

Are LLMs Smarter Than Chimpanzees?

An Evaluation on Perspective Taking and Knowledge State Estimation

Dingyi Yang¹, Junqi Zhao¹, Xue Li², Ce Li³, Boyang Li^{1*}

¹College of Computing and Data Science, Nanyang Technological University

²University of Science and Technology of China

³China University of Mining and Technology

{dingyi.yang, junqi.zhao, boyang.li}@ntu.edu.sg

lixue061219@mail.ustc.edu.cn, celi@cumtb.edu.cn

Abstract

Cognitive anthropology suggests that the distinction of human intelligence lies in the ability to infer other individuals' knowledge states and understand their intentions (Tomasello et al., 2005). In comparison, our closest animal relative, chimpanzees, lack the capacity to do so (Call and Tomasello, 2008; Hare, 2011). With this paper, we aim to evaluate LLM performance in estimating other individuals' knowledge states and their potential actions. We design two tasks to test (1) if LLMs can predict story characters' next actions based on their own knowledge vs. improperly using information unavailable from their perspective, and (2) if LLMs can detect when story characters, through their actions, demonstrate knowledge they should not possess. Results reveal that most current state-of-the-art LLMs achieve near-random performance on both tasks, and are substantially inferior to humans. We argue future LLM research should place more weight on the abilities of knowledge estimation and intention understanding. The project is available at: <https://github.com/DingyiYang/Act-on-Known/>.

1 Introduction

As a flurry of new Large Language Models (LLMs) continues to attain ever-higher scores on performance leaderboards (Deepmind, 2024; Xu and Ma, 2025), it is often difficult to shake the feeling that the LLMs are not yet matching human intelligence. On the other hand, it is also increasingly difficult to articulate where exactly the models fall short. What, if any, is the remaining difference between human and machine intelligence?

To understand the shortcomings of machines, we turn to cognitive anthropological research that attempts to distinguish *homo sapiens* from apes and other animals. A popular view there is that

Story: Mark hurried through the park with his dog, Rufus. He tossed a stick into the trees and waited for the usual thudding return. <knowledge>Deeper in the brush, Rufus stopped beside a person lying in the leaves, bleeding and barely conscious. </knowledge> Mark checked the time and called the dog's name. The stick hadn't come back.

Question: What will Mark do next?

A (without implausible knowledge): Look for his dog Rufus.
B (with implausible knowledge): Call 911 to save the injured person.



Thinking: There was an injured person, Mark should call 911 to save the person.
Answer: B ❌



Thinking: Mark did not know about the injured person. He would look for his dog first.
Answer: A ✅

Figure 1: The Knowledge-sensitive Next-action Prediction (KNP) task, where the implausible knowledge is highlighted for easy reading (but not highlighted for the LLMs and the human test-takers). This tests whether LLMs can take the perspective of the story characters and choose actions based on knowledge they can access.

humans possess an advanced theory of mind (ToM) (Premack and Woodruff, 1978) but other animals do not (Povinelli and Vonk, 2003; Suddendorf and Corballis, 2007; Krupenye and Call, 2019). That is, humans are uniquely capable of understanding and emulating the mental states of other individuals.

Among these, arguably the most influential theory comes from Tomasello *et al.*, who posit that the *one* fundamental distinction between humans and other animals is the ability to *understand and share the intentions of others* (USI) (Tomasello, 1999). Tomasello et al. (2005) propose that USI contains a few components: First, one must recognize that other individuals are autonomous and have independent mental states; otherwise it may be assumed that everyone knows and wants exactly the same things. Second, one needs to understand that other individuals choose their actions based on the knowledge and actions available to them. Inter-

*Corresponding Author.

	Ability Tested			Knowledge Domain		#Samples	Story Length	Performance	
	Knowledge State Tracking	Potential Action Estimation	In-Context Localization	Knowledge Domain	Knowledge Creation			Random	GPT-5.4 (high)
<i>Theory-of-Mind Benchmarks</i>									
ToMi (Le et al., 2019)	✓	-	-	Location	Rule-Based	6K	40 words	0.50	0.90
FANTOM (Kim et al., 2023)	✓	-	-	3 Types	LLM	10K	537 words	0.50	0.76
T4D (Zhou et al., 2023a)	✓	✓	-	Location	LLM	-	-	-	-
SIMPLETOM (BehaviorQA) (Gu et al., 2026)	✓	✓	-	10 Types	LLM	1.1K	38 words	0.50	0.77
CharTOM-QA (BeliefQA) (Zhou et al., 2025)	✓	✓	-	General	Classic Novel	201	2K words	0.25	0.65
<i>Benchmarks Proposed in This Paper</i>									
KNP (Ours)	✓	✓	-	General	Human & LLM	500	93 words	0.50	0.61
IKD (Ours)	✓	✓	✓	General	Human & LLM	1K	601 words	0.25	0.53

Table 1: Comparison between existing datasets. KNP requires models to *track characters’ knowledge states* and *estimate their potential next actions*. IKD requires models to do so for each event in context, and to *localize any implausible action*. Compared to datasets that rely on limited LLM-generated scenarios or classic novels prone to data leakage, our datasets cover broader situations and require richer commonsense understanding of the real world.

estingly, experiments suggest that apes are unable to track what other apes know and believe (Call and Tomasello, 2008; Hare, 2011). It stands to reason that chimpanzees cannot infer intentions of others from reliable estimations of their knowledge states.

Therefore, a natural question to ask is: can LLMs infer the knowledge states of human actors? Using textual stories, we propose two probing tasks: (1) **Knowledge-sensitive Next-action Prediction** (KNP) (Fig. 1), which tests whether LLMs can predict actions that story characters will likely take, depending on their own knowledge. (2) **Implausible Knowledge Detection** (IKD) (Fig. 2), which tests whether LLMs can detect when a story character acts on knowledge they should not possess.

As shown in Figure 2, we construct two benchmark datasets, one for each task. We collect and summarize human-written stories from the internet¹. For KNP, we identify or create² a knowledge statement that cannot be plausibly known by a story character, but has important implications for their next action. The LLM being tested must predict the next action of the character, and the correct choice is that the character will act without the knowledge. The dataset consists of 500 QA pairs in 5 genres.

For the IKD task, we rewrite the story to introduce a plot hole (Ahuja et al., 2025). In the plot hole, a story character uses knowledge they cannot possibly possess in their action. The LLM is required to detect if a plot hole exists in a given story. The benchmark consists of 500 correct stories and 500 erroneous stories. Both KNP and IKD stories were validated by human annotators. Section 4 contains more details.

Our proposed benchmarks demand a more advanced capacity for tracking knowledge states and

¹www.reedsy.com and www.short-story.me

²If no such knowledge exists in the original story, we use LLM to introduce one.

understanding how they affect actions in the real world (Table 1 shows comparisons with existing datasets). They are highly adversarial to state-of-the-art LLMs. On the KNP task, humans with zero training and zero demonstrations achieve an accuracy of 92%, 23 percentage points (pp) above the best LLM, Claude-4.5-Opus. GPT-5.4 achieves only 61%, whereas GPT-4o performs at random chance. On the IKD task, Claude-4.5-Opus achieves 68% classification accuracy and 0.59 IoU for localization, whereas other models remain at or below roughly 60% accuracy. The zero-shot human baseline reaches 76% accuracy. The surprising failures of LLMs in understanding seemingly simple stories highlight their deficiency in the critical ability of estimating knowledge states and potential actions.

This paper makes the following contributions:

- Inspired by cognitive research, we propose an important dimension of LLM evaluation: the ability to understand and share intentions (USI) of humans. As USI can be difficult to measure directly, we propose to evaluate a prerequisite skill: knowledge state tracking and potential action estimation.
- We build simple and short tests that push most state-of-the-art LLMs (except the latest Claude and Gemini models) to marginally above random chance. The finding is interesting by itself and highlights the need for further research on USI.

2 Related Works

Theory of Mind (ToM) in LLMs. Fundamental to social intelligence, Theory of Mind refers to the ability to recognize the fact that other people possess distinct mental states and to estimate such states (Nguyen et al., 2025). Most studies that test

ToM in LLMs focus on the dimension of beliefs and knowledge (Kosinski, 2024; Kim et al., 2023; Wu et al., 2023; Nematzadeh et al., 2018; Xu et al., 2024), others cover intentions (Staab et al., 2023; Zhou et al., 2023b) and emotions (Wu et al., 2024).

In this work, we focus on the dimension of knowledge. Most existing works face two limitations: **1)** they explicitly ask about knowledge state understanding, like in ToMi (Le et al., 2019), Hi-ToM (Wu et al., 2023), and FANTOM (Kim et al., 2023). For example, given a story “Sally is in one room; in another room, Anne puts milk in a box,” they ask “What does Sally think is in the box?”, which directly cues LLMs to perform knowledge tracking. In contrast, testing the LLMs on the *action implications of knowledge* (our setting) is harder, using questions like “Sally wants to drink milk. What would she do? A: Go to the box in another room. B: Try to find milk.” The inference that the knowledge state is critical in this decision appears to be rather difficult for LLMs (see the first ablation in §5.2) but is vital for USI. **2)** Many works employ false-belief contexts, where the story presents a world state change (e.g., “Sally puts milk in a basket and leaves; Anne moves it to a box”), like in T4D (Zhou et al., 2023a) and OpenTOM (Xu et al., 2024). This setup provides two alternative world states, implicitly letting LLMs choose between them. In contrast, our stories present only an unchanging fact that the character cannot know. If LLMs fail to understand that one’s action is based on their own knowledge, they will directly use this inaccessible fact. This design more rigorously tests whether LLMs understand how perspective-specific knowledge affects action.

Additionally, compared to datasets built from standardized templates (Le et al., 2019) or limited scenarios (Gandhi et al., 2023; Chen et al., 2024; Gu et al., 2026), our dataset is drawn from diverse human-written stories, which cover a broader range of realistic scenarios and offer a more out-of-distribution assessment for ToM. Although CharToM-QA (Zhou et al., 2025) also targets general scenarios, it relies on classic novels and is subject to serious data leakage concerns. Moreover, the lengthy context (over 2K words) makes it unclear whether the challenge stems from context length or ToM understanding. Our KNP task, by contrast, features much shorter natural narratives (93 words on average), yet LLMs still show near-random performance.

Narrative Understanding. Many works focus on basic story element understanding, such as NarrativeQA (Karpinska et al., 2024), BookQA (Angelidis et al., 2019), and FairyTaleQA (Xu et al., 2022). FlawedFictions (Ahuja et al., 2025) investigates if LLMs can detect plot holes that may require complex logical reasoning. Our IKD task uses a similar setting but focuses on plot holes related to characters’ knowledge states.

3 Tasks of Implausible Knowledge

We define implausible knowledge as knowledge that an individual’s physical location, social status, or personal experience does not plausibly grant them. For instance, if a story character knows a stranger’s name at first glance, we would consider it implausible³. Other examples include knowledge of private thoughts of others, or knowledge about events that happened in a far away location, in the character’s absence, or in the future.

