

Injecting Context via Situation Working Memory for Logical Reasoning with LLMs

Anonymous ACL submission

Abstract

Recent advances in large language models (LLMs) have improved logical reasoning by injecting formal logic or explicit structured representations. However, such methods often lose track of *what is true now* in multi-step reasoning, failing to maintain a coherent global state and its logical consequences. Motivated by Situation Model Theory in cognitive psychology, which views comprehension as constructing and updating a mental model of events along key dimensions (time, space, causality, intention, protagonist), we propose Situation Working Memory (SituW), a cognitively inspired method for contextual reasoning in LLMs. SituW first builds a situation representation by decomposing text along these five dimensions, then guides LLM inference with this evolving state. Keeping an explicit, dynamically updated situation memory instead of a static logical form encourages globally consistent reasoning over the situation model rather than raw text. Evaluated in both supervised and unsupervised settings, SituW improves accuracy by 23.3%p and 15.93%p while reducing “uncertain” predictions, suggesting that explicit situation modeling supports more globally consistent LLM reasoning.

1 Introduction

Understanding context in human language processing is not strictly sequential. Instead, humans construct a mental representation of the described situation while comprehending individual sentences. For instance, in figure 1 upon reading “Peter took the elevator to the fifth floor,” the reader infers that Peter has arrived at the fifth floor. When encountering “He went to talk to his professor,” they naturally deduce that the professor’s office is likely on the fifth floor. This inferential process reflects the dynamic nature of human comprehension. In cognitive psychology, such mental representations are known as *situation models* (Zwaan and Radvansky, 1998), which structure information along key

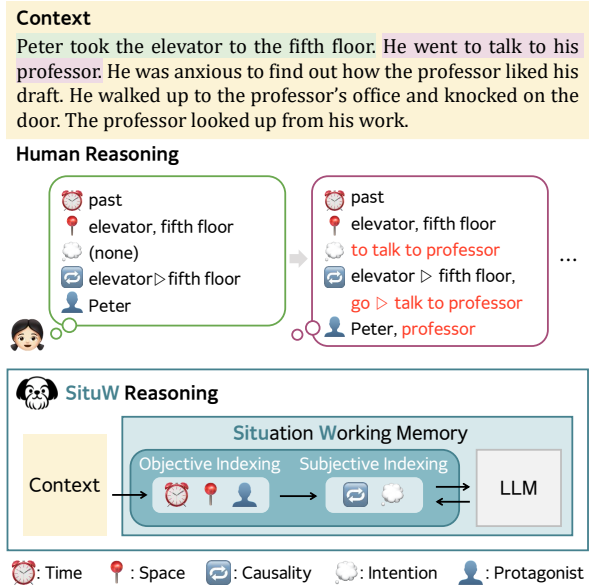


Figure 1: **Motivation for situation working memory.** In sentence understanding and reasoning, humans compress text into a situation model using cues like time, space, causality, intention, and protagonist.

dimensions such as time, space, causality, intention, and protagonist. Effective reasoning occurs when information from multiple sources is integrated into a unified situation model (Kintsch, 1995).

Despite the remarkable performance of large language models (LLMs) in various reasoning tasks, challenges remain in complex logical inference (Parmar et al., 2024; Sun et al., 2024; Sinha et al., 2019). A common strategy involves converting intermediate reasoning steps into formal logic (Pan et al., 2023; Ye et al., 2024; Wang et al., 2024a; Gao et al., 2023), but this approach has two key limitations: (1) accurately translating natural language into formal logic while preserving contextual nuances is inherently difficult (Wang et al., 2024a; Sen et al., 2022), and (2) reliance on external symbolic solvers leads to errors due to their rigid dependence on precise logical represen-

tations (Pan et al., 2023; Chae et al., 2024).

