and **mistake hiding** (Class 4). This categorization represents both altruistic white lies (Classes 1, 2, and 3) and Pareto white lies (Classes 0 and 4), ensuring comprehensive coverage of realistic social interactions. We constructed 100 seed stories, the data distribution across these categories is provided in Appendix A.1.

Generation Pipeline To facilitate conversation generation, we designed a set of scenario elements and combined them with seed stories as input for our 4-step generation prompt template, in which each step generates one element of the white lie triplet sequentially (conversation generation template is provided in Appendix A.2). The generation process employed GPT-4o (Hurst et al., 2024) in a human-in-the-loop methodology. Following FANToM (Kim et al., 2023), we assigned character names sampled from the top 20% most popular US birth names (Burnsworth, 2022) and simulated scenarios where characters temporarily leave the conversation, allowing remaining participants to discuss without the absent participants' knowledge. We expanded the leaving reasons from Kim et al. (2023) into a more comprehensive list (samples provided in Appendix A.3).

This stepwise generation pipeline and controlled asymmetry serve two primary purposes. First, it enables the model to generate white lie conversations even when its understanding of white lies is limited, by preventing the model from developing its own interpretation of white lies, thus avoiding conversations generated based on the model's potentially flawed understanding and reducing generation bias (see Appendix A.5 for detailed explanation of how we address GPT-4o's white lie understanding limitations). Second, this approach enables controlled information asymmetry by systematically managing each participant's involvement at different conversation segments, where multiple dialogue segments were generated and carefully selected to ensure each segment contained only the intended information and roles before being combined into complete conversations.

3.2 Question-Answer Pair Generation

We developed a systematic templated generation approach for all question types (introduced in Section 2.3), where templates are populated with white lie triplet elements and role information, enabling efficient question generation. All templates and examples are provided in Appendix A.6. Additionally, we systematically generated wrong options for MCQs to ensure each question has one correct answer and several high-quality but misleading distractors. For most question types, we automated this process using formalized operators, while justification questions required few-shot prompting to generate semantically diverse wrong options. Examples are provided in Appendix A.7.

3.3 Strict Quality Control of TactfulToM

We employed a multi-stage approach for strict quality control. For constructing seed stories, graduate students reviewed all white lie triplets to ensure logical consistency. During the generation, we created multiple versions (typically requiring 3-4 iterations) of each conversation segment and selected the best version to achieve natural dialogue flow while preserving the intended information asymmetries (more details can be found in Appendix A.4).

For final validation, we recruited 21 annotators from the Prolific platform² who met high-standard selection criteria. From an initial pool of 79 candidates, we conducted a qualification test designed to verify the ability to evaluate conversation coherence, resulting in the selection of 21 qualified annotators. Each conversation was reviewed by three independent annotators who flagged potential issues with coherence, safety, or white lie authenticity. While we received occasional flags, no conversation received majority votes for removal (detailed screening criteria, qualification test design, and validation results can be found in Appendix B.1).

3.4 Statistics

TactfulToM comprises 100 conversations spanning 5 white lie categories and 3 difficulty levels (based on falsifiability), containing 6.7K questions across multiple formats: multiple-choice, binary, and free-form questions. The dataset features an average of 23.66 tokens per question, with conversations averaging 32.6 turns and 33.0 tokens per turn. Table 1 presents a comparison of basic statistics between TactfulToM and other ToM benchmarks.

²<https://www.prolific.com/>

Dataset	Total #Convs	Total #Questions	Avg. Length per Turn (#Tokens)	Avg. #Turns per Conv.
ToMATO	753	5.4K	41.6	16.0
FANToM	256	10K	31.4	24.5
TactfulToM (Ours)	100	6.7K	33.0	32.6

Table 1: A comparison of dataset statistics between TactfulToM, ToMATO (Shinoda et al., 2025a) and FANToM (Kim et al., 2023). Conv. refers to conversation.

4 Experiments

4.1 Model Choice

We test nine LLMs from four families, including vanilla and reasoning models (indicated by *): **GPT**: gpt-4o-2024-08-06 (Hurst et al., 2024), o1-2024-12-17* (Jaech et al., 2024), o3-mini-2025-01-31*³; **DeepSeek**: DeepSeek-V3-0324 (DeepSeek-AI, 2024), DeepSeek-R1-Turbo* (DeepSeek-AI, 2025); **Llama**: Llama-3.3-70B-Instruct (Grattafiori et al., 2024); **Qwen**: Qwen2.5-72B-Instruct (Yang et al., 2025), QwQ-32B* (Qwen Team, 2025).

For the vanilla non-reasoning models, we employed two types of prompt templates: one that guided them to generate a direct answer, and another that guided them to generate the answer after a Chain-of-Thought (CoT) process. For the reasoning models, we only used the direct-answer prompt template. We presented the prompt templates in Table 7, Appendix B.3. In addition to CoT, we also tested a compatible ToM-specific method (Shoes-of-Others prefixing (Shinoda et al., 2025b)) for our white lie scenarios, with results in Appendix B.4.

4.2 Metrics

We employ four question formats across our evaluation framework: MCQs, binary, list-type, and free-form responses. Comprehension, Justification, Lie Ability, Belief, and Fact questions (except Lie Ability: MCQs only), while Lie Detectability, Info Accessibility, and Answerability questions use the binary and list formats. For structured responses (MCQs, binary, and list), we use accuracy as the primary evaluation metric and conduct detailed analyses of error patterns. For freeform responses, we determine the closest option using three complementary methods: cosine similarity (all-MiniLM-L6-v2⁴), token-F1, and LLM-as-judge (DeepSeek-v3). Given the varying chance levels across formats,

³<https://openai.com/index/openai-o3-mini/>

⁴<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

we report the MCQs and list format results, while using free-form responses for in-depth analysis.