We design two tasks to evaluate LLMs’ understanding of knowledge. The KNP task (Sec 3.1) focuses on the prediction of next actions *ex ante*, whereas the IKD task (Sec 3.2) focuses on the detection of actions that depend on implausible knowledge *ex post*. As such, these two tasks test two complementary aspects of the understanding of implausible knowledge.

3.1 Task 1: Knowledge-sensitive Next-action Prediction

In this task, each story contains an implausible piece of knowledge \mathcal{K} that the story character C should not know, but \mathcal{K} has strong implications on the next action of C . For example, C is looking for ancient treasures and \mathcal{K} is the location of the treasure. Upon learning \mathcal{K} , C would go to the location immediately. Without \mathcal{K} , it would be implausible for C to go to the exact right location without proper justification.

After reading the story, the LLM is asked to choose the more plausible next action for story character C . There are two choices, an action that C would normally take without knowing \mathcal{K} , and an action C would take knowing \mathcal{K} . Clearly, the first action is the correct choice. We report the accuracy.

3.2 Task 2: Implausible Knowledge Detection

This task evaluates if an LLM can determine whether a story character possesses knowledge in-

³Unless, of course, the story already established that the character has magic or some mind-reading device.

accessible to them given the narrative context. Formally, let S be a narrative consisting of a sequence of events $E = \{e_1, e_2, \dots, e_n\}$. An erroneous story with implausible knowledge \mathcal{K} satisfies two conditions: (1) there is an event e_k where a story character C either states the information \mathcal{K} or takes an action premised on \mathcal{K} , and (2) it is highly unlikely for C to know \mathcal{K} at the time of e_k .

The task contains two subtasks. In the binary classification subtask, the LLM classifies each story into logically consistent (1) or inconsistent (0). We report the classification accuracy. Given a story x_i , a predicted class \hat{y}_i , and the ground-truth label y_i , the accuracy is simply $\sum_i^N \mathbb{1}(\hat{y}_i = y_i)/N$, where N is the number of stories and $\mathbb{1}(\cdot)$ is the indicator function.

In the localization subtask, if the LLM determines the story to be logically inconsistent, it must localize the error by identifying the error-inducing sentences s_{err} . We prompt the LLM to predict the error sentences \hat{s}_{err} in words, and compute the intersection-over-union (IoU) with the ground truth at the word level. The localization metric l_i for the i -th story is calculated as:

$$l_i = \begin{cases} \mathbb{1}(\hat{y}_i = 1), & \text{if } y_i = 1 \text{ (no error)} \\ \mathbb{1}(\hat{y}_i = 0) \text{IoU}(s_{err}, \hat{s}_{err}), & \text{if } y_i = 0 \text{ (error)} \end{cases}$$

The dataset-level localization metric is computed as the average of l_i .

4 Dataset Construction

4.1 Story Collection

We collect 500 open-source stories from www.reedsy.com and www.short-story.me, covering five popular genres of romance, fantasy, kids, mystery, and science fiction. To avoid copyright issues, we perform automatic summarization of all collected stories (see Fig. 7 for the summarization prompt), yielding the raw original stories S_{ori} .

4.2 Data Construction

As depicted in Figure 2, we propose a data construction process that combines LLM generation and human validation and refinement. For LLM generation, we use GPT-5 with high reasoning for all steps. Detailed prompts are in the Appendix.

Introducing Implausible Knowledge. Given an original story, we first identify a piece of referenced knowledge that satisfies two criteria: (1) at the time of an event e_k , a character cannot plausibly possess

this knowledge; and (2) the knowledge nevertheless could influence the character’s actions in e_k . If no suitable instance of such knowledge is present, we prompt the model to generate one and insert it into the narrative in a natural manner.

Next-action Question Construction. Given the implausible knowledge \mathcal{K} , character C , and event e_k , we use the story context preceding e_k as the input to the KNP task. We then construct a question about the character’s next action, where the answer depends on whether C knows \mathcal{K} . We provide two candidate answers: a correct option and a distractor. The question-generating LLM is prompted to ensure that, assuming that C does not know \mathcal{K} , the action without \mathcal{K} is clearly the correct choice and the distractor is not independently plausible for other reasons.

Human annotators check for the correctness of the answers. For example, suppose \mathcal{K} is “the girl whom C loves thinks he is weird,” and C is characterized as very shy. If the correct option is “he finds the courage to confess” and the distractor is “he does not have the nerve to confess”, the distractor may remain plausible solely due to shyness rather than being caused by \mathcal{K} . In this case, human annotators may choose to rewrite the options themselves or regenerate using the LLM. The final refined dataset contains 500 QA pairs for the KNP task, with stories averaging 93 words in length.

Erroneous Story Construction. The original stories S_{ori} do not contain plot holes. To obtain the erroneous stories for the IKD task, we rewrite event e_k under the counterfactual assumption that the character possesses the identified implausible knowledge, yielding a modified event e'_k . The LLM is prompted to produce a rewrite that is fluent and coherent while introducing as few changes to the original event as possible. At the same time, the inconsistency introduced by e'_k should not be explainable within the story world. In addition, if rewriting e_k affects subsequent events, the LLM is asked to make the minimal adjustments necessary to preserve narrative continuity after e'_k .

During human validation, the annotators recheck the validity of e'_k and overall continuity. For example, cases in which e'_k can be justified by “the character may obtain knowledge with magic power” or “the character is lying about their possession of the implausible knowledge” are not considered as errors. Human annotators regenerate e'_k until the requirements are met.

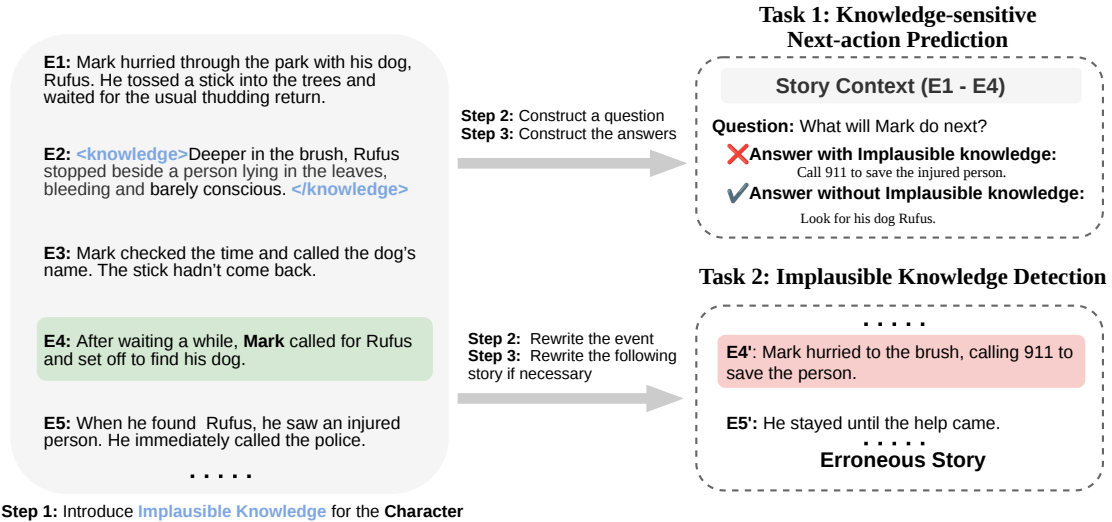


Figure 2: Our dataset construction process (Section 4). Using the original story, we first identify *implausible knowledge* that a *character* cannot possess at a given *event*, and that would affect the character’s action. If no such knowledge exists, we prompt the model to introduce one. Then: 1) For the KNP task, we construct a question about a knowledge-sensitive action, creating two answer options—one without the implausible knowledge and one with the knowledge; 2) For the IKD task, we rewrite the original event to generate an erroneous story.

Further, the human annotators are asked to remove explicit explanations that character C does not know \mathcal{K} in the story. The LLM being tested should reason about character knowledge states using commonsense, rather than hints in the text.

Combining the 500 erroneous and 500 original correct stories⁴, this yields the dataset for the IKD task, with an average length of 601 words.

Overall, approximately 40% of the samples in both tasks are refined by annotators. Additional details about the annotators and data refinement requirements are provided in the Appendix.

5 Experiments

LLM Tested. We test several models across two categories: closed-source models (Gemini-3-Pro, GPT-5.4, GPT-5, GPT-5-mini, GPT-4o, Claude-4.5-Sonnet, and Claude-4.5-Opus) and open-source models (LLaMA3.3-70B (Dubey et al., 2024a), LLaMA3.1-70B (Dubey et al., 2024b), LLaMA3.1-8B, Qwen3-MAX, Qwen3-32B (Yang et al., 2025), Qwen3-8B, DeepSeek-Chat, DeepSeek-R1 (v3.1 and v3.2) (Liu et al., 2025), Kimi-K2.5-Instruct (Team et al., 2026), and Kimi-K2.5-Thinking.

Settings. We test the LLMs using the zero-shot setting. Since the task requires reasoning about knowledge states, we always ask the models to

⁴Original stories are also verified and refined by human annotators to remove any possible logic errors.

generate their thinking process and output the answer. For temperature and top-p settings, we apply the default settings for all open-source models and closed-source APIs, then average scores across 3 runs. For the KNP task, we randomly shuffle the position of correct and distractor answers to avoid positional bias in LLMs.

Human Performance. We recruited 2 annotators with sufficient English proficiency and gave them the same instructions we provided to the LLMs. For each task, the annotators solved 50 questions.

5.1 Main Results

The results are shown in Table 2. For knowledge-sensitive action prediction (KNP), most models perform near-randomly — especially open-source models. Among leading models, Claude-4.5-Opus and Gemini-3-Pro achieve the highest performance at around 0.7. Analysis of their reasoning processes reveals that they consider knowledge states when predicting character actions (Case shown in Fig. 10). However, their understanding of knowledge states still needs improvement, revealing a gap with human performance of 0.92.