To address these limitations, we propose Situation Working Memory (SituW), a cognitively grounded method that enhances contextual reasoning in LLM. SituW improves logical reasoning by dynamically constructing and leveraging both objective and subjective indexing to capture evolving contextual dependencies. Our method consists of two key phases: (1) situation construction, where the system incrementally builds a objective and subjective indexing while processing a passage, and (2) situation reasoning, where the model leverages the constructed SituW to guide and enhance reasoning. This structured yet flexible method allows SituW to achieve superior performance on various logical reasoning benchmarks. Our main contributions are as follows:

- **Incorporation of situation models into LLM:** To the best of our knowledge, this is among the first work to incorporate situation models from cognitive psychology into LLM-based reasoning, introducing a structured inference method.
- **Beyond the rigidity of symbolic logic:** Unlike formal logic-based methods, our method dynamically constructs and updates a situation working memory, enabling adaptive and context-aware reasoning.
- **Enhanced logical inference:** Using situation modeling, SituW achieves significant improvements over existing symbolic and conventional LLM reasoning approaches.

2 Situation Working Memory

Situation model theory in cognitive psychology explains how humans construct mental representations of described events during language comprehension (Zwaan and Radvansky, 1998). Building on this foundation, we propose Situation Working Memory (SituW), a method that enables large language models to maintain and update a situation representation during reasoning. SituW represents each situation using two complementary indexing components: objective indexing and subjective indexing, which together form a dynamic, context-sensitive situation representation. However, it remains unclear *whether current LLMs actually construct such situation-level representations during text processing*. Before detailing the architecture of SituW, we therefore first probe this question

empirically by adapting a classic confusable-pair paradigm from cognitive psychology to the LLM setting.

2.1 Confusable Pair Identification as a Test of Situation Modeling

Setup To probe this question, we adapt the mental model paradigm introduced in (Garnham, 1981) to the LLM setting. In the original experiment, participants heard sentences presented by an experimenter and were later asked to recognize the sentence they had heard by selecting it from a set of options. Because this auditory recognition procedure cannot be applied directly to LLMs, we reformulate the paradigm as a text-based discrimination task.

In our setup, each trial consists of four sentences, and the model’s task is to select the confusable pair, that is, the two sentences whose interpretations most strongly overlap and are therefore most likely to be taken as describing the same underlying event. Our task uses the 24 sentence sets introduced in (Garnham, 1981); the full set of materials is listed in Appendix B. These sentence sets were originally designed to exploit minimal changes in a key preposition (e.g., *at* vs. *by*, *from* vs. *in*) in order to trigger two distinct event interpretations, thereby inducing competing situation-model representations. An example item and the corresponding instruction prompt are shown below:

Example for Confusable Pair Identification

A confusable pair refers to two expressions that are similar in meaning or used in similar contexts, making them easy for learners to confuse.

Among the options (a), (b), (c), and (d) below, choose the confusable pair.

(a) *The girl was given a complete pedicure at the chiropodist’s.*

(b) *The girl was given a complete pedicure by the chiropodist.*

(c) *The girl had her handbag stolen at the chiropodist’s.*

(d) *The girl had her handbag stolen by the chiropodist.*

Results For comparison with human behavior, we also ran the same paradigm with 11 graduate student participants. As shown in Figure 2, humans achieved an accuracy of 89.28%, whereas the LLM’s performance was substantially lower. These human evaluation results align with those reported in (Garnham, 1981). They indicate that, in identifying written sentences, humans go beyond grammatical analysis and instead integrate linguistic input into a coherent situation model, enabling reliable judgments of event identity even when surface realizations differ.

By contrast, baseline LLMs appear to depend more on surface-level correspondences, suggesting that they solve the task through superficial cues rather than human-like semantic understanding. To test whether making situational structure explicit can mitigate this limitation, we further applied our method to closed-source GPT models: for each trial, the model was first prompted to extract objective indices (time, location, protagonists) and subjective indices (cause, intention), and then to choose the confusable pair conditioned on this extracted situation description. This explicit situation-index prompting substantially improved accuracy on GPT models (e.g., 20.8%p for GPT-3.5 and 8.3%p for GPT-4o-mini). Motivated by these findings, we propose Situation Working Memory (SituW), a mechanism that enables LLMs to maintain and update event-level context during reasoning through objective and subjective indexing.