4.3 Human Performance

We evaluated human performance through 12 participants (9 annotators and 3 graduate students) on 15 sets of questions (drawn through stratified random sampling from the 100 question sets to ensure balanced representation across all five question categories and three complexity levels). To remove redundancy, we selected one format for each question type as follows: Comprehension [binary], Justification [MCQs], Lie Ability [MCQs], Lie detectability [list], belief [MCQs], Information Accessibility [list], and Answerability [list]. We collected 2-3 responses from different tessees for each set. Participants received the same instructions as the models in order to compare them equally. More details can be found in Appendix B.1.

4.4 Results

Figure 3 displays the full results of examined LLMs on TactfulToM. We categorize the results according to question types mentioned in Section 2.3 and use different colors to represent different models. Detailed scores are provided in Table 6 in Appendix B.2.

Overall Performance GPT-4o and DeepSeek families consistently outperformed all other model families. DeepSeek models demonstrate a slight edge over GPT-4o on several tasks, including justification and Lie Ability questions. However, compared to humans who achieved an accuracy rate of over 85% on all tasks, all current models still exhibit a substantial gap in our benchmark.

Vanilla vs. CoT Prompting vs. Reasoning Models CoT prompting shows inconsistent benefits across model families. GPT models show minimal improvements or even degrade performance with CoT prompting, particularly on lie detectability tasks. GPT reasoning models also unexpectedly underperformed their regular models. DeepSeek models exhibited a different pattern, with reasoning variants outperforming both vanilla models and CoT-prompted versions across most question categories. Llama and Qwen families demonstrated no consistent pattern in response to either CoT prompting or reasoning-specialized models. These findings suggest that current reasoning enhancement techniques provide inconsistent benefits for ToM

reasoning involving white lies, indicating the need to improve performance in this domain.

LLMs Struggle with True White Lie Understanding As described in Section 2.3, true white lie understanding requires models to identify falsity while recognizing prosocial motivation. However, as shown in Figure 4, model performance drops significantly on this combined task, with even the best models achieving $< 50\%$ accuracy. This suggests that models may succeed on individual dimensions by chance or through pattern matching, without integrating the complementary aspects required for genuine understanding. DeepSeek-v3 performs best but remains far from human-level competence. Given that psychological research shows second-order ToM reasoning as a necessary condition for white lie understanding (Happé, 1994), this result encourages further investigation into the second-order ToM reasoning capabilities of current LLMs.

LLMs Can Track Mental States But Fail to Apply Them in White Lie Contexts Our analysis reveals a performance gap between Info-State questions and White Lie Reasoning questions. While models track beliefs reasonably well, they struggle with questions requiring the application of these representations, particularly lie detectability where accuracy drops significantly. This pattern is consistent across all model families. This suggests two possibilities: either current LLMs possess mental state tracking abilities but cannot integrate these states to understand behavioral capabilities in white lie scenarios, or their apparent success in belief tracking may be superficial, lacking genuine second-order ToM reasoning needed to determine conditions for detecting deception.

4.5 In-depth Analysis

Common Sense Falsehoods Are Easier for Models Our analysis reveals performance differences across different white lie classes as shown in Figure 5. While Info-State questions show consistent performance, White Lie Understanding and Reasoning questions vary significantly. Models perform exceptionally well in Classes 1 and 2; this pattern suggests models use commonsense knowledge as a shortcut rather than engaging in genuine contextual reasoning. For Class 1 scenarios involving globally recognized falsehoods (e.g., “Santa is real”), models can directly identify the statement as false without complex belief reasoning. Simi-




Figure 3: The answer accuracy of different LLM families on our benchmark: Comprehension[MCQs], Justification[MCQs], Lie Ability[MCQs], Lie Detectability[list], Belief[MCQs], Info Accessibility[list], and Answerability[list]




Figure 4: The proportion of model performance types in Justification questions.

larly, Class 2 scenarios featuring symbolic explanations of sensitive topics (e.g., death) are recognizable through common patterns in the data. In contrast, scenarios requiring situation-specific reasoning without obvious common sense cues pose significantly greater challenges, highlighting that models still largely rely on statistical regularities rather than sophisticated ToM capabilities when navigating white lie understanding.

Surface-Level Detection vs. Motivation Understanding Table 2 shows a clear drop in model performance across three falsifiability levels (described in Section 2.2). DeepSeek-v3’s comprehension, for instance, falls from Level-1 (79.41%) to Level-3 (34.78%). This reveals a critical insight: models excel at detecting lies through explicit contradictions but struggle to infer deception directly from motivations. A manual examination of DeepSeek’s reasoning reveals the model primarily identifies lie detectability by checking

Model	Level	Comp	Justi	B-2	LieAb	LieDe
GPT-4o	L-1	79.41	55.88	69.02	51.06	8.82
	L-2	70.37	59.26	73.58	63.08	7.41
	L-3	60.87	43.48	71.50	47.06	4.35
DeepSeek	L-1	79.41	88.24	62.63	65.96	26.47
	L-2	51.85	66.67	69.81	64.62	18.52
	L-3	34.78	65.22	55.56	58.82	39.13

Table 2: Performance (%) of GPT-4o and DeepSeek-v3 across levels. Abbreviations: Comp=Comprehension, Justi=Justification, B-2=2nd-order Belief, LieAb=Lie Ability, LieDe=Lie Detectability.

which characters have access to what information. This pattern explains DeepSeek-v3’s counterintuitive improvement in lie detectability for Level-3 (39.13%) compared to Level-1 (26.47%): without explicitly stated truths, the model faces less confusion about characters’ information access but fails to recognize that genuine white lie detection requires understanding protective intentions, not merely contradiction recognition.