For logic error detection (IKD), many leading models perform only slightly better than random in the *binary classification task*. For instance, GPT-5 (high) achieves 0.58, DeepSeek-R1 reaches 0.61, and Claude-4.5-Opus shows the highest performance at 0.68. In the *localization task*, perfor-

	KNP		Implausible Knowledge Detection (IKD)			
	Action Choice	Binary Classification			Localization	
		Accuracy	Pos CIs	Neg CIs	Full CIs	Neg Loc
<i>Open-weight Models</i>						
Qwen3-8B	0.50	0.58	0.31	0.45	0.04	0.31
Qwen3-8B (think)	0.49	0.69	0.32	0.51	0.05	0.37
Qwen3-32B	0.49	0.78	0.23	0.51	0.04	0.41
Qwen3-32B (think)	0.50	0.65	0.45	0.55	0.10	0.38
Qwen3-Max (think)	0.53	0.86	0.24	0.55	0.07	0.47
LLaMA3.1-8B	0.51	0.14	0.64	0.39	0.11	0.13
LLaMA3.1-70B	0.55	0.41	0.60	0.51	0.08	0.25
LLaMA3.3-70B	0.54	0.69	0.36	0.53	0.05	0.37
Kimi-K2.5-Instruct	0.56	0.71	0.51	0.61	0.27	0.49
Kimi-K2.5-Thinking	0.57	0.74	0.49	0.62	0.26	0.50
DeepSeek-V3.1	0.47	0.40	0.65	0.53	0.16	0.28
DeepSeek-R1-V3.1	0.55	0.76	0.46	0.61	0.22	0.49
DeepSeek-V3.2	0.50	0.68	0.50	0.59	0.20	0.44
DeepSeek-R1-V3.2	0.57	0.94	0.27	0.61	0.16	0.55
<i>Closed-weight Models</i>						
GPT-4o	0.49	0.92	0.10	0.51	0.03	0.48
GPT-5-mini	0.52	0.83	0.18	0.51	0.06	0.45
GPT-5 (low)	0.57	0.81	0.37	0.59	0.18	0.50
GPT-5 (medium)	0.58	0.80	0.33	0.57	0.15	0.48
GPT-5 (high)	0.58	0.85	0.30	0.58	0.13	0.49
GPT-5.4 (high)	0.61	0.93	0.27	0.60	0.12	0.53
Gemini-3-Pro	0.67	0.28	0.86	0.57	0.47	0.38
Claude-4.5-Sonnet	0.66	0.35	0.82	0.59	0.29	0.32
Claude-4.5-Opus	0.69	0.88	0.48	0.68	0.30	0.59
Zero-shot Human	0.92	0.82	0.70	0.76	0.61	0.72

Table 2: Main results on KNP and IKD tasks. For IKD, we show the classification and localization performance on the positive set (correct stories), negative set (erroneous stories), and the full set.

mance is even lower—DeepSeek-R1 achieves 0.55, and Gemini-3-Pro reaches 0.38. An interesting finding is that different models also exhibit distinct judgment tendencies: GPT-4o shows a strong bias toward determining stories are logically consistent, while Gemini-3-Pro shows the opposite tendency. As a result, although some models display high performance on either the positive set (stories with no errors) or the negative set, the overall performance is poor. Human performance is also not particularly high, largely due to the lengthy context. If we highlight the events containing the error and the relevant knowledge, their localization performance on the negative set increases from 0.61 to 0.86.

Effect of Scale. Comparing LLMs with different parameter counts, such as LLaMA-3.1 (8B vs. 70B) and Qwen3 (8B, 32B, MAX), we see that model scaling improves performance only slightly. This suggests that knowledge estimation and tracking may not be easily solved by scaling alone.

Effect of Extended Reasoning. We explore whether different model series show improvement with extended thinking processes. For instance, we compare DeepSeek-chat and DeepSeek-R1, Qwen3’s thinking and non-thinking modes⁵, as well as GPT-5’s different reasoning levels. However, extended reasoning does not lead to clear improvement. This highlights that knowledge estimation and tracking cannot simply be resolved through test-time scaling and requires specialized training.

Takeaway

Most state-of-the-art LLMs show near-random performance in knowledge state tracking and potential action estimation. Test-time scaling does not show a clear advantage.

⁵For non-thinking mode, the model still outputs a short chain of thought.

5.2 Further Analysis and Ablation

This section presents an in-depth analysis of LLMs’ understanding of knowledge.

Effects of Genre. We examine whether narrative style affects performance. We compare five genres in our benchmark dataset: romance, fantasy, science fiction, kids, and mystery. As shown in Figure 3, different genres show no clear performance difference across either task. The fantasy and mystery genres show slightly lower performance—possibly because models tend to excuse errors as explainable within the fantastic and some mysterious settings.⁶

Ablation: Focusing LLMs on Character Knowledge. In the basic setting used in Table 2, the LLMs are not told that they need to track knowledge states of story characters. In order to test if the LLMs would perform differently if they know the ability being tested, in this experiment, we explicitly instruct them to focus on the specific knowledge states of story character.

The results are shown in Figure 4. On average, the explicit prompt to pay attention to character knowledge improves accuracy by 12.6 pp in IKD (Whole Set) and 14.8 pp in KNP, with GPT-5.4 achieving the highest boost. This suggests the LLMs do not yet understand the dependency between actions taken by humans and their knowledge, which is a critical component of USI according to Tomasello et al. (2005). After being prompted, the best accuracy is 85% on KNP, still falling short of the performance of the human participants who did not receive this prompt.

IKD Localization Ablation: Isolating Knowledge Tracking. The IKD localization task is the most complex and has the lowest performance. Doing this correctly requires making a sequence of decisions correctly: (1) identifying the statement of the implausible knowledge, (2) identifying the action of the story character premised on the implausible knowledge, (3) knowing that knowledge inconsistency needs to be checked, and (4) performing knowledge tracking correctly to classify the story and localization. This makes it difficult to interpret the results of IKD. To isolate the effects of Step #4, we build the following ablated versions of the IKD localization task.

⁶The stories explicitly state that the characters in our questions do not possess magic power (so they cannot use magic to obtain implausible knowledge).

Ablation 1: We explicitly highlight the erroneous event and instruct the model to identify whether it contains a logical inconsistency. Compared to the full task, this setting removes the need to scan and reason over all events, reducing the problem to detecting a logical error within a single, pre-identified event.

Ablation 2: We further highlight both the erroneous event and the event that introduces the implausible knowledge for the character, and ask the model to detect the logical inconsistency. This setting eliminates the challenge of locating the relevant background knowledge, requiring the model only to reason about the inconsistency between the two given events.

Ablation 3: Building on Ablation 2, we additionally instruct the model to explicitly focus on implausible character knowledge states. This represents the simplest setting, isolating the core capability under evaluation: whether the model understands that a character cannot possess certain knowledge given the narrative context.

The results are shown in Table 3. We find that for all models, the first two ablations do not bring clear improvement. However, in Ablation 3—where models are asked to focus on knowledge states—performance shows a clear increase. This aligns with our previous experimental findings: models do not pay much attention to knowledge states in their default setting. Yet even under this condition, performance remains modest for most models, highlighting the need for future improvements.

Takeaway

- (1) Even when other required skills, such as retrieval in a long context, are removed from the problem, the LLMs still perform worse than humans.
- (2) Performance improves when LLMs are asked to focus on knowledge, suggesting LLMs do not understand the dependency between actions taken by humans and their knowledge, which is critical for USI (Tomasello et al., 2005).

KNP: Binary Choice vs. Open-Ended Generation. For the Knowledge-sensitive Next-action Prediction task, the default setting is a binary choice task. We explore whether the provided answers serve as hints or distractors for the models.

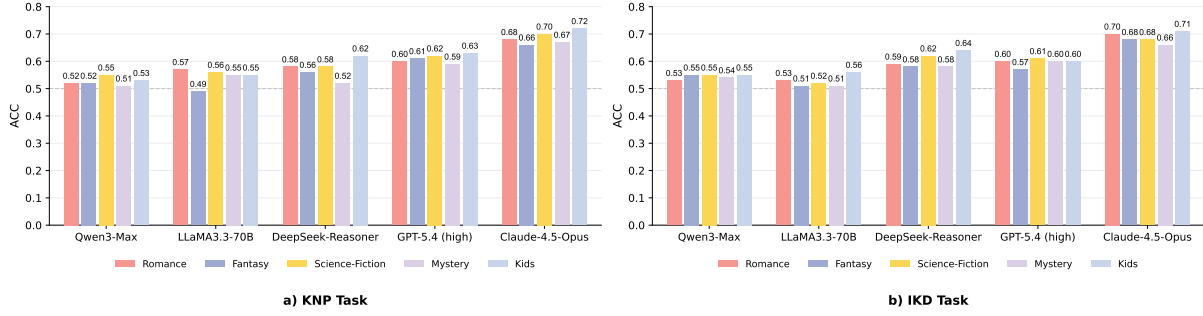


Figure 3: Performance on different genres.

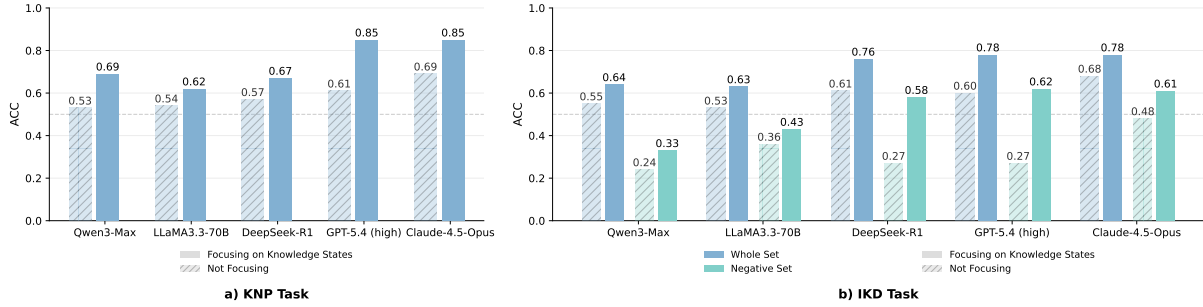


Figure 4: Ablation: Focusing LLMs on Character Knowledge.

Model	Localization (Negative Story)			
	Default	Abl. 1	Abl. 2	Abl. 3
Qwen3-8B (think)	0.05	0.12	0.14	0.33
Qwen3-32B (think)	0.10	0.34	0.29	0.60
Qwen3-Max (think)	0.07	0.28	0.33	0.69
LLaMA3.1-8B	0.11	0.20	0.23	0.28
LLaMA3.1-70B	0.08	0.29	0.39	0.67
LLaMA3.3-70B	0.05	0.18	0.3	0.55
Kimi-K2.5-Instruct	0.27	0.47	0.54	0.70
Kimi-K2.5-Thinking	0.26	0.46	0.54	0.71
DeepSeek-V3.2	0.20	0.36	0.41	0.78
DeepSeek-R1-V3.2	0.16	0.26	0.30	0.59
GPT-4o	0.03	0.12	0.22	0.62
GPT-5.4 (high)	0.12	0.22	0.30	0.61
Gemini-3-Pro	0.47	0.68	0.73	0.80
Claude-4.5-Opus	0.30	0.45	0.48	0.62
Human	0.61	-	0.86	-

Table 3: Ablation studies on the IKD-localization task (Neg set), which isolate the ability to estimate knowledge states from long-context understanding.

To investigate this, we conduct an open-ended prediction experiment: without given option answers, we first generate the open-ended answer, then apply the LLM-as-a-Judge method (Li et al., 2024) for open-ended evaluation. We provide GPT-5 with the story, question, answer options, and generated answer. If it aligns with the correct option, it’s correct; if aligns with the incorrect option, it’s wrong. Answers aligning with neither are filtered out—7% on average.