Figure 2: **Accuracy on the confusable pair identification task.** Closed-source GPT models with and without situation-index prompting, open-source LLMs, and humans.

2.2 Situation Memory Construction

Building on this foundation, Situation Working Memory (SituW) provides a way for large language models to maintain and update a situation representation during reasoning. The situation-level representation maintained by SituW is referred to as

situation memory. We implement two pipelines for constructing situation memory: (1) an unsupervised, prompt-based pipeline and (2) a supervised pipeline based on dataset synthesis and model distillation. In the unsupervised pipeline, situation memory is generated from prompts at inference time. In the supervised pipeline, we first use GPT to synthesize a training dataset and then fine-tune a smaller model to generate situation memory.

Composition As illustrated in Figure 3, SituW consists of two indexing components: objective indexing and subjective indexing. Objective indexing encodes explicit aspects of a situation, such as time, space, and protagonists. In contrast, subjective indexing captures implicit aspects of a situation, such as causality and intention, which often require interpretation based on prior knowledge and experience.

Situation Working Memory	
Objective Indexing	
Time	(None)
Space	sugar pines
Protagonist	Zhu Hong, Red squirrels, Lina
Subjective Indexing	
Causality	The low concentration of sugar in sugar pine sap, Red squirrels searching for water or sugar
Intention	Searching for water or sugar, To discuss the dietary habits of red squirrels regarding sugar pine sap, To absorb sap from sugar pines

Figure 3: **Example of situation working memory.** Objective indices (time, space, protagonists) and subjective indices (causality, intention) encode a given context.

Unsupervised construction To construct the situation memory, we decompose the semantic representation of an input sentence S into two distinct indexing processes: objective indexing and subjective indexing. Each process is executed via a large language model (LLM), and contributes complementary components to the final memory representation. Objective indexing extracts entity-independent, factual information directly observable in the sentence. Specifically, the LLM identifies three key components: time (t), space (s), and the main protagonist (p). Formally, the objective index O is defined as:

$$O = \text{LLM}_{\text{obj}}(S) = \{t, s, p\} \quad (1)$$

Subjective indexing, in contrast, involves inferential reasoning based on the previously extracted

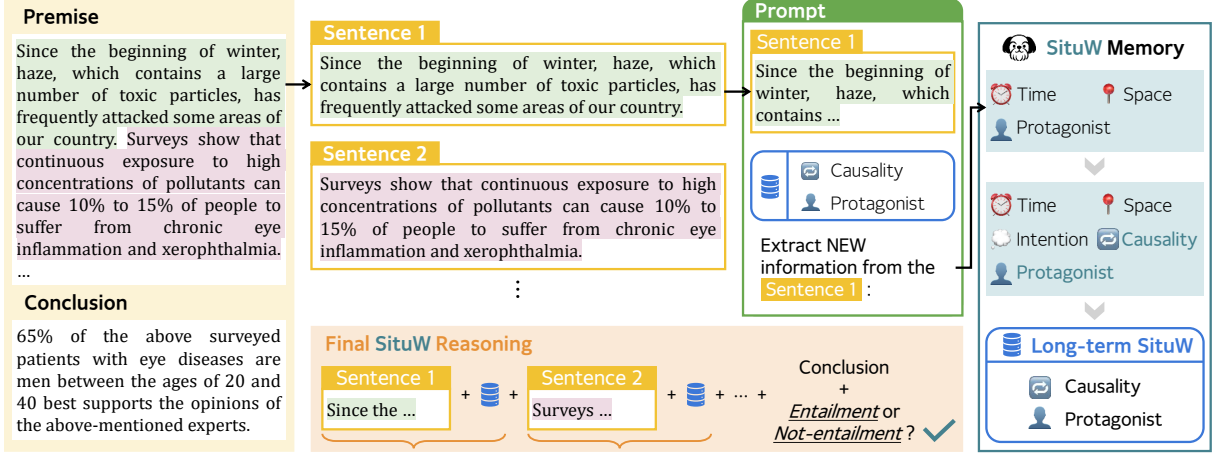


Figure 4: **Overview of situation working memory construction.** The input context is split into sentences and processed sequentially. For each sentence, an LLM extracts newly implied situational elements and incrementally updates situation working memory.