Models Struggle with Genuine Motivation Understanding Without Guidance To assess models’ true comprehension of white lie motivations, we examined the Justification question’s free-form responses where no options provide hints. As shown in Table 3, models’ performance drops significantly from MCQs to free-form responses. DeepSeek’s falls from 75% to approximately 30% across different metrics. This gap suggests multiple-choice accuracy is inflated by provided options, as models struggle to independently infer the prosocial intentions behind white lies. Even the best-performing models fail to identify emo-




Figure 5: The performance of models across different classes.

Model	MCQs	FreeForm		
		Cos. Sim.	Token-F1	LLM-Judge
GPT-4o	53.57	22.62	27.38	16.67
DeepSeek	75.00	29.76	35.71	26.19
Qwen	57.14	19.05	9.52	25.00
Llama	46.43	20.24	10.71	23.81

Table 3: The accuracy of the model’s CoT performance in Justification tasks under different task formats and evaluation methods.

tional protection motivations in most free-form responses, highlighting significant limitations in their unprompted emotional reasoning.

5 Related Work

ToM in Psychology ToM has been categorized into seven abilities by the Abilities in Theory of Mind Space framework (Beaudoin et al., 2020): Intentions, Desires, Emotions, Knowledge, Percepts, Beliefs, and Mentalistic Understanding of Non-literal Communication. Non-literal communication understanding enables the interpretation of non-literal language such as irony, sarcasm, and white lies, where the intended meaning diverges from the literal. This requires second-order ToM ability, which is typically assessed through false-belief tasks and nested belief attribution (Wimmer and Perner, 1983; Quesque and Rossetti, 2020) in psychology experiments. As Beaudoin et al. (2020) demonstrated, accurately interpreting such expressions depends on the listener’s capacity to infer communicative intent and consider the speaker’s emotional motivations. These forms of pragmatic inference are especially relevant in white lies, where the goal may be to avoid harm or maintain relationships (Erat and Gneezy, 2012). This reflects a broader understanding of ToM as a key mechanism for navigating complex social communication, supported by evidence from developmental, clinical, and neurocognitive studies (Baron-Cohen et al., 1985; Langley et al., 2022).

ToM in LLMs Most existing ToM evaluations of LLMs focus on false-belief tests, such as the benchmarks ToMi (Nematzadeh et al., 2018), ToM-QA (Le et al., 2019), and FANToM (Kim et al., 2023), primarily testing whether models can track belief states when objects are moved or information changes. Some recent works go beyond false belief tests by incorporating broader mental states testing (Chen et al., 2024; Shinoda et al., 2025a) and explore ToM in applied social scenarios (Chan et al., 2024; Bara et al., 2021; Shapira et al., 2023b), or improving evaluation methods (Gandhi et al., 2023). However, white lies remain largely understudied. While ToMBench (Chen et al., 2024) included white lie tests, it offers only 20 non-conversational samples, too limited for comprehensive evaluation. This limited understanding of white lie capabilities poses risks as LLMs are increasingly deployed in emotional support and caregiving applications where such skills are essential.

6 Conclusion

We present TactfulToM, an English ToM benchmark designed to evaluate LLMs’ understanding of white lies through complex social scenarios. Our comprehensive evaluation reveals that even state-of-the-art LLMs underperform compared to humans in white lie understanding and reasoning, particularly in understanding the emotional motivation behind it. This performance gap raises ethical questions about LLMs’ development: should LLMs understand white lies merely to interpret human behavior, or also to potentially generate them? The dilemma lies in choosing between strict truthfulness and social grace that might involve benign deception. TactfulToM provides a foundation for improving LLMs’ social reasoning of white lie understanding, but we must carefully consider whether aligning LLMs completely with human social behaviors, including prosocially-motivated deception, is truly desirable for human-AI interaction.

Limitations

The main limitations of this paper are:

Limited to White Lies This dataset is primarily focused on white lie scenarios in order to analyze LLMs’ ToM capabilities in such contexts. We do not extensively explore LLMs’ other second-order ToM abilities; however, we hope that the methodology proposed in this paper can provide insights for future researchers seeking to construct related datasets.

Lack of Prior Impression In real-life situations, people typically possess prior knowledge and impressions of others. In our dataset, we deliberately constrained the scenarios such that the white lie triplets are not previously known to any of the involved roles, with the exception of the liar who initiates the deception. While this design choice helps isolate the ToM reasoning process, it does not fully capture the complexity of real-world social interactions. We consider incorporating this aspect of human cognition in our future work.

Limited Culture and Language Our benchmark includes only English-language data. However, in some other languages and cultures, communication tends to be more indirect, which may lead to different patterns of ToM reasoning in white lie scenarios.

Societal and Ethical Considerations

We acknowledge that our focus on white lies and Theory of Mind may raise concerns about anthropomorphizing AI systems. However, our research does not advocate for developing AI systems capable of telling white lies. Rather, we aim to systematically evaluate LLMs’ social reasoning capabilities within specific informational contexts. Our results demonstrate that current models fall significantly short of human-like understanding in these scenarios, primarily relying on pattern matching rather than genuine understanding of mental states or intentions. We recognize the ethical complexities surrounding deception, even when prosocially motivated, and the particular sensitivity of developing AI systems with capabilities that could involve any form of misrepresentation.

All annotators participating in our data collection and validation were recruited through Prolific. We established fair compensation standards based on estimated task duration, ensuring payment rates

above minimum wage requirements. We maintained transparent communication channels with annotators, promptly addressing questions and incorporating feedback to improve task instructions. All annotator data was anonymized, with only minimal identifiers stored securely and not included in the released dataset. We were careful to design our task instructions clearly, providing sufficient context without biasing responses. Annotators were informed about the academic research nature of the task and how their contributions would be used. When selecting annotators, we sought diversity across demographic factors to minimize potential biases in our data collection process, though we acknowledge that online recruitment platforms have inherent demographic limitations.