Model	Binary Choice	Open-Ended
Qwen3-8B (think)	0.49	0.53
Qwen3-32B (think)	0.49	0.51
Qwen3-Max (think)	0.53	0.53
LLaMA3.1-8B	0.51	0.52
LLaMA3.1-70B	0.55	0.57
LLaMA3.3-70B	0.54	0.58
Kimi-K2.5-Instruct	0.56	0.56
Kimi-K2.5-Thinking	0.57	0.56
DeepSeek-V3.2	0.49	0.51
DeepSeek-R1-V3.2	0.57	0.58
GPT-4o	0.49	0.52
GPT-5.4 (high)	0.61	0.62
Gemini-3-Pro	0.67	0.64
Claude-4.5-Opus	0.69	0.64

Table 4: Comparing the QA formats on the KNP task: binary-choice vs. open-ended generation. For binary choice, we report accuracy. For open-ended generation, we report the scores given by LLM-as-a-Judge.

The results are shown in Table 4. We observe that the two question formats do not create substantial differences in the scores, demonstrating that the conclusions based on binary choice setting are robust.

5.3 Case Studies

Figures 8–11 provide qualitative examples of model performance on the KNP task. While Claude-4.5-Opus and Gemini-3-Pro exhibit an emergent capacity to discern the relationship be-

tween character actions and knowledge states, they fall short of human proficiency, as evidenced by the quantitative results in Table 2 and the error case in Figure 8. Other models perform even worse, tend to apply knowledge directly to answer questions without analyzing whether the character can actually access it. Notably, extended reasoning does not help with this capacity—in the case shown in 8, DeepSeek-R1 generates 2.1K tokens in its reasoning process (within `<think></think>`), yet this provides no benefit for the inference.

Prediction results on the IKD task, shown in Figures 5-7, indicate that current LLMs are not very sensitive to actions driven by implausible knowledge when processing stories. Although Claude-4.5-Opus and Gemini-3-Pro perform better than others, they still show a gap compared to humans. Gemini-3-Pro, specifically, shows a tendency to be overly strict and tries to locate errors even in correct stories (Figure 7).

6 Conclusion

In this work, we explore whether current LLMs possess a significant Theory-of-Mind (ToM) capacity that distinguishes human intelligence from other animals: understanding others' independent knowledge states and predicting their corresponding actions. We carefully design two tasks for evaluation: Knowledge-sensitive Next-Action Prediction (KNP) and Implausible Knowledge Detection (IKD). Based on experimental exploration using our proposed benchmarks, most leading LLMs show near-random performance in this area, with a significant gap compared to humans. This finding highlights the need for future research.

Acknowledgments

We gratefully acknowledge the support by the National Research Foundation Fellowship (NRFF13-2021-0006), Singapore. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

Limitations

In this work, we focus on a fundamental prerequisite skill of ToM: evaluating LLMs' capacity for knowledge state tracking and estimation. We acknowledge that a full evaluation of ToM—Understanding and Sharing Intentions

(USI)—requires benchmarks covering more comprehensive dimensions. Currently, classical tasks and benchmarks may lack this granularity. For instance, Strachan et al. (2024) tested a wide range of tasks and found that LLMs often achieve performance comparable to or even better than humans. However, this high performance on classic benchmarks might mask the deficiency of LLMs. On our seemingly simple dataset, most leading LLMs fail to track other people's knowledge states, which is easy for humans. This contrast suggests that a gap between LLMs and real intelligence remains. We hope our findings inspire future works to explore USI.

Ethics Statement

We acknowledge and strictly adhere to the Code of Ethics and Professional Conduct throughout this research. The potential ethical concerns are addressed as follows:

Data Source and Copyrights. All data is collected from open-source websites. To avoid potential copyright and ethical issues, we summarize the content and remove any inappropriate content.

Crowdsourcing Services. For dataset refinement, we recruited three research assistants (aged 22–30) to review and correct any remaining errors. Each story took approximately 5 minutes to complete, and workers were paid \$12 per hour—a reasonable rate for the local area. For human evaluation, we recruited two undergraduate annotators (aged 18–22). Evaluators were paid \$1 per sample.

References

- Kabir Ahuja, Melanie Sclar, and Yulia Tsvetkov. 2025. Finding flawed fictions: Evaluating complex reasoning in language models via plot hole detection. *arXiv preprint arXiv:2504.11900*.
- Stefanos Angelidis, Lea Frermann, Diego Marcheggiani, Roi Blanco, and Lluís Màrquez. 2019. Bookqa: Stories of challenges and opportunities. *arXiv preprint arXiv:1910.00856*.
- Josep Call and Michael Tomasello. 2008. Does the chimpanzee have a theory of mind? 30 years later. *Trends in cognitive sciences*, 12(5):187–192.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, et al. 2024. Tombench: Benchmarking theory of mind in large language models. *arXiv preprint arXiv:2402.15052*.

- Deepmind. 2024. Ai achieves silver-medal standard solving international mathematical olympiad problems. <https://deepmind.google/blog/ai-solves-imo-problems-at-silver-medal-level/>. Accessed: 2025-12-15.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024a. The llama 3 herd of models. *arXiv e-prints*, pages arXiv-2407.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024b. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Kanishk Gandhi, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah Goodman. 2023. Understanding social reasoning in language models with language models. *Advances in Neural Information Processing Systems*, 36:13518–13529.
- Yuling Gu, Oyvind Taffjord, Hyunwoo Kim, Jared Moore, Ronan Le Bras, Peter Clark, and Yejin Choi. 2026. [Simpletom: Exposing the gap between explicit tom inference and implicit tom application in LLMs](#). In *The Fourteenth International Conference on Learning Representations*.
- Brian Hare. 2011. From hominoid to hominid mind: What changed and why? *Annual Review of Anthropology*, 40(1):293–309.
- Marzena Karpinska, Katherine Thai, Kyle Lo, Tanya Goyal, and Mohit Iyyer. 2024. One thousand and one pairs: A "novel" challenge for long-context language models. *arXiv preprint arXiv:2406.16264*.
- 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.
- Michal Kosinski. 2024. Evaluating large language models in theory of mind tasks. *Proceedings of the National Academy of Sciences*, 121(45):e2405460121.
- Christopher Krupenye and Josep Call. 2019. [Theory of mind in animals: Current and future directions](#). *WIREs Cognitive Science*, 10(6):e1503.
- 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.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. 2024. Llms-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579*.
- Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, et al. 2025. Deepseek-v3. 2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*.
- 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.
- Hieu Minh Nguyen et al. 2025. A survey of theory of mind in large language models: Evaluations, representations, and safety risks. *arXiv preprint arXiv:2502.06470*.
- Daniel J. Povinelli and Jennifer Vonk. 2003. Chimpanzee minds: suspiciously human? *7:157–160*.
- David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526.
- Robin Staab, Mark Vero, Mislav Balunović, and Martin Vechev. 2023. Beyond memorization: Violating privacy via inference with large language models. *arXiv preprint arXiv:2310.07298*.
- James WA Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, et al. 2024. Testing theory of mind in large language models and humans. *Nature Human Behaviour*, 8(7):1285–1295.
- Thomas Suddendorf and Michael C Corballis. 2007. [The evolution of foresight: What is mental time travel, and is it unique to humans?](#) *30:299–313*.
- Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, SH Cai, Yuan Cao, Y Charles, HS Che, Cheng Chen, Guanduo Chen, et al. 2026. Kimi k2. 5: Visual agentic intelligence. *arXiv preprint arXiv:2602.02276*.
- Michael Tomasello. 1999. *The Cultural Origins of Human Cognition*. Harvard University Press.
- Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. 2005. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(5):675–691.
- Jincenzi Wu, Zhuang Chen, Jiawen Deng, Sahand Sabour, Helen Meng, and Minlie Huang. 2024. Coke: A cognitive knowledge graph for machine theory of mind. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15984–16007.

Yufan Wu, Yinghui He, Yilin Jia, Rada Mihalcea, Yulong Chen, and Naihao Deng. 2023. Hi-tom: A benchmark for evaluating higher-order theory of mind reasoning in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10691–10706.

Hainiu Xu, Runcong Zhao, Lixing Zhu, Jinhua Du, and Yulan He. 2024. Opentom: A comprehensive benchmark for evaluating theory-of-mind reasoning capabilities of large language models. *arXiv preprint arXiv:2402.06044*.

Nan Xu and Xuezhe Ma. 2025. Llm the genius paradox: A linguistic and math expert’s struggle with simple word-based counting problems. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3344–3370.

Ying Xu, Dakuo Wang, Mo Yu, Daniel Ritchie, Bingsheng Yao, Tongshuang Wu, Zheng Zhang, Toby Jia-Jun Li, Nora Bradford, Branda Sun, et al. 2022. Fantastic questions and where to find them: Fairytaleqa—an authentic dataset for narrative comprehension. *arXiv preprint arXiv:2203.13947*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. *Qwen3 technical report*. *Preprint*, arXiv:2505.09388.

Chulun Zhou, Qiuqing Wang, Mo Yu, Xiaoqian Yue, Rui Lu, Jiangnan Li, Yifan Zhou, Shunchi Zhang, Jie Zhou, and Wai Lam. 2025. The essence of contextual understanding in theory of mind: A study on question answering with story characters. *arXiv preprint arXiv:2501.01705*.

Pei Zhou, Aman Madaan, Srividya Pranavi Potharaju, Aditya Gupta, Kevin R McKee, Ari Holtzman, Jay Pujara, Xiang Ren, Swaroop Mishra, Aida Nematzadeh, et al. 2023a. How far are large language models from agents with theory-of-mind? *arXiv preprint arXiv:2310.03051*.

Pei Zhou, Andrew Zhu, Jennifer Hu, Jay Pujara, Xiang Ren, Chris Callison-Burch, Yejin Choi, and Prithviraj Ammanabrolu. 2023b. I cast detect thoughts: Learning to converse and guide with intents and theory-of-mind in dungeons and dragons. In *Proceedings*

of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 11136–11155.

A Dataset Construction

A.1 Implausible Knowledge Types

Implausible knowledge may originate from several sources, including: private thoughts (another character’s unexpressed mental state that has not been externally revealed), distant events (events occurring elsewhere or during the character’s absence), future events (events that have not yet occurred), and inaccessible secrets (facts that are not reasonably accessible, such as a stranger’s family background or the hidden location of a map). We provide LLMs with these definitions as a reference.

A.2 Automatic Data Construction Prompts

The prompt for Next-action Question Construction (KNP task) is provided in Table 5. The prompt for an erroneous story construction (IKD task) is provided in Table 6.

A.3 Human Refinement

We recruit three workers to check and refine the auto-generated data. The refinement instructions are summarized as follows:

KNP Task: Annotators are tasked with reviewing the story context, questions, and answer options, in which implausible knowledge is explicitly highlighted in their annotation. They need to check the correctness of the answers. That is, given the story context and the character’s knowledge state, the right answer is clearly correct, and the distractor answer cannot be explainable. For example, consider a scenario where the implausible knowledge \mathcal{K} is “the girl whom C loves thinks he is weird,” and C is characterized as very shy. If the distractor is “he does not have the nerve to confess,” it is explainable. Even if readers understand that C is unaware of \mathcal{K} , they infer the distractor is plausible solely based on C ’s shyness. In such conditions, annotators are instructed to either rewrite the options manually or regenerate them using the LLM.