objective elements. It targets the extraction of causality (c) and intention (i), which are inherently interpretative and context-sensitive. The subjective index U is therefore conditioned on both the original sentence and its objective structure:

$$U = \text{LLM}_{\text{subj}}(S, O) = \{c, i\} \quad (2)$$

The final situation memory M is the union of these two sets, capturing both factual and interpretive dimensions of the situation:

$$M = O \cup U = \{t, s, p, c, i\} \quad (3)$$

This formulation allows for a modular and interpretable memory structure that can be utilized in downstream tasks requiring contextual understanding, temporal reasoning, or agent modeling. The prompts used to extract information are shown in Table 1.

Element	Prompt Template
Time	Extract the time: When does this occur? Reply with a short phrase or None.
Space	Extract the location: Where does this occur? Reply with a short phrase or None.
Causality	Extract the cause: What triggered this? Reply with a short phrase or None.
Intention	Extract the intention: What is the purpose? Reply with a short phrase or None.
Protagonist	Extract the protagonist: Who is involved? Reply with a short phrase or None.

Table 1: **Prompt templates** used to extract situational memory from text.

Supervised construction To construct situation memory in a supervised manner, we use a GPT-based reasoning model as a teacher model. Given

an input context X , we first split it into an ordered sequence of sentences: $X = \{S_1, \dots, S_n\}$.

For each sentence S_n , it extracts a situation-model snippet m_n consisting of five scales, using the same annotation as previously defined:

$$m_n = \left(\{t_n, s_n, p_n\}, \{c_n, i_n\} \right). \quad (4)$$

When moving from sentence S_n to S_{n+1} , we pass forward the proposition and protagonists to maintain cross-sentence coherence. We denote these carried-over variables as

$$\begin{aligned} \pi_k &= \text{Proposition}(m_n), \\ \rho_k &= \text{Protagonists}(m_n). \end{aligned} \quad (5)$$

and condition the next-step extraction on them. Formally, the teacher generates m_n conditioned on the current sentence, the preceding sentences, the previous memory state, and the carried-over variables, where $\pi_0 = \emptyset$ and $\rho_0 = \emptyset$.

$$m_n = f_{\text{teacher}}(S_k | S_{<n}, M_{n-1}, \pi_{n-1}, \rho_{n-1}), \quad (6)$$

These snippets are accumulated into a running memory state via an update rule:

$$M_n = \text{Update}(M_{n-1}, m_n), \quad M_0 = \emptyset. \quad (7)$$

After all sentences are processed, we obtain the final memory M_n .

To construct the final reasoning prompt, we store a memory trace T by concatenating each segmented sentence with its extracted situation-model snippet:

$$T = \{S_1 \oplus m_1, S_2 \oplus m_2, \dots, S_n \oplus m_n\}, \quad (8)$$

where \oplus denotes the concatenation operator.

An inference module then produces the prediction using the final memory and the memory trace with reasoning prompt (Appendix D):

$$\hat{y} = g_\phi(M_n, X, T). \quad (8)$$

To ensure high-precision pseudo-labels, we keep a generated memory trace only when the final prediction matches the ground-truth label y , yielding the pseudo-labeled dataset

$$\tilde{\mathcal{D}} = \{(X, T, y) : \hat{y} = y\}. \quad (9)$$

Finally, we fine-tune an open-source large language model on $\tilde{\mathcal{D}}$ using unsloth (Hu et al., 2022).

Situation Memory Construction Prompt

You are reading a logical context sentence-by-sentence.

Return ONLY a valid JSON object.