Our dataset is intended for research purposes only. While we have taken measures to ensure the conversations do not contain offensive content, research using generative models always carries a risk of unexpected outputs, particularly in free-form reasoning contexts. We encourage responsible use of our benchmark and dataset for advancing understanding of social reasoning in AI systems while remaining mindful of potential misapplications.

Acknowledgments

This work was supported by JST FOREST Grant Number JPMJFR232R and JST K Program Grant Number JPMJKP24C2. Thanks to EPFL College of Humanities for providing financial support for Yiwei Liu to attend the conference and present this work. We would like to thank the anonymous reviewers for their valuable feedback and comments that helped improve this paper, also the participants who contributed to the annotation.

References

- Hojjat Abdollahi, Mohammad H Mahoor, Rohola Zandie, Jarid Siewierski, and Sara H Qualls. 2022. Artificial emotional intelligence in socially assistive robots for older adults: a pilot study. *IEEE Transactions on Affective Computing*, 14(3):2020–2032.
- Cristian-Paul Bara, Sky CH-Wang, and Joyce Chai. 2021. *MindCraft: Theory of mind modeling for situated dialogue in collaborative tasks*. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1112–1125, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

- Simon Baron-Cohen, Alan M Leslie, and Uta Frith. 1985. Does the autistic child have a “theory of mind”? *Cognition*, 21(1):37–46.
- Cindy Beaudoin, Élizabel Leblanc, Charlotte Gagner, and Miriam H Beauchamp. 2020. Systematic review and inventory of theory of mind measures for young children. *Frontiers in psychology*, 10:2905.
- Ryan Burnsworth. 2022. Popular names by birth year 1880-2022. <https://www.kaggle.com/datasets/ryanburnsworth/popular-names-by-birth-year-1880-2022>.
- Chunkit Chan, Cheng Jiayang, Yauwai Yim, Zheyue Deng, Wei Fan, Haoran Li, Xin Liu, Hongming Zhang, Weiqi Wang, and Yangqiu Song. 2024. **NeegotiationToM: A benchmark for stress-testing machine theory of mind on negotiation surrounding**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4211–4241, Miami, Florida, USA. Association for Computational Linguistics.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, and Minlie Huang. 2024. **ToMBench: Benchmarking theory of mind in large language models**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15959–15983, Bangkok, Thailand. Association for Computational Linguistics.
- DeepSeek-AI. 2024. Deepseek-V3 technical report. *Preprint, arXiv:2412.19437*.
- DeepSeek-AI. 2025. Deepseek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *Preprint, arXiv:2501.12948*.
- Sanjiv Erat and Uri Gneezy. 2012. White lies. *Management science*, 58(4):723–733.
- Kanishk Gandhi, J-Philipp Fränken, Tobias Gerstenberg, and Noah D. Goodman. 2023. Understanding social reasoning in language models with language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The Llama 3 herd of models. *Preprint, arXiv:2407.21783*.
- Yuling Gu, Oyvind Taffjord, Hyunwoo Kim, Jared Moore, Ronan Le Bras, Peter Clark, and Yejin Choi. 2024. SimpleToM: Exposing the gap between explicit ToM inference and implicit ToM application in LLMs. *Preprint, arXiv:2410.13648*.
- Francesca GE Happé. 1994. An advanced test of theory of mind: Understanding of story characters’ thoughts and feelings by able autistic, mentally handicapped, and normal children and adults. *Journal of autism and Developmental disorders*, 24(2):129–154.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. GPT-4o system card. *arXiv preprint arXiv:2410.21276*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. OpenAI o1 system card. *arXiv preprint arXiv:2412.16720*.
- Hyunwoo Kim, Melanie Sclar, Xuhui Zhou, Ronan Bras, Gunhee Kim, Yejin Choi, and Maarten Sap. 2023. **FANToM: A benchmark for stress-testing machine theory of mind in interactions**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14397–14413, Singapore. Association for Computational Linguistics.
- Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. 2017. Building machines that learn and think like people. *Behavioral and brain sciences*, 40:e253.
- Christelle Langley, Bogdan Ionut Cirstea, Fabio Cuzolin, and Barbara J Sahakian. 2022. Theory of mind and preference learning at the interface of cognitive science, neuroscience, and AI: A review. *Frontiers in artificial intelligence*, 5:778852.
- Matthew Le, Y-Lan Boureau, and Maximilian Nickel. 2019. **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, Hong Kong, China. Association for Computational Linguistics.
- Aida Nematzadeh, Kaylee Burns, Erin Grant, Alison Gopnik, and Tom Griffiths. 2018. **Evaluating theory of mind in question answering**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2392–2400, Brussels, Belgium. Association for Computational Linguistics.
- David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526.
- François Quesque and Yves Rossetti. 2020. What do theory-of-mind tasks actually measure? theory and practice. *Perspectives on Psychological Science*, 15(2):384–396.
- Qwen Team. 2025. **QwQ-32B: Embracing the power of reinforcement learning**.

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](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13960–13980, Toronto, Canada. Association for Computational Linguistics.

Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2023a. [Clever hans or neural theory of mind? stress testing social reasoning in large language models](#). *Preprint, arXiv:2305.14763*.

Natalie Shapira, Guy Zwirn, and Yoav Goldberg. 2023b. [How well do large language models perform on faux pas tests?](#) In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10438–10451, Toronto, Canada. Association for Computational Linguistics.