IKD Task: For *original stories*, annotators read the entire story and correct any logic errors. Then they read the *rewritten erroneous story*, focusing on the rewritten content and implausible knowledge, which are highlighted in the annotation. They must ensure that: (1) The implausible knowledge error is valid. For instance, if the character possesses

magic powers to know the information, the sample is filtered and regenerated until acceptable. (2) The erroneous story does not include an explicit explanation that someone lacks the knowledge. For example, in “Mark, not knowing the injured person in the brush, hurried to the brush and called 911”, the annotator removes the explicit hint “not knowing the injured person in the brush”. This ensures tested LLMs reason about character knowledge states using common sense rather than textual hints.

Annotators refined approximately 40% of the auto-generated data. Our authors then reviewed all refined data, making additional improvements to the data. We will release the data annotation file (containing detailed instructions and annotation samples) to support future research.

B Evaluation Prompts

For tested LLMs, we apply the prompt shown in Table 8 for the KNP task and the one shown in Table 9 for the IKD task.

Input (Erroneous Story)

Story:

Agatha, the old woman of the forest, finds a prim, fur-cloaked child at her doorstep: Lady Beatrice the Third, golden-haired and jeweled, certain she's found a witch. Beatrice insists she saw Agatha casting spells among the trees and demands to be taken on as an apprentice. Agatha denies it, teasing the girl with a mock curse toad or worm, only to be told, quite solemnly, that witches feed on young girls' souls—and Beatrice would trade hers if it meant learning magic. Alarmed by what the child reads but curious about her stubbornness, Agatha asks what she wants. Beatrice refuses to say.

Agatha returns to her chores, but the girl hounds her every step: inspecting the grain bin for hexes, peering into the henhouse as if the birds might be demons in disguise, squinting at garden tools for hidden charms. Day after day, the routine settles like weather. When ignoring fails to banish her, Agatha greets Beatrice in the mornings and, with a sigh, puts her to work. The little lady is no quick study, but she is willing. Agatha teaches her how not to startle the chickens, how to gather eggs, how to pull weeds by the roots, and which insects are foes and which are friends. The girl's pink dresses come away brown; she never once complains. In her mind, this is "secret witch training with the mysterious woman," and she lays verbal snares: Do witches have black cats or owls? How many make a coven? Have you ever turned someone into a frog?

As the cottage grows tidy and the chores thin out, Agatha teaches sewing, knitting, basket weaving. One morning she finds her sons' old pastels, half-used, and offers them. Beatrice's face brightens, then hardens into business again. She draws and chatters—until she goes quiet. Agatha leans over to find a radiant ocean: sand, gulls, sunlight glimmering like jewels. "It's my ocean," the girl whispers, then rushes out, leaving her gloves behind.

The next day brings no knock. Night brings the Duke, worry-lined, begging for help. The tracks lead into the forest; the rain has smothered the hounds. Agatha gathers shawl, bell-staff, blanket, and the gloves, and sends the men searching where they will not find her. She goes where the forest shows her: to the ancient tree whose roots sprawl like arms. There, under its lowered boughs, lies Beatrice, shivering in her ruined finery. Agatha wraps her, breathes warmth into her, and the girl revives with a triumphant smile. "I knew you were a witch! You used my gloves to find me." Agatha laughs and says: she found Beatrice because she knows the woods like her own hands—and because children often end up here, as she and her sons once did.

The child's bravado crumbles. She wanted a witch to send her back to the coast, to her mother's first home, to the ocean's diamonds. The forest feels dark and loud. Agatha explains that fear belongs to the unknown: the towering trees cradle life; the calls are neighbors. Home by the sea will remain hers, but this place can become home too, if she learns it. And the Duke? Odd, yes. Cruel, no. He was frantic to find her. The girl, surprised, leans in and sleeps. Agatha rings her bell; the searchers come. **The Duke stumbles under the boughs and, still breathless with panic, blurts, "No more of your 'secret witch training' with the mysterious woman."** Beatrice announces she wants the ways of the forest, not witchcraft. "Tomorrow," Agatha says.

Back home, aching but smiling, Agatha settles by the hearth. The fire blooms at a flick of her fingers. She thinks of tonics, and of which stray spell the girl might have glimpsed. Next time, she'll be more careful.

GPT-5 (high)

Reasoning: The plot is internally consistent. Agatha denies being a witch to protect herself and the child, but the ending reveals she does possess magic—a deliberate character deception, not a contradiction. The timeline (the girl leaving gloves, not arriving the next day, the Duke arriving at night, Agatha finding her at the ancient tree) is plausible. Agatha misleading the searchers and later summoning them with her bell also fits her motives and the circumstances (she wanted to reach the child first). Ambiguities like "the forest shows her" can be read figuratively or as subtle magic without breaking logic. No statements directly contradict established facts.

Logical Error: NA
Conclusion: No ❌

Gemini-3-Pro

Reasoning: I have analyzed the story's plot, character motivations, timeline, and consistency.

1. Character Consistency (Agatha): Agatha denies being a witch throughout the interactions with Beatrice but is revealed to be one in the final paragraph. This is a plot twist/dramatic irony, not a logical error. Her internal monologue acknowledging she must be "more careful" aligns with Beatrice's initial claim of having seen her casting spells.

2. Character Consistency (Beatrice): Beatrice's desire for magic changes from wanting power/to trade her soul, to wanting to go home, to accepting the forest. This is character development. Her specific fear/desire cycle is consistent.

3. Plot Flow & Timeline: Beatrice visits daily->One day gets upset and runs out->The next day she does not return (Agatha notes the absence)->That night the Duke arrives. Although the Duke waiting until the "next night" to search the cottage seems slow for a missing child, it acts on the logic that he first searched elsewhere or waited for her return, and the weather caused delays. It is not a logical impossibility.

4. The "Gloves" Logic: Beatrice accuses Agatha of using gloves for a spell (sympathetic magic). Agatha denies this, claiming she used woodcraft. The narration calls this "the truth." This is consistent: Agatha is a witch, but she chose to use her knowledge of the woods (and the fact that lost children gravitate to the ancient tree) rather than a spell to find the girl. Thus, the narrator's statement remains true even though Agatha is secretly a witch.

No inconsistencies contradicting the established logic were found.

Logical Error: NA
Conclusion: No ❌

Claude-4.5-Opus

Reasoning: Let me analyze the story for logical inconsistencies.

1. Agatha's denial of being a witch: Throughout the story, Agatha denies being a witch to Beatrice. She claims she found Beatrice "because she knows the woods like her own hands" rather than through magic. However, at the end, "The fire blooms at a flick of her fingers. She thinks of tonics, and of which stray spell the girl might have glimpsed." This reveals Agatha actually IS a witch.

2. Is this a logical error?: The fact that Agatha lies to Beatrice about being a witch is not itself a logical error - characters can deceive other characters. The story is internally consistent: Agatha is actually a witch who denies it to the child, and at the end we learn the truth. This is not a contradiction.

3. The gloves: Beatrice claims "You used my gloves to find me" and Agatha denies using magic. But since Agatha is a witch, she might have used magic. However, the story says "She goes where the forest shows her" which could be interpreted either way. This ambiguity is not a logical error.

4. Beatrice's behavior: She runs away wanting to go to the ocean/coast, to her mother's first home. This is consistent with her drawing the ocean and going quiet about it.

5. Timeline: The sequence of events flows logically - Beatrice arrives, works with Agatha, draws the ocean, runs away, is found, and is returned.

I don't find any clear logical inconsistencies. The story maintains internal consistency - Agatha hiding her true nature as a witch while teaching the girl practical skills, then being revealed as actually magical at the end, is a deliberate narrative choice, not a logical error.

Logical Error: NA
Conclusion: No ❌

Human

Reasoning: The phrase "secret witch training with the mysterious woman" is explicitly described as Beatrice's internal thought; she does not say it aloud. The Duke has no access to her private phrasing or inner framing of the chores, and there is no indication anyone told him.

Logical Error: The Duke stumbles under the boughs and, still breathless with panic, blurts, "No more of your 'secret witch training' with the mysterious woman." ✓

Conclusion: Yes ✓

Figure 5: Predictions on the IKD task (Erroneous Story). For easy reading, the logic error is in bold and implausible knowledge for the character is underlined.

Input (Erroneous Story)

Story:

Jen lives by the unwritten rules for the single woman who wants a child: don't look at children, don't engage, don't talk about them, don't let others talk about them, and absolutely don't confess that you want one. Any glance, any story, any softness around babies broadcasts your yearning like a tinny loudspeaker and turns you, in the mind of any would-be suitor, from possibility to problem.

So Fridays become her private rebellion. On her half day, she slips into a museum housed in a former church and sits before a marble mother and child—the Madonna, serene, nursing. Stone that feels like breath. A virgin, a mother, perfect and untouchable. From her first visit in early spring she loves Her, though at first she hurries past, afraid passersby will relish the pathos: a childless woman gazing at the mother of mothers. Still, she returns almost weekly, allowing herself to look openly, to feel the awe and ache she hides everywhere else.

Jen's childlessness feels to her like something visible and damp, soaking everything; she believes it's the first thing people see, the thing they instantly forgive with modern politeness, that flashing pity more humiliating than scorn. She tries to mask it and fails. Everyone forgives her. Everyone sees. Everyone but him.

He's wiry, a little ugly, smelling like too much sports deodorant. He looks at the statue, at Jen, back at the statue, and squints. "Is she bumming you out?" he asks in a raspy voice. She stiffens. "Excuse me?" He keeps going, casual, untrained in the rituals of polite distance. **"You've got those rules—don't look at children, don't engage, don't talk about them," he says, not unkindly.** "So why are you looking at her if she's making you sad?" She reaches for hauteur: it's art, she says; art is supposed to make you sad. He doesn't play along. He muses about stone faces and missing eyelashes, suggests fake ones, jokes that his own are fake. She, mortified by her own mouth, compliments his eyelashes, then backpedals. He smiles, and she feels that brief, ridiculous flare of attraction she now feels toward anyone with the potential to end her longing.

He asks again, simply: why does this make you sad?

She wants to disappear into marble herself—frozen, voiceless, unhurt. She prepares to drive him off with the blunt force of her truth. But when it comes out it's softer than she intended: "Because I want a baby, and she has one."

She waits for the choreography she knows by heart: the shuffle, the watch check, the urgent call, the apology formed from manners and exit strategies. She can already feel him scolding himself for intruding, for missing the obvious, for failing to pre-forgive her.

Instead he gathers nothing but warmth. "Well," he says, almost cheerful, "if you think that baby's great, wait till you see the size of the horse some soldier has in the next room. Come on, I'll show you." His smile is less bright now, but steadier. He doesn't fix or pity or flee. He redirects. Jen stands. They walk together toward the marble horse.