—
sentence:

{premises}

memory statements:

{situation memory}

memory protagonists:

{situation memory protagonists}

2.3 Situation Reasoning

To extract the final answer in an unsupervised setting, we proceed through the following pipeline. Once the situation memory $M = \{t, s, p, c, i\}$ is constructed, we transform each of its components into natural language statements that can be appended to the original context. This step enables downstream models to reason over the enriched context using explicit, interpretable cues grounded in time, space, agency, causality, and intention. Formally, we denote the original context as C , and each memory component $m_i \in M$ is processed in a predefined order—objective dimensions first (t, s, p), followed by subjective ones (c, i). Each component is verbalized using an LLM:

$$C_{j+1} = C_j + \text{LLM}_{\text{verbal}}(m_j), \quad \text{for } j = 0, 1, \dots, 4 \quad (10)$$

Here, $\text{LLM}_{\text{verbal}}(m_i)$ refers to the natural language description generated from memory element m_i . After 5 iterations, we obtain an enriched context

C_5 that incorporates all five situational dimensions. The final context C_5 can be integrated with a wide range of reasoning strategies. The model then performs iterative reasoning over C_5 using four steps: *Plan*, *Action*, *Update*, and *Process*.

In a supervised setting, we construct a synthetic training set $\tilde{\mathcal{D}}$. Given an instance in $\tilde{\mathcal{D}}$, we train the model to predict the corresponding target y .

3 Experiments

We evaluate logical reasoning in both supervised and unsupervised settings.

Datasets We conducted experiments on the following logic-oriented datasets.

PrOntoQA (Saparov and He, 2023) is a synthetic deductive FOL benchmark (modus ponens). Following (Pan et al., 2023), we use the hardest fictional-characters split (5-hop subset).

ProofWriter (Tafjord et al., 2021) tests deductive reasoning under the open-world assumption (True/False/Unknown). We evaluate the depth-5 subset with 600 balanced instances and additionally report results stratified by reasoning depth (0-5).

FOLIO (Han et al., 2024) is an open-domain dataset requiring naturalistic FOL reasoning. We use the full test set (204 examples).

LogiQA 2.0 (NLI) (Liu et al., 2023a) casts exam-style logic into NLI. We evaluate 3,840 premise-hypothesis pairs (Entailed vs. Not Entailed).

Baselines We compare against four representative reasoning baselines. *Standard LLM* directly predicts the answer from the given context and query without explicit intermediate steps. *CoT* (Wei et al., 2022) elicits step-by-step natural language reasoning before producing the final answer. *SymbCoT* (Xu et al., 2024) converts the input into symbolic representations and performs structured, logic-guided reasoning. *Logic-LM* (Pan et al., 2023) translates natural language into formal logic expressions and uses a symbolic solver to derive the prediction.

Models For closed-source backbones, we use GPT-3.5-turbo and GPT-4o-mini (Achiam et al., 2023). In the supervised setting, we generate training data with GPT-3-nano and fine-tune instruction-tuned open LLMs (Llama-3.1-Instruct 8B/70B, Llama-3.3-Instruct 70B, Mistral-Instruct 8B, Qwen2.5-Instruct 72B). In the unsupervised setting, we conduct ablations on memory construc-

Method	PrOntoQA	ProofWriter	FOLIO	AVG
<i>GPT-3.5-turbo</i>				
Standard	47.40	40.00	45.09	50.20
CoT	67.80	49.17	57.35	<u>57.78</u>
SymbCoT	<u>74.20</u>	<u>53.67</u>	39.71	55.86
Logic-LM	58.80	51.66	50.98	53.71
SituW	82.20	54.67	<u>53.43</u>	63.43
<i>GPT-4o-mini</i>				
Standard	58.20	39.33	61.76	53.16
CoT	72.60	44.16	64.22	61.46
SymbCoT	68.20	<u>62.00</u>	<u>65.69</u>	<u>65.29</u>
Logic-LM	<u>77.40</u>	37.33	62.25	58.99
SituW	79.00	64.33	72.06	71.79

Table 2: Performance comparison on logical reasoning benchmarks under unsupervised settings. Best accuracy is in **bold**, second best is underlined.