Kazutoshi Shinoda, Nobukatsu Hojo, Kyosuke Nishida, Saki Mizuno, Keita Suzuki, Ryo Masumura, Hiroaki Sugiyama, and Kuniko Saito. 2025a. [Tomato: Verbalizing the mental states of role-playing llms for benchmarking theory of mind](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 1520–1528.

Kazutoshi Shinoda, Nobukatsu Hojo, Kyosuke Nishida, Yoshihiro Yamazaki, Keita Suzuki, Hiroaki Sugiyama, and Kuniko Saito. 2025b. [Let’s put ourselves in sally’s shoes: Shoes-of-others prefixing improves theory of mind in large language models](#). *Preprint, arXiv:2506.05970*.

Tomer Ullman. 2023. [Large language models fail on trivial alterations to theory-of-mind tasks](#). *Preprint, arXiv: 2302.08399*.

Alex Wilf, Sihyun Lee, Paul Pu Liang, and Louis-Philippe Morency. 2024. [Think twice: Perspective-taking improves large language models’ theory-of-mind capabilities](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8292–8308, Bangkok, Thailand. Association for Computational Linguistics.

Heinz Wimmer and Josef Perner. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children’s understanding of deception. *Cognition*, 13(1):103–128.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. [Qwen3 technical report](#). *Preprint, arXiv:2505.09388*.

Zhining Zhang, Chuanyang Jin, Mung Yao Jia, Shunchi Zhang, and Tianmin Shu. 2025. [Autotom: Scaling model-based mental inference via automated agent modeling](#). *Preprint, arXiv:2502.15676*.




Figure 6: Proportion of different classes in TacfulToM.

A TactfulToM Construction

A.1 Dataset Distribution

The proportion distribution of different classes within the TacfulToM dataset is shown in Figure 6.

A.2 Prompt Templates for Conversation Generation

Prompt templates for generating the conversation are listed below as Step_1, Step_2, Step_3, and Step_4.

Step_1

In this conversation, four characters are participating: A: the liar name, B: the target name, C: the accomplice name and D: the observer name. Start the conversation with all participants discussing Topic for the scenario. They are relationship descriptor, A: the liar name and C: the accomplice name are parents, B: the target name is kid, and D: the observer name is aunt/uncle.

Step_2

Now, B: the target name leaves the conversation because of the reason: leave reason B, while D: the observer name leaves the conversation because of the reason: leave reason D_1.

With B: the target name and D: the observer name absent, A: the liar name and C: the accomplice name continue their conversation, shifting the topic to 'real_reason_c'. Ensure the dialogue flows naturally without directly repeating these exact words.

Remember to indicate that B: the target name and D: the observer name are leaving the conversation.

Step_3

Now, only D: the observer name returns and rejoins the conversation after previously leaving. Note that B: the target name has not returned to the conversation yet.

Then the conversation continues and naturally shifts to situation_topic. Do NOT catch up or recap details from the previous conversation.

And then 'truth_c' is naturally revealed in the conversation.

After engaging briefly, D: the observer name leaves the conversation again because of the reason: leave reason D_2.

IMPORTANT:

- Clearly indicate through dialogue only (e.g., "I'm back," or "I have to leave again now.") that D: the observer name first returns and later departs again.
- Do NOT use third-person narrative or stage directions.
- Avoid mentioning details from previous conversations.

Step_4

Now B: the target name returns to the conversation after leaving the conversation.

First have B: the target name explicitly indicate the return through dialogue. Do NOT catch up or recap details from the previous conversation.

And then situation naturally unfolds. Make sure the dialogue flows naturally without directly repeating these exact words.

In response, A: the liar name and C: the accomplice name tells B: the target name that 'the lie'.

IMPORTANT:

- Do NOT use third-person narrative or stage directions.
- Avoid mentioning details from previous conversations.

A.3 Leaving Reason List

bathroom break
coffee break

forgot something important
forgot to print some documents
forgot to receive a package
forgot to return a package
forgot to run errands
forgot to submit documents
have a meeting starting soon that I need to prepare for
have a previous engagement that I need to attend to quickly
have a work-related emergency that requires my immediate attention
have an unexpected visitor at my door
have errands to run
have to attend to someone who just walked in
have to check on something
have to go to the restroom
have to pick up a prescription
have to pick up dry cleaning
have to print or scan documents
have to receive a delivery
have to recharge laptop
have to return a borrowed item
have to take care of a family matter
have to take care of an unexpected task
have unexpected visitor
his/her pet needs attention
his/her family is calling
incoming delivery
must respond to a phone call
need to check on a friend or family member who needs assistance
need to finish a task that's time-sensitive
need to get a phone call
need to get some coffee
need to go to the toilet
need to grab a snack or a drink
need to have a quick chat with someone else
need to make a phone call
need to make a quick trip to the drug store
need to make a quick trip to the grocery store
need to pick up a package
need to receive a parcel
need to recharge cellphone
need to register for an event
need to schedule a haircut or salon appointment
need to schedule another appointment
need to step out for a moment
need to submit some papers
need to take care of some paperwork or documents

need to take care of some personal matters
 need to take care of something urgent
 need to troubleshoot something
 parking meter expiring
 remembered something that needs to be taken care of
 remembered to receive a package
 remembered to submit some papers
 remembered to take care of some paperwork or documents
 remembered to take care of some personal matters
 remembered to take care of something urgent
 want to go grab a drink
 want to go grab a coffee
 want to go take some fresh air
 want to go to the bathroom
 need to move my car
 have to take an urgent call from my boss
 need to check my emails quickly
 have to respond to an important message
 need to restart my computer
 have to take a quick medication
 need to handle a minor household emergency
 have to refill my water bottle
 have to feed my pet
 have to water my plants
 have to take a brief walk to clear my mind
 need to step outside briefly to meet someone
 have to adjust the thermostat
 need to quickly tidy up my workspace
 need to quickly verify something important
 have to quickly arrange something for a later meeting
 need to briefly step out to confirm travel arrangements
 have to take care of an urgent email
 have to briefly assist a coworker
 need to briefly leave to verify appointment details
 have to briefly tend to something outside
 have to quickly reschedule an upcoming meeting
 need to briefly attend to my child
 need to grab a quick snack
 need to quickly stretch my legs
 have to briefly troubleshoot my internet connection
 have to briefly step out for privacy reasons
 need to quickly tidy the room before another meeting