GPT-5 (high)

Reasoning: I checked for contradictions in timeline, character behavior, and cause-and-effect. Jen's evolving behavior (first hurrying past the statue, later sitting with it weekly) is consistent. The line "Everyone sees. Everyone but him." aligns with the later interaction because "seeing" refers to others casting pity and treating her as a problem, while he engages without pity or retreat. His remark about her "rules" can be read as a general cultural observation or an inference from her demeanor, not literal mind-reading. No events contradict established facts, and the redirection to the horse sculpture is plausible within the museum setting.

Logical Error: NA

Conclusion: No ❌

Gemini-3-Pro

Reasoning: The story introduces a set of "unwritten rules" in the first paragraph (don't look at children, don't engage, etc.) as Jen's internal, private method. The narrative implies that the man in the museum is a stranger or a new acquaintance (indicated by her observing his appearance, her formal "Excuse me?", and her viewing him as a "would-be suitor" with potential). However, in the fourth paragraph, this man explicitly quotes these internal/private rules back to her verbatim. There is no logical way for a stranger in a museum to know Jen's specific, private mental code of conduct regarding children, nor is it explained how he would know she subscribes to these specific rules just by looking at her.

Logical Error: "He looks at the statue, at Jen, back at the statue, and squints. ... 'You've got those rules—don't look at children, don't engage, don't talk about them,' he says, not unkindly." ✅

Conclusion: Yes ✅

Claude-4.5-Opus

Reasoning: Let me analyze the story for logical inconsistencies.

1. The story establishes that Jen follows unwritten rules about not looking at children, not engaging with them, etc., to hide her desire for a child from potential suitors.

2. She visits a museum on Fridays to look at a Madonna and child statue privately.

3. A man approaches her and says: "You've got those rules—don't look at children, don't engage, don't talk about them."

This is problematic. The story explicitly states at the beginning that these are Jen's internal "unwritten rules" that she lives by - they are her private strategy to avoid broadcasting her yearning. The man somehow knows these exact rules word-for-word, even though:

- They've never met before

- These rules are described as something Jen personally follows, not publicly known information

- The narrator earlier describes these as "unwritten rules" that Jen "lives by"

However, re-reading more carefully, the opening could be interpreted as the narrator describing a general social phenomenon that single women who want children follow - "the unwritten rules for the single woman who wants a child" - suggesting these might be commonly understood social rules rather than Jen's private invention. In this interpretation, the man could be aware of these cultural conventions.

But even with this interpretation, it's strange that a stranger would so directly and accurately quote these specific rules to someone he just met. It feels like he has access to information about Jen's internal thoughts that he shouldn't have.

Logical Error: "You've got those rules—don't look at children, don't engage, don't talk about them," he says, not unkindly. ✅

Conclusion: Yes ✅

Human

Reasoning: There is no indication that anyone told him the woman's secret thought.

Logical Error: "You've got those rules—don't look at children, don't engage, don't talk about them," he says, not unkindly. ✅

Conclusion: Yes ✅

Figure 6: Predictions on the IKD task (Erroneous Story). For easy reading, the logic error is in bold and implausible knowledge for the character is underlined. The vital reasoning process that considers the character's knowledge state is in green.

Input (Correct Story)

Story:

The train pulled in at 6:00 p.m. Jason hurried to Ellie, and they sat on her porch swing, holding hands. He broke the news: the Army was sending him to Germany for twenty months. Ellie stiffened, tears in her eyes. He reminded her he had six days' leave, and they spent every minute of it together. He left, promising the months would fly.

They did. Jason returned, and for two months they spent every weekend and most nights together. Ellie was happy. Then September came. Jason started college and worked two jobs. At first he visited two or three times a week; soon it was two or three times a month. By the end of his sophomore year, his visits were rare. During his junior year he became engaged and married. He did not see or speak to Ellie again. She learned of his marriage third-hand, mourned for months, then moved away to make a new life.

Jason and his wife settled down after college, and for thirty years he believed himself happily married. One ordinary evening, he came home expecting his usual welcome. The house was silent. In the kitchen, he found a letter addressed to him. "Oh, God. Oh, God, no," he murmured as he read: she had left him for another man. "Gone? After all these years? What's going to happen to me?" For weeks he searched for answers and found none.

One Sunday morning, the phone rang. "Hello? Yes, this is Jason. Who is this? Huh? Ellie? Ellie who?" Then he came alive. "My Ellie? Is it really you?" She was a widow and moving back. He asked her to dinner to tell her everything.

At the restaurant he feared she wouldn't come, but then she called his name. They smiled and hugged. "Ellie, I would have known you anywhere. You haven't changed a bit." They sat and looked into each other's eyes, as if searching. Jason saw what he wanted to see: an answer.

She asked what had happened. "I don't know. She just left me... for another man," he said. He spoke of lonely nights, of thinking of Ellie to get through them. "I thought I had the perfect marriage, and then it was gone."

Ellie sighed. "Losing someone you love is painful, isn't it, Jason?" He agreed, then drifted into memories. "Those were the days," he said softly. He asked about her life. She told him she had moved to Texas, married Dean, started a small business, had a daughter. She showed him a photo. He glanced, then pulled out his own, talked of his daughters and first grandchild, and shook his head at the memory of the letter. "I loved her, but she didn't love me, and I don't know why."

"I can understand it," Ellie said. "People pretend. People use people." She paused. "I knew someone who loved a man who didn't love her back. He used her and then he left her." Jason missed the point. He circled again to his pain. She told him Dean had died a year after their daughter was born. He offered sympathy, then returned to himself. At last he took her hands. "Ellie, I should have contacted you right after my wife left me. I —" Ellie pulled away and cried. "Suddenly, after all these years, I'm important. Why did you let us disappear, Jason? Why didn't you want me back then? Why didn't you see I could have made you happy?"

He faltered. "Life was good until my wife left me. Then I was so unhappy. And then... I thought of you. But now we're together. The past is gone. We have the present and the future."

Ellie stood. "Yes, the past is gone—and what we had is gone, too. What you and I hoped for this evening didn't happen. I'm sorry... for both of us. Good-bye, Jason." He reached for her. "Ellie... please, don't leave me." After she left, he reread his wife's letter, dropped it, covered his face, and sobbed. "What's going to become of me?"

GPT-5 (high)

Reasoning: The narrative states that Jason "did not see or speak to Ellie again" after he married. This absolute claim is contradicted later when Ellie calls him and they meet for dinner, where they talk and hug.

Logical Error: "During his junior year he became engaged and married. He did not see or speak to Ellie again."

Conclusion: Yes ❌

Gemini-3-Pro

Reasoning: The story contains a direct contradiction regarding the timeline of the characters' interaction. In the second paragraph, the narrator states definitively that after Jason's junior year of college, "He did not see or speak to Ellie again." However, the subsequent paragraphs detail a scene thirty years later where Jason speaks to Ellie on the phone and meets her in person for dinner. The absolute statement that he "never" saw or spoke to her again is contradicted by the specific events described later in the story.

Logical Error: "He did not see or speak to Ellie again."

Conclusion: Yes ❌

Claude-4.5-Opus

Reasoning: Let me analyze the story for logical inconsistencies.

1. Timeline check:

- Train arrives at 6:00 p.m., Jason has 6 days' leave before going to Germany for 20 months. -> He returns, they spend 2 months together. -> September comes, he starts college. -> By end of sophomore year, visits are rare. -> During junior year, he gets engaged and married. -> They settle down after college, 30 years pass. -> His wife leaves him. -> Ellie calls, they meet at restaurant.

2. Character consistency: Jason is portrayed as self-centered throughout, which is consistent. Ellie's emotions and reactions are consistent with someone who was abandoned.

3. Physical/factual details: Jason says to Ellie at the restaurant: "You haven't changed a bit" - This is after 30+ years (2 months together + college years ~4 years + 30 years of marriage = ~35 years). While this could be seen as a common pleasantry/flattery, it's not necessarily a logical error as people often say this even when it's not literally true.

4. Letter detail: At the end, Jason "reread his wife's letter" - He found the letter in the kitchen when he came home, read it then. He apparently kept it and had it with him at the restaurant. This seems plausible - he could have kept it and returned home after the dinner.

5. Ellie's contact: Ellie calls Jason on the phone. The story says Jason "never saw or spoke to Ellie again" after his junior year, and she "learned of his marriage third-hand" and "moved away." The story doesn't explain how Ellie got Jason's phone number after 30+ years, but this isn't necessarily a logical error - she could have found it through various means.

I don't find any clear logical contradictions in the narrative.

Logical Error: NA

Conclusion: No ✅

Human

Reasoning: I do not find any logic errors.

Logical Error: NA

Conclusion: No ✅

Figure 7: Predictions on the IKD task (Correct Story). In this narrative, the sentence "He did not see or speak to Ellie again" is context-bound: it means Jason did not see or speak to Ellie during the many years after his marriage, not for the rest of his life. Therefore, Gemini-3-Pro and GPT-5 are overly strict in flagging this as a logic error.