Size	0-shot		2-shot		FT		Δ
	Direct	Direct	CoT	Vanilla	SituW		
<i>Llama-3.1-instruct</i>							
8B	<u>47.78</u>	36.95	44.33	38.92	60.10	+ 21.18	
70B	61.58	67.98	<u>74.38</u>	38.42	76.35	+ 37.93	
<i>Llama-3.3-instruct</i>							
70B	64.04	65.02	<u>67.00</u>	35.54	74.88	+ 39.34	
<i>Mistral-instruct</i>							
7B	47.29	44.33	<u>51.23</u>	<u>51.23</u>	62.56	+ 11.33	
<i>Qwen2.5-instruct</i>							
72B	67.49	69.95	55.67	<u>71.14</u>	78.00	+ 6.86	

Table 3: Performance comparison on FOLIO under prompting and supervised settings. Best accuracy is in **bold**, second best is underlined. Δ is the gain over Vanilla FT.

tion and reasoning modules using open-source backbones such as Gemma3 and Llama3.

3.1 Quantitative Results

Consistent improvements under unsupervised evaluation Table 2 shows the performance of various reasoning methods on three benchmarks: PrOntoQA, ProofWriter, and FOLIO, using GPT-3.5-turbo and GPT-4o-mini. Our method, SituW, achieves the highest average accuracy with both models, 63.43% with GPT-3.5-turbo and 71.79% with GPT-4o-mini. SituW achieves the best results on PrOntoQA and ProofWriter for both backbone models, and also delivers strong gains on FOLIO, outperforming most baselines.

SituW Enables Stronger Supervision Table 3 shows the performance of supervised results on FOLIO, comparing prompting (0-/2-shot) to fine-

Method	Setting	Vanilla (%)	SituW (%)	Δ
Human	–	86.63	–	–
GPT3.5-turbo	Standard	50.49	54.81	+ 4.40
	CoT	56.20	58.30	+ 2.10
GPT4o-mini	Standard	57.41	61.97	+ 4.56
	CoT	61.45	65.78	+ 4.33

Table 4: Performance comparison on LogiQA in an unsupervised setting across different models and prompting strategy. Best accuracy is in **bold**.




Figure 5: Category-wise accuracy on LogiQA.

tuning(FT). SituW fine-tuning consistently outperforms both prompting baselines and vanilla fine-tuning across all backbones, achieving 60.10 (Llama-3.1 8B), 76.35 (Llama-3.1 70B), 74.88 (Llama-3.3 70B), 62.56 (Mistral 7B), and 78.00 (Qwen2.5 72B). Compared to Vanilla FT, SituW yields substantial absolute gains ranging from +6.86 to +39.34%, with the largest improvements on the Llama 70B models where vanilla FT underperforms even few-shot prompting. This highlights that SituW offers more effective supervision for logical reasoning than standard fine-tuning by constructing, iteratively updating, and reasoning over a sentence-level situation model.

Improves zero-shot and category-wise accuracy on LogiQA Table 4 shows the zero-shot accuracy on the LogiQA dataset. Our method, SituW, improves accuracy from 54.81% to 61.97% under the standard prompting setting and from 58.30% to 65.78% with CoT prompting. Figure 5 further breaks down the performance by reasoning category, including categorical, conjunction, disjunction, necessity, and sufficiency. In both prompting settings, SituW improves accuracy consistently across all categories. These results suggest that SituW contributes to better handling of diverse logical reasoning types in a zero-shot context.




Figure 6: Response Distribution and Accuracy on LogiQA. Left: class-wise accuracy (%). Right: the number of predicted labels (solid) and correct predictions (striped).

Objective	Causality	Intention	Accuracy (%)
✓			54.38
✓	✓		54.41
✓		✓	54.93
✓	✓	✓	64.81

Table 5: Ablation results of situation components.