have to quickly update someone about my status
 need to briefly review notes or materials
 have to briefly leave to answer an urgent text
 have to briefly assist someone else in the household
 have to quickly plug in my device to charge
 need to briefly leave to retrieve an important item
 have to step away briefly to close a window or door
 need to quickly ensure I turned off an appliance
 have to briefly step away for a personal reason
 need to briefly step away due to allergy or health symptoms

Table 4: Predefined reasons for characters leaving the conversation.

A.4 Quality control of generation

Generation filter: Each dialogue has four parts. Because later parts are conditioned on earlier ones, we observed quality decrease in later steps. To maintain high standards, conversation parts (steps 2-4) typically required 3-4 iterations or more to achieve natural, realistic dialogue flow. This iterative approach ensured that only high-quality conversations reached the annotation phase.

A.5 Clarification on the usage of gpt-4o for dataset generation

The generation pipeline relies on GPT-4o, while our pre-experimental results show it struggles with detecting white lies. This raises concerns about the appropriateness of using LLMs for generating conversational data involving white lies. However, this is precisely why we designed our white lie triplet and 4-step generation pipeline—to avoid having the model understand what constitutes a white lie and then generate conversations based on the model’s potentially flawed understanding of white lies. Our white lie triplet decomposes white lies into three core elements (the lie, the truth, and the real reason). The 4-step generation loop uses one element at each step, requiring the model to understand only individual components rather than sophisticated white lie reasoning. This method ensures the model builds conversations without being instructed to “write a white lie.”

A.6 Question Generation Templates

The question generation templates we used are provided below as examples in Table 5.

Question Type	Template
Comprehension	“Was what {liar} told {target}, {lie}, true?”
Justification	“Why did {liar} and {accomplice} tell {target} {lie}?”
Lie Ability	“Please choose the correct answer from the options below.” Correct: “{liar} tell {target} that {lie}’ because {liar} believes that {target} is unaware that {real reason}.” Wrong 1: {liar} tell {target} that {lie}’ because {liar} believes that {target} believes {real reason}.” Wrong 2: “{liar} tell {target} that {lie}’ because {liar} is unaware that {target} is unaware that {real reason}.” Wrong 3: {liar} tell {target} that {lie}’ because {liar} is unaware that {target} believes {real reason}.”
Lie Detectability	“Was what {liar} told {target}, {lie}, true?” List all the characters who know the precise correct answer to this question.”
Fact	“Who discussed {real reason}?” “Who discussed {truth}?”
Belief	“What does {character} believe about {real reason}?” “What does {character} believe about {truth question}?” “What does {liar} believe about {target}’s belief regarding {real reason}?”
Info Accessibility (List)	“List all the characters who know this information: {real reason}”
Info Accessibility (Binary)	“Does {character} know this information: {real reason}?”
Answerability (List)	“Who discussed {real reason}?” “List all the characters who know the precise correct answer to this question.”
Answerability (Binary)	“Who discussed {real reason}?” “Does {character} know the precise correct answer to this question?”

Table 5: Question generation templates Examples for different question types in TactfulToM.

A.7 Wrong Option Design

Belief Statement Options For second-order belief statements, we formalized four logically distinct cases using belief operators:

- $Bel_Z(\varphi)$: Person Z believes proposition φ
- $\neg Bel_Z(\varphi)$: Person Z is unaware of (or does not believe) φ

A second-order belief statement takes the form $Bel_X(\cdot)$, where the inner argument concerns Y’s

epistemic state about proposition p:

$Bel_X(Bel_Y(p))$	(X thinks Y thinks p)
$Bel_X(\neg Bel_Y(p))$	(X thinks Y is unaware of p)
$\neg Bel_X(Bel_Y(p))$	(X is unaware that Y thinks p)
$\neg Bel_X(\neg Bel_Y(p))$	(X is unaware that Y is unaware of p)

When “X thinks Y thinks p” is supported by the dialogue, we use $Bel_X(Bel_Y(p))$ as the correct answer. The remaining expressions serve as distractors representing three error types:

- Wrong attribution of Y’s first-order belief
- Wrong attribution of X’s meta-belief
- Simultaneous error in both belief layers

Justification Options For justification questions, we employed few-shot prompting with the following criteria:

- Correct answer must reflect the genuine prosocial motivation (e.g., sparing feelings, maintaining harmony)
- Wrong options:
 - Mutually exclusive
 - Plausible alternative explanations
 - Consistent with the dialogue context

Example:

Correct: “Because Jamey wants to politely decline without making Pearl feel bad about choosing an expensive restaurant.”

Wrong options:

- “Because Jamey actually has to work this weekend.”
- “Because Jamey dislikes Pearl and doesn’t want to spend time with him.”
- “Because Jamey already has dinner plans with someone else.”

A.8 An Example from TactfulToM

We provided a full conversation sample from TactfulToM below for reference:

Pearl: So, I was thinking about food and I’m curious, what’s everyone’s favorite cuisine? I absolutely love Italian, especially a good risotto. It just feels like a warm hug in a bowl!

Jamey: Oh, Italian is great! But for me, it's definitely Thai food. I love the bold flavors and the perfect balance of sweet, sour, and spicy. Pad Thai is my absolute favorite.