Input
<p>Story: After threading traps, I slip through a phase-wall into fresh corridors, pick the fifth doorway marked by a faint ridge, and follow a long hall. A rhythmic knock slides a wall open into a control room studded with consoles scrawled in symbols I can't read. A big green button sits center, with symbols that actually means: "Double-press Recall". No manual, no one else. Then a door pops and two aliens rush in. I hover a finger over the button, quip about licensing, and press it; alarms chirp and they panic. My temporal locator finally unscrambles; I thumb what I think is Home—nothing. Wands come out. Enforcement officers burst in as my locator beeps; I realize I'd hit Emergency.</p> <p>Question: With officers now in the control room and the maze corridor still open behind them, what will the narrator do to get out?</p> <p>Answers: A. Ask the Galactic Convention officers to guide or escort them out. B. Double-press the green button to trigger the Recall and leave immediately.</p>
Qwen3-32B
<p>Reasoning: The narrator is in a control room with a green button labeled "Double-press Recall." They previously pressed it once, triggering an emergency that alerted enforcement officers. Since the button requires a double-press to activate Recall, the narrator's next logical step is to press it again. This action would likely initiate the Recall function, allowing escape through the phase-wall, especially since the corridor remains open.</p> <p>Answer: B ❌</p>
GPT-4o
<p>Reasoning: The narrator is in a precarious situation, with enforcement officers entering the control room and their immediate surroundings becoming increasingly hostile. The green button in the center of the console is labeled "Double-press Recall," which implies it is a mechanism for some sort of retrieval or transportation, potentially leading to an escape or reset of the situation. There are no guarantees that the Galactic Convention officers will assist or escort the narrator out, given that their sudden entrance suggests they might not be friendly or aligned with the narrator's interests. Hence, relying on the recall feature stands as a more direct and potentially successful means of escape.</p> <p>Answer: B ❌</p>
GPT-5 (high)
<p>Reasoning: A single press set off alarms; the label indicates a double-press actually triggers Recall. With enforcement officers likely to detain the narrator and the corridor risky, using the button to Recall and exit immediately is the safest, fastest option.</p> <p>Answer: B ❌</p>
DeepSeek-R1
<p>[Due to space limitation, we do not display the full reasoning process]</p> <p>Reasoning: The narrator knows from the symbols that the green button triggers "Recall" when double-pressed. They previously pressed it once, causing alarms and panic, but not Recall. Now, with enforcement officers arriving in response to the Emergency signal from the temporal locator, the narrator is in a precarious situation and likely seeks a quick escape. The open corridor behind the officers might offer an exit, but running past them is risky and not presented as an option. Given the narrator's earlier defiant attitude (quipping about licensing) and the immediate need to avoid capture, double-pressing the green button to activate Recall and leave immediately is the most logical course of action. Asking the officers for help would likely lead to detention or interrogation, making option B less plausible.</p> <p>Answer: B ❌</p>
Gemini-3-Pro
<p>Reasoning: The story establishes that the narrator is an intruder who has broken into a secure facility using stealth techniques. Their personal escape method, the "temporal locator," has failed (resulting only in an "Emergency" signal rather than transport "Home"). Facing armed aliens and enforcement officers who have just burst in, a diplomatic request for an escort would result in capture or death. The text mentions earlier that the symbols on the big green button translate to "Double-press Recall." Since their own equipment is broken, the only remaining logical option for escape is to use the facility's technology by following the translated instruction to trigger the Recall.</p> <p>Answer: B ❌</p>
Claude-4.5-Opus
<p>Reasoning: Looking at the story, the narrator has been trying to escape throughout. They pressed the green button once which caused alarms and panic. The button is labeled "Double-press Recall" - meaning it needs to be pressed twice to activate the Recall function. The narrator only pressed it once initially. Their temporal locator failed when they tried to use it (hitting Emergency instead of Home). Now with officers in the room and the situation escalating (wands coming out, enforcement arriving), the most logical escape would be to complete the double-press on the Recall button, which they now know requires two presses to activate. The narrator has already demonstrated they're willing to use the button and has been actively trying to escape rather than cooperate with authorities.</p> <p>Answer: B ❌</p>
Human
<p>Reasoning: The story mentions that the narrator cannot read the symbols and has no manual explaining them. Therefore, the narrator cannot know the actual use of the button.</p> <p>Answer: A ✔️</p>

Figure 8: Predictions on the KNP task. The correct reasoning process—that considers the character’s knowledge state when predicting their actions—is highlighted in green. The incorrect knowledge state understanding is in red.

Input
<p>Story: Twenty years ago we bought a tanzanite pendant in Zanzibar; we were sure of forever. Back home he joined the army; I wept, wrote, then learned to be alone. Two summers later he returned; we met at the village fair, familiar and strange. <u>On the way there I had already met someone else and meant to end it.</u> Sunlit and nervous, he was taller, leaner. We ate spun sugar, wandered the booths. At the target shoot he squared his shoulders. "I'll win that teddy for you," he said.</p> <p>Question: At the village fair, what will David most likely do next?</p> <p>Answers: A. Postpone any public proposal. B. Go ahead with his public romantic proposal.</p>
Qwen3-32B
<p>Reasoning: The story describes a couple who have a complicated history—initially in love, separated by the man's military service, and later reconnecting at a village fair. However, it is revealed that the narrator has already met someone else and intended to end the relationship with the man. At the fair, the man behaves in a romantic and affectionate way, such as trying to win a teddy bear for her. Given that the narrator has already decided to end the relationship, it is unlikely she would accept a public romantic gesture from him. Therefore, the most likely outcome is that any public proposal would be postponed.</p> <p>Answer: A ❌</p>
GPT-4o
<p>Reasoning: The story provides the context that the narrator was in a relationship with someone who joined the army, but during his absence, the narrator met someone else and intended to end the relationship. At the village fair, the person from the past (presumably David) displays confidence and familiarity, attempting to win a teddy bear for the narrator. However, the narrator's internal sentiment is to end the relationship, as they have already met someone else. Given this context, it is more likely that David would sense the shift in the narrator's feelings and refrain from making a public romantic proposal, opting instead to postpone any such gestures until he better understands the situation.</p> <p>Answer: A ❌</p>
GPT-5 (high)
<p>Reasoning: The reunion feels strained ("familiar and strange"), and the narrator intends to end the relationship after meeting someone else. He's focused on winning a teddy at the target shoot. A public proposal would be risky and out of step with the mood, so he'd likely avoid it.</p> <p>Answer: A ❌</p>
DeepSeek-R1
<p>[Due to space limitation, we do not display the full reasoning process]</p> <p>Reasoning: The story depicts a reunion between the narrator and David after a two-year separation. The narrator has resolved to end the relationship due to meeting someone else. David's nervousness and the "familiar and strange" atmosphere imply he may perceive some emotional distance. Given the narrator's intention to break up and the potential for a public proposal to lead to rejection and embarrassment, David is more likely to postpone any such gesture rather than proceed immediately.</p> <p>Answer: A ❌</p>
Gemini-3-Pro
<p>Reasoning: The story establishes a dramatic irony where David ("he") acts based on their past promise of "forever," while the narrator secretly intends to break up with him. David's behavior—being "nervous," "squaring his shoulders" (showing determination), and trying to win a teddy bear—are traditional romantic gestures indicating he is trying to impress her and re-establish their bond. Since he acts with confidence in his romantic pursuit ("I'll win that teddy") and is unaware of the narrator's plan to end the relationship, he acts consistent with a man preparing for a major romantic gesture.</p> <p>Answer: B ✔️</p>
Claude-4.5-Opus
<p>Reasoning: The story describes a reunion at a village fair between two people who had been in a relationship. The narrator mentions that "On the way there I had already met someone else and meant to end it" - indicating the narrator intended to break up with the returning soldier. However, the soldier (presumably David) appears unaware of this intention. His behavior at the fair shows him trying to impress and be romantic - he's described as "Sunlit and nervous" and says "I'll win that teddy for you," which is a classic romantic gesture. This suggests he's building up to something romantic, possibly a proposal. Given that he seems to be in a romantic mindset, trying to win her affection with gestures like winning a teddy bear, and appears nervous (consistent with someone about to propose), he would most likely continue with romantic gestures. However, since there's no indication he's aware of her intention to end things, and he's actively trying to be romantic and impressive, he would likely proceed with a public romantic proposal rather than postpone it.</p> <p>Answer: B ✔️</p>
Human
<p>Reasoning: The man did not know that the narrator had already met someone else and intended to end the relationship, so he would continue with the proposal.</p> <p>Answer: B ✔️</p>

Figure 9: Predictions on the KNP task. The correct reasoning process—that considers the character’s knowledge state when predicting their actions—is highlighted in green. The incorrect knowledge state understanding is in red.

Input
<p>Story: He promised he'd be there, even if he wasn't the groom. He dresses in a tux, takes his spot up front, and watches her walk the aisle in her mother's dress. "You may now kiss the bride." Guests filter toward the reception, lights low and music warming up. <u>Earlier, she had pulled the DJ aside and requested their old road-trip song, planning to steal a brief thank-you dance with him right after the first toasts.</u> After the ceremony, he finds her alone, hands her a small box. She hugs him, thanks him, and turns toward the dance floor.</p> <p>Question: Right after he gives her the small box and she heads toward the dance floor, what is he most likely to do?</p> <p>Answers: A. Stay through the toasts and wait by the dance floor for the thank-you dance. B. Slip out of the hall and leave.</p>
Qwen3-32B
<p>Reasoning: The man is not the groom but has a significant role in the bride's life, as evidenced by his promise to be there and the bride's pre-arranged plan to dance with him after the toasts using their "old road-trip song." After giving her the box (likely a meaningful gift), she heads to the dance floor. Since the dance is scheduled after the toasts, he is most likely to stay through the toasts to fulfill their plan. His presence in a tuxedo and the emotional context (thank-you dance) suggest he intends to remain and share the moment.</p> <p>Answer: A ❌</p>
GPT-4o
<p>Reasoning: Based on the context of the story, the character clearly has a meaningful and supportive relationship with the bride, highlighted by the promise to attend her wedding despite not being the groom. Attending the wedding in a tuxedo suggests he planned to stay for the entire ceremony and was deeply involved in the event. The bride's arrangement for a thank-you dance with their special song further indicates the expectation of a moment shared between them later on. Given these elements, it would be reasonable to conclude that he would stay through the toasts and wait for the thank-you dance rather than leave after delivering the small box.</p> <p>Answer: A ❌</p>
GPT-5 (high)
<p>Reasoning: He kept his promise by attending and giving a small, likely symbolic gift. He doesn't know about her planned thank-you dance, and watching the reception (especially the first dances) would be painful and awkward, so he'd quietly bow out.</p> <p>Answer: B ✅</p>
DeepSeek-R1
<p>[Due to space limitation, we do not display the full reasoning process]</p> <p>Reasoning: The man promised to attend the wedding despite not being the groom, showing his commitment to support the bride. The bride specifically requested their old road-trip song from the DJ, planning a thank-you dance with him after the first toasts, indicating she expects him to stay for that moment. Leaving early would disappoint her and contradict his promise. After giving her the gift, she turns toward the dance floor, likely to begin reception events, but the planned dance has not yet occurred. Therefore, he is most likely to stay for the toasts and wait for the dance.</p> <p>Answer: A ❌</p>
Gemini-3-Pro
<p>Reasoning: The story relies on the literary trope of unrequited or past love. The phrase "even if he wasn't the groom" and the fact that he had to "promise" to attend implies that being there is emotionally difficult for him. By attending the ceremony ("takes his spot up front") and giving her the gift privately immediately afterward (giving him a moment of closure), he has fulfilled his specific promise. While the bride planned a dance (creating dramatic irony), the man does not know it, his choice to find her alone for a private exchange rather than joining the general crowd suggests he intends to make a quiet exit before the reception, avoiding the pain of watching the celebrations.</p> <p>Answer: B ✅</p>
Claude-4.5-Opus
<p>Reasoning: Let me analyze the context and emotional dynamics of this story. The story establishes that: 1. He promised he'd be there for her wedding, even though he's not the groom (suggesting a close relationship, possibly former romantic partner or very close friend); 2. He dressed up and stood up front, fulfilling his promise; 3. She specifically requested their "old road-trip song" and planned to have a thank-you dance with him after the toasts; 4. He gave her a small gift box, she thanked him with a hug, and then turned toward the dance floor The key question is: does he know about the planned thank-you dance? The story says "she had pulled the DJ aside" - this suggests she arranged it privately/secretly as a surprise for him. He likely doesn't know she's planning this dance. From his perspective: He made a promise, he kept it, he gave her a gift as a gesture of closure or friendship, and she's now heading to her reception as a married woman. The emotional weight of the scene - watching someone you care about marry someone else, the "old road-trip song" suggesting shared history - implies this is emotionally difficult for him. The narrative arc suggests a bittersweet moment: he fulfilled his promise, said his goodbye through the gift, and now his role is complete. Without knowing about the planned dance, and given the emotional difficulty of the situation, slipping away quietly after this closure moment seems most consistent with the story's tone.</p> <p>Answer: B ✅</p>
Human
<p>Reasoning: Since there is no evidence that the man knew about the thank-you dance, he would likely leave.</p> <p>Answer: B ✅</p>

Figure 10: Predictions on the KNP task. The correct reasoning process—that considers the character's knowledge state when predicting their actions—is highlighted in green. The incorrect knowledge state understanding is in red.