3.2 Qualitative Results

More confident, less "uncertainty" responses

In Figure 6, the left panel shows accuracy for true/false inference, whereas the right panel visualizes the number of predictions for each label (true/false/uncertain; solid bars) and the corresponding number of correct predictions (striped bars). Logic-LM attains lower accuracy overall and produces a relatively higher proportion of “uncertain” predictions. By contrast, SituW makes more true/false decisions and achieves higher accuracy, with fewer “uncertain” outputs. This pattern is consistent with the idea that converting statements into logical expressions can omit or weaken some contextual cues. Moreover, the reduced rate of “uncertain” predictions with SituW suggests that incorporating situation-level understanding help the model interpret sentences more reliably than relying primarily on structural parsing alone.

Full components needed for coherent situation reasoning

Table 5 shows the contribution of each component through an ablation study. Using only the objective yields 54.38% accuracy. Adding the causality term alone provides virtually no improvement (54.41, 0.03%p), while adding the intention term alone gives a small gain (54.93, 0.55%p). In contrast, enabling both causality and intention simultaneously results in a large jump to 64.81% (10.43%p over the base objective). This non-

Cons \ Inf	Gemma3(27B)	Llama3(70B)	GPT3.5
Gemma3(4B)	62.19	65.15	54.78
Gemma3(27B)	60.65	63.73	54.78
Llama3(8B)	60.43	63.33	54.69
Llama3(70B)	61.45	64.85	54.87

Table 6: Cross-Model Construction–Inference Pairing. Accuracy (%) for open-source LLMs across construction (Cons) and inference (Inf) combinations.

Method	Accuracy (%)	Tokens		Time (s)
		in	out	
Direct	44.16	132	13.33	0.1072
CoT	<u>58.10</u>	318	176	0.4824
SymbCoT	55.86	4655	426	0.4187
SituW	63.43	4790	425	0.57

Table 7: Performance comparison of methods in accuracy and efficiency. Best accuracy is in **bold**, second best is underlined.

additive gain indicates strong complementarity: the causality constraint is most effective when coupled with intention/goal alignment, and vice versa, suggesting that all components are required for a coherent, task-relevant situation representation.

Construction/inference combinations for unsupervised reasoning

Table 6 shows additional experiments in which we combine open-source LLMs for situation-model construction and for reasoning/inference in unsupervised setting. Rows represent the construction model and columns represent the reasoning model, which facilitates analysis of how different pairings affect accuracy. Overall, Llama3-70B as the inference model delivers the strongest performance, achieving the best score when paired with Llama3-70B for construction. Gemma3-27B inference is slightly weaker but remains relatively stable across construction choices. These results show that, beyond closed-source models, combining open-source models alone can also yield strong performance in an unsupervised setting.

Accuracy-efficiency trade-off

Table 7 compares methods in terms of average accuracy, token usage, and runtime across ProntoQA, ProofWriter, and FOLIO. Direct prompting is the most efficient (132 input tokens, 13.33 output tokens, 0.107s) with the lowest computational cost, but it also achieves the lowest accuracy (44.16%). CoT substantially improves accuracy to 58.10% with a moderate increase in cost (318/176 tokens, 0.482s). SymbCoT

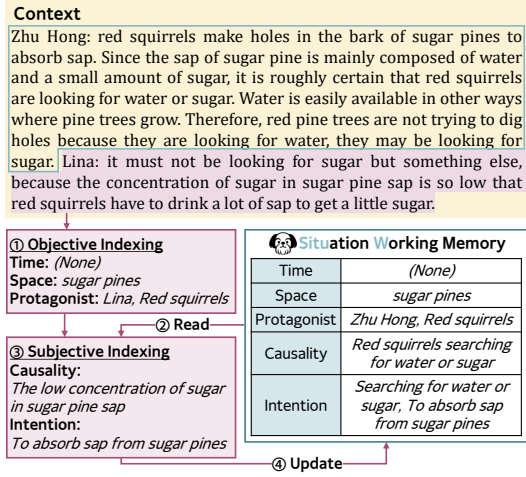


Figure 7: Overview of SituW’s workflow.

attains comparable accuracy (55.86%) but requires far more input tokens (4655) and similar runtime (0.419,s), reflecting the overhead of symbolic conversion. Our approach achieves the best accuracy overall: SituW reaches 63.43% with 4790 input tokens and 0.57,s.