Donyae: Oh, all great choices! I'm all in for Mexican food. Tacos, enchiladas, and don't even get me started on the guacamole. It's just so satisfying.

Pearl: Yum, I could eat some good guac every day. Has anyone tried making it from scratch? It's surprisingly easy and so much better than store-bought.

Jamey: I have! I actually make a mean guacamole—and you're right, Pearl, fresh ingredients make all the difference. It's a hit whenever I have friends over.

Marithza: I bet! Guacamole's a crowd pleaser for sure. On another note, I tried making homemade sushi once. Let's say, rolling those little vinegared rice gems isn't as easy as it looks!

Donyae: I can imagine! But doing it yourself must be fun. Maybe we should have a cooking night sometime, where everyone brings a dish from their favorite cuisine.

Pearl: That's a fantastic idea, Donyae! We could even make it like a friendly mini competition. Each dish gets a few points for taste and presentation. What do you guys think?

Jamey: I'm totally in! It'll be a fun way to try different cuisines and maybe even learn a few new recipes from each other.

Donyae: Awesome! I just remembered, I need to submit some papers before the end of the day. I'll catch you guys later, though. This cooking night sounds like such a blast—keep me posted!

Pearl: Oh, I need to make a quick phone call. Let's definitely plan this out soon! Bye for now, everyone.

Marithza: Bye, Donyae and Pearl! We'll keep you in the loop. So, Jamey, what kind of place were you thinking about for our cooking night? We could do it at someone's house or rent out a small venue.

Jamey: Honestly, I love the idea of just doing it at someone's house. It keeps things cozy and informal. Kind of like why I prefer casual dining places over the upscale spots. There's something special about a laid-back atmosphere.

Marithza: I know what you mean. Sometimes those upscale places can feel a bit too... stiff, maybe? I guess it's nice to just unwind and not worry about which fork to use—chat over some good food without all the fancy stuff.

Jamey: Exactly! Plus, it can be pretty expensive

going out to those places too often. I'd much rather hit up a cozy little diner or a local joint where you can really taste the love put into the cooking.

Marithza: For sure! And there's something so genuine about finding that hole-in-the-wall spot with the best food. It kind of feels like discovering a secret everyone should know about.

Donyae: Hey everyone, I'm back! I finally submitted all the papers. What did I miss?

Marithza: Welcome back, Donyae! We were just talking about where to hold our cooking night. Jamey and I thought doing it at someone's house might be fun and cozy.

Donyae: That sounds perfect! I think the home setting adds a nice personal touch. Plus, it's easier to relax and enjoy each other's company.

Jamey: Absolutely, home-cooked vibes really can't be beat. Do you have a dish in mind for when we do this?

Donyae: I'm thinking something classic like enchiladas. They are always a hit and not too complicated. By the way, I have to leave again now—I have a meeting starting soon that I need to prepare for.

Marithza: No worries, Donyae. Thanks for popping back in! Let us know how the meeting goes.

Jamey: Yeah, hope all goes well, Donyae! We'll continue planning and keep you updated.

Pearl: Hey everyone, I'm back from my call! What did I miss in the world of culinary planning?

Jamey: Welcome back, Pearl! We were chatting about turning our houses into five-star restaurants for our cooking night.

Pearl: Sounds exciting! But speaking of five-star, there's this new Italian restaurant that just opened downtown. Would you all be interested in checking it out this weekend?

Jamey: That sounds amazing, Pearl, but unfortunately, I have to work this weekend, so I can't make it. Maybe next time?

Pearl: Oh no, that's a bummer! We'll definitely catch up soon then. Maybe for that cooking night we talked about—we can even bring some Italian-inspired dishes to you instead.

Marithza: I'm up for the restaurant visit if it's still on. I've been dying to try their truffle pasta from what I've heard.

Pearl: Awesome, Marithza! Let's make it a date then. We'll let Donyae know and hopefully, she can join us too.

Jamey: You guys enjoy it! Be sure to save me a slice of that truffle pasta, at least in spirit.

Marithza: We will! And we’ll definitely share all the delicious details with you. Catch up soon, Jamey!

Pearl: For sure, Jamey. Good luck with work, hope the weekend goes smoothly!

B Experiments

B.1 Annotation details

Annotator Recruitment Our participants for coherence validation were recruited on the Prolific platform⁵. They met the following criteria: English as their first and primary language, fluency in English, and completion of an undergraduate degree (BA/BSc/other). Additionally, participants had an approval rate between 95–100%. We posted a qualification screening test with 79 initial candidates. The screening test required participants to evaluate conversation coherence for three dialogues, two of which were intentionally scrambled for a few sentences. Applicants were required not only to identify incoherent conversations but also to provide explanations for why they considered the conversations incoherent. This screening process filtered the pool down to 21 reliable annotators to ensure the quality of human annotation, who served as participants for the coherence validation. We paid participants 9 pounds per hour, which is considered appropriate according to Prolific’s standards. This payment level ensures fair compensation for participants, who were predominantly located in English-speaking regions, thereby encouraging high-quality annotations.

For human performance annotation, the participants consisted of annotators (crowdworkers) recruited via Prolific and graduate students. We recruited a total of 9 qualified annotators through the Prolific platform. We additionally included 3 graduate students to diversify the participant pool and enhance annotation quality. The inclusion of graduate students allowed us to validate the reliability of crowdsourced annotations and maintain consistency across all raters.

Annotation Process for Coherence Each conversation is reviewed for coherence, safety, and white lie authenticity using binary assessment with written explanations for flags. Each conversation was reviewed by three annotators, and we required a 2/3 majority approval rate to retain conversations

⁵<https://www.prolific.com/>

in the final dataset. While 7 out of 100 dialogues received individual flags, none reached majority disagreement.