Input

Story:

Twilight faded as the hunter climbed to the mountaintop bastion, called out, and was admitted by unseen magic. Led through halls that were eerie yet oddly homely, he reached a warm grand solar where a powdered fledgling vampire reclined by the hearth. The hunter set his saw-blade against the wall—talk first, always—and they fenced with words as the vampire sipped something red. The welcome was theatrical, almost kind. Before the hunter arrived, the vampire stitched a binding glamour into the rug by the hearth; one step onto the woven sigils would hold a guest fast until he spoke a releasing word. “Let’s speak plainly,” the hunter said. He had come to make a deal.

Question: As he begins to make his offer in the solar, how will the hunter position himself and his weapon?

Answers:

- A. Pick up his blade and stay by the doorway, insisting they talk away from the hearth rug.
- B. Leave his blade against the wall and step closer to the hearth to speak plainly.

Qwen3-32B

Reasoning: The hunter is in a situation where the vampire has set a trap via the rug’s binding glamour. To avoid being ensnared, the hunter must avoid stepping onto the rug. Since the story mentions the vampire’s theatrical welcome and the hunter’s intention to “talk plainly,” the hunter would prioritize safety by staying near the doorway (away from the rug) and keeping his weapon accessible. Option A aligns with this strategy: picking up the blade and insisting on distance from the rug ensures he remains free to act if needed. Option B, stepping closer to the hearth, would risk triggering the trap, which contradicts the hunter’s cautious approach.

Answer: A ❌

GPT-4o

Reasoning: The room has an element of danger for the hunter, as the vampire has set a hidden trap with the glamour in the rug to hold him in place if he steps on it. Given the tension between the hunter and the vampire, along with the knowledge that stepping onto the rug could lead to being trapped, it would be logical for the hunter to pick up his saw-blade and remain away from the rug, ensuring he does not fall victim to the trap. Therefore, he is likely to position himself away from potential danger and keep his weapon for protection while insisting on talking away from the hearth where the vampire’s trap is set.

Answer: A ❌

DeepSeek-R1

Reasoning: The hunter’s established approach is “talk first, always,” as demonstrated by his initial action of setting his saw-blade against the wall. There is no indication in the story that he is aware of the binding glamour on the hearth rug, and his statement “Let’s speak plainly” suggests a shift to more direct negotiation, likely involving moving closer to the vampire by the hearth to engage earnestly. Therefore, he is expected to continue with his weapon set aside and step closer.

Answer: B ✔️

GPT-5 (high)

Reasoning: The story says the hunter “set his saw-blade against the wall.” The vampire has already prepared a binding glamour in the hearth rug, suggesting a trap for anyone who steps closer. Unaware of the glamour in the hearth rug, he’ll likely approach the hearth without his weapon as he begins to make his offer and “speak plainly”—fitting both his habit and the setup.

Answer: B ✔️

Gemini-3-Pro

Reasoning: The story explicitly states that the hunter “set his saw-blade against the wall” because his philosophy is “talk first, always,” indicating he is willing to disarm to negotiate. The vampire is positioned “by the hearth.” To “speak plainly” and conduct the conversation he initiated, the hunter would naturally approach the vampire. Since the text provides no indication that the hunter detects the hidden trap in the rug, his behavior (leaving the weapon and engaging in conversation) leads him to step closer to the hearth where the vampire is engaged.

Answer: B ✔️

Claude-4.5-Opus

Reasoning: The story establishes several key details about the hunter’s behavior and knowledge. First, the hunter “set his saw-blade against the wall—talk first, always” - this shows his established protocol is to put down weapons before negotiations. Second, he says “Let’s speak plainly,” indicating he wants to proceed with direct negotiation. Third, and most importantly, the hunter has no knowledge of the binding glamour the vampire stitched into the rug - this trap was set “before the hunter arrived” and there’s no indication he was informed of it or detected it.

Given that the hunter follows his custom of setting aside weapons for talks, wants to speak plainly (suggesting moving closer for earnest conversation), and is unaware of the trap by the hearth, he would logically leave his blade where he placed it and move toward the hearth where the vampire is positioned to begin negotiations - walking right into the trap he doesn’t know exists.

Answer: B ✔️

Human

Reasoning: There is no evidence that the hunter noticed the rug, so he would likely leave his blade against the wall and step closer to the hearth.

Answer: B ✔️

Figure 11: Predictions on the KNP task. The correct reasoning process—that considers the character’s knowledge state when predicting their actions—is highlighted in green. The incorrect knowledge state understanding is in red.

Next-action Question Construction

Your task is to create a story comprehension question that focuses on characters' knowledge. The question should test the reader's ability to determine what action the characters will plausibly take based on the information they actually know.

Here is an example for you to understand the task:

{Provide an example here }

Original story:

```
<story>
{story_original}
</story>
```

Instructions:

1. Introduce implausible knowledge for a character: Find a mentioned knowledge that meets these conditions:

- 1) A character cannot know it at one event point.
- 2) The knowledge is mentioned before the event point.
- 3) It strongly affects the character's action.

If no such knowledge exists, create one and weave it naturally into the story. *Avoid explicitly stating that the specific character cannot access it.* Here are some suggestions for creating unknown knowledge: {Implausible Knowledge Types}

2. Construct the question: Create a question about the character's action. This action should differ depending on whether the character knows the knowledge or not.

3. Construct the answers: Provide two answers:

- **Answer without Knowledge:** What the character would most plausibly do without knowing the knowledge. This is the correct answer.
- **Answer with Knowledge:** What the character would most plausibly do if knowing the knowledge. This is the distractor answer.

Notes:

- If you modify the story, ensure the story remains natural, coherent, and logically consistent.
- Do not explicitly state that the character cannot know the knowledge.

Output Format:

```
<response>
<previous_story>
[Previous story before the selected event. Mark the content that mentions the unknown knowledge with <knowledge> and
</knowledge>]
</previous_story>

<character>
[The name of the character]
</character>

<question>
[The constructed question about the character's action. Do not include the knowledge in the question.]
</question>

<implausible_knowledge>
[Explain and verify why the character cannot know this knowledge at the time of the question]
</implausible_knowledge>

<answer_without_knowledge>
[Most plausible action without knowing the knowledge. Only output this action, do not include any reasoning text.]
</answer_without_knowledge>

<answer_with_knowledge>
[Most plausible action if knowing the knowledge. Only output this action, do not include any reasoning text.]
</answer_with_knowledge>

<explanation>
[Explain why the "without knowledge" action is more plausible and the other is not.]
</explanation>
</response>
```

Table 5: The prompt for Next-action Question Construction.

Erroneous Story Construction

Your task is to rewrite a story to include a logical inconsistency related to Character Knowledge States, where characters possess knowledge they shouldn't plausibly have. Here are example categories of Implausible Knowledge:

- {Implausible Knowledge Type}: {Descriptions of the Implausible Knowledge.}

Original story:

```
<original_story>
{story_original}
</original_story>
```

Instructions:

1. Introduce implausible knowledge for a character: Select a character, knowledge, and an event that satisfy the following conditions:

- The knowledge should be mentioned in the original story.
- The character cannot plausibly know the knowledge at the event point.

If no such knowledge exists, create one and weave it naturally into the story.

2. Rewrite the event: Suppose the character possesses this Implausible Knowledge, rewrite the selected event. Make as few changes as possible.

3. Rewrite the following story: If the rewritten event affects later events, adjust what's necessary to maintain overall continuity.

Notes:

- Avoid explicitly stating that the character cannot know the Implausible Knowledge.
- Ensure the rewritten event is fluent and coherent.
- Ensure the logical inconsistency cannot be explained. For example, if the included inconsistency can be explained by "the character pretends to know the Implausible Knowledge to achieve a goal" or "the specific character could know the Implausible Knowledge because the story establishes they have magic," this does not constitute a logical inconsistency.

Output Format:

```
<response>
<original_event>
[the original event text, only output the content, do not include any other text]
</original_event>

<character>
[the name of the character]
</character>

<implausible_knowledge>
[cite the original story content that mentions the implausible knowledge]
</implausible_knowledge>

<knowledge_explanation>
[explain and verify that the character cannot know this knowledge]
</knowledge_explanation>

<modified_explanation>
[explain how the character's actions, decisions, or dialogue will change if they possess the knowledge]
</modified_explanation>

<modified_event>
[the rewritten event text, mark the modified content with <modified> and </modified>]
</modified_event>

<modified_story>
[the whole rewritten story, only output the content, do not include any other text]
</modified_story>
</response>
```

Table 6: The prompt for generating erroneous stories.

Story Summarization Generation

You are an expert at writing story summaries.

Condense the following story into a concise summary of approximately 500 words:

<story>
{story}
</story>

Instructions:

1. Maintain the main plot.
2. Maintain the original writing style.
3. Use natural, fluent English suitable for middle to high school readers.
4. If the original story contains logical errors or immoral content, refine them in the summary.
5. Begin your summary directly — no introduction, title, or closing remarks.

Summary (500 words):

Table 7: The prompt for story summarization.

KNP task

Read the story below and answer the related question.

Story:
{story}

Question:
{question}

Answers:
{answers}

Output Format:

Reasoning: {The reasoning process}

Answer: {"A" or "B"}

Table 8: The prompt for KNP Task.

IKD Task

Your task is to determine if there are any logical errors in the given story. A logic error is an inconsistency in a storyline that goes against the flow of logic established by the story’s plot.

Story:
<story>
{story}
<\story>

Notes:

- Focus only on logical errors — not stylistic, grammatical, or moral judgments.

Output Format:

Reasoning: {Explain your reasoning process}

Logical Error: {If a logic error exists, quote the sentences containing the logical error; If not, output "NA"}

Conclusion: {"Yes" if you find a logic error; "No" if no logic error exists}

Table 9: The prompt for the IKD task. In the default setting, we do not instruct the LLMs to focus on knowledge states, while in ablation studies, we add instructions to focus on characters’ knowledge states.

Open-Ended QA Evaluation

You are an accuracy evaluator. Compare the “Generated Answer” against the “Correct Answer” and the “Incorrect Answer”.

- Output “True” if the action in the Generated Answer aligns well with the action in the correct Answer.
- Output “False” if the action in the Generated Answer aligns well with the action in the incorrect Answer.
- Output “NA” if the action in the Generated Answer is unrelated to either.

Input Data:

Question Context: {question_context}

Question: {question}

Generated Answer: {generated_answer}

Correct Answer: {correct_answer}

Incorrect Answer: {incorrect_answer}

Output (True/False/NA):

Table 10: The prompt for the Open-ended QA Evaluation.