Human Performance Evaluation For conversation-wise sampling, we randomly sampled 15 out of 100 conversation sets using stratified sampling to maintain representation across all five classes and three difficulty levels (described in Section 2.2). The number of sampled sets (15/100) was determined by referencing the human evaluation ratio in similar ToM benchmarks - for example, FANToM (Kim et al., 2023) evaluated 32 out of 256 sets (12.5%), while our ratio (15%) provides comparable coverage. For each sampled conversation set, we collected 2-3 human responses to ensure reliability.

For question type-wise collection, human data was collected for only one format of each question type (e.g., only MCQ format, not freeform). This design choice reflects our human evaluation methodology, where participants received one conversation along with all associated questions for that conversation, making it redundant to ask the same conceptual question in multiple formats. Our model-human comparisons are therefore conducted only on equivalent question formats.

B.2 Model Performance on All Tasks

Detailed scores of the model performance on all tasks are provided in Table 6.

B.3 Prompt Templates for Model Evaluation

Prompt templates used for model evaluation are listed in Table 7.

B.4 ToM Method Testing

In addition to CoT, there are also some ToM-specific enhancement methods (Wilf et al., 2024; Sclar et al., 2023; Zhang et al., 2025; Shinoda et al., 2025b). But these methods most were designed for Sally-Anne-style belief-tracking tasks and cannot be directly applied to our white lie scenarios due to three key challenges:

- **Multiple asymmetries:** Each white lie scenario involves an information triplet (truth, lie, real reason), requiring models to track multiple asymmetries simultaneously, unlike Sally-Anne tasks that focus on a single hidden fact.
- **Graduated information access:** Different characters have varying levels of access to

Answer Type	Prompt Template
Binary	You are an expert in social reasoning. Answer the following question with 'Yes' or 'No'. Remember: Your answer should ONLY include 'Yes' or 'No' with nothing else. # Context: {} # Question: {} (Let's think step by step:)
MCQs	You are an expert in social reasoning. Answer the following question with the option number of the most appropriate answer. Remember: Your answer should ONLY include the option number with nothing else. # Context: {} # Question: {} # Options: {} (Let's think step by step:)
List	You are an expert in social reasoning. List the required items and split them with commas. Remember: Your answer should ONLY include the required items splited by commas with nothing else. # Context: {} # Question: {} (Let's think step by step:)
Freeform	You are an expert in social reasoning. Answer the following question with a single sentence. # Context: {} # Question: {} (Let's think step by step:)

Table 7: The prompt templates for model evaluation. The CoT prompt template additionally includes the instruction “Let’s think step by step:”.

tracking and white-lie comprehension, highlighting the need for specialized methods to handle the complex information structure and multi-role beliefs inherent in white-lie scenarios.

Question	Standard GPT-4o	GPT-4o with SoO
Prefixing Belief	61.56	63.94
LieAbility	60.62	59.36

Table 8: Results (%) by the *Let’s Put Ourselves in Sally’s Shoes* method (Shinoda et al., 2025b).

B.5 More Analysis

We provide additional error analysis by tracking the specific wrong options selected by models across different question types, which provide further insights into the failure modes of various LLMs when handling white lie scenarios. The distribution of error types for Belief Understanding (Figure 7), and

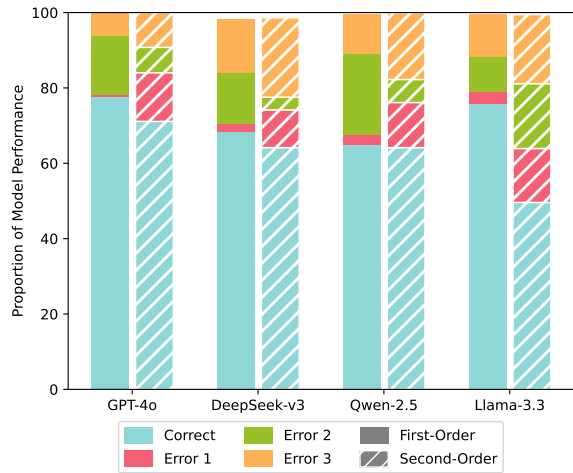


Figure 7: The proportion of model performance types in BeliefQA questions.

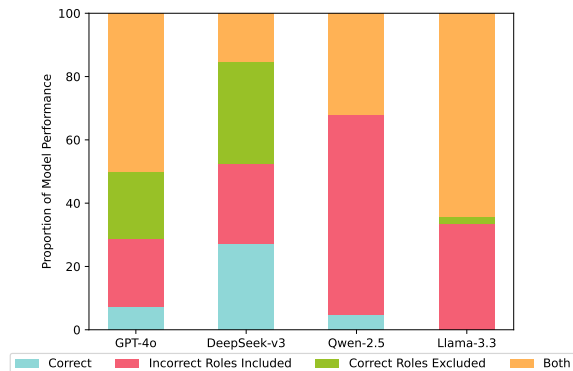


Figure 8: The proportion of model performance types in Lie Detectability questions of list format.

Role-Specific Performance in Lie Detection (Figure 8) reveal systematic patterns in how models misunderstand white lie contexts.

C AI Usage and Resources

In this project, we used LLMs for assistance. During paper writing, we used models from GPT and Claude families to help us refine and enhance our expressions. For programming, we also relied on models from GPT family to generate reference code, which we subsequently modified to complete our tasks.

Icons used in Figure 1 and Figure 2 were generated using Recraft⁶ and GPT-4o (Hurst et al., 2024) or sourced from Icons8⁷.

⁶<https://www.recraft.ai/>

⁷<https://icons8.com/icons>