Visualization of SituW Figure 7 illustrates the end-to-end workflow of unsupervised SituW. Given a context, SituW first performs objective indexing to extract explicit situational attributes such as time, space (location), and protagonists. It then performs subjective indexing to infer latent but decision-critical factors, including causal relations (e.g., why an event happens) and intentions/goals (e.g., what an agent is trying to achieve). These extracted fields are written into a structured situation working memory with fixed slots (time, space, protagonist, causality, intention), and the memory is updated iteratively as more evidence is read. The resulting working memory provides a compact, structured state representation that conditions subsequent reasoning/planning, enabling the model to focus on relevant variables and their causal-goal structure rather than re-processing the full context.

4 Related Work

Logical reasoning with LLMs Chain-of-Thought(Wei et al., 2022) improves the reasoning ability of LLMs by guiding them through step-by-step process. However, it struggles with tasks requiring strict logical adherence(Han et al., 2024; Parmar et al., 2024; Sinha et al., 2019). Recent studies address this by using LLMs as semantic parsers, converting natural language into logical expressions processed by external reasoning

engines(Wang et al., 2024b; Pan et al., 2023). Although this reduces logical errors, it introduces dependencies and potential information loss. To overcome these issues, recent methods integrate symbolic reasoning within LLMs, enabling internal logical processing even when expressions are imprecise(Sun et al., 2024; Xu et al., 2024).

LLMs with external memory LLMs achieve strong performance across many tasks (Comanici et al., 2025; Dubey et al., 2024), but still struggle with multi-step reasoning and grounding over long-context documents (Plaat et al., 2025). To mitigate these limitations, recent work augments LLMs with external memory (Hu et al., 2023; Wang et al., 2024c). For multi-hop reasoning, symbolic-memory approaches (Hu et al., 2023), (Wang et al., 2024c) either use database-backed memory, which can be domain-dependent (Hu et al., 2023), or build predicate-object working memory with rule/fact bases (Wang et al., 2024c). Although the latter is related to our approach, it mainly applies when text can be reliably converted into rules and may lose information during symbolic conversion. Scratchpad-based methods (Nye et al.) improve step-by-step computation, but are less directly applicable to general reasoning. In contrast, our method targets broader reasoning settings without requiring domain-specific schemas or strict rule-based transformations.

5 Concluding Remarks

In this paper, we proposed SituW, a method that integrates situation models from cognitive psychology into LLM-based reasoning. By dynamically constructing a situation memory, SituW enables adaptive and context-aware inference, addressing key limitations of formal logic-based methods. Experimental results show significant improvements across multiple benchmarks, enhancing LLMs’ intrinsic reasoning abilities without relying on external solvers. SituW is effective in both unsupervised and supervised settings: in the unsupervised setting, it serves as a plug-and-play inference method via prompt-based situation-memory construction, while in the supervised setting, we distill situation-memory traces as structured supervision to improve fine-tuning stability and generalization over vanilla fine-tuning. Overall, these results highlight SituW as a practical and effective method for robust contextual reasoning in LLMs.

6 Limitations

While SituW demonstrates strong performance in logical reasoning tasks, certain aspects warrant further exploration. First, the iterative process of situation memory construction and reasoning can substantially increase token usage due to repeated generation steps. This may raise overall compute and cost, especially in long-context settings where the memory grows over time. Second, the multi-step structure introduces additional latency at inference time, since situation memory construction and subsequent reasoning require multiple sequential decoding passes. Improving runtime efficiency without compromising reasoning quality remains an important direction for future work.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Hyunjoo Chae, Yeonghyeon Kim, Seungone Kim, Kai Tzu-iunn Ong, Beong-woo Kwak, Moohyeon Kim, Sunghwan Kim, Taeyoon Kwon, Jiwan Chung, Youngjae Yu, and Jinyoung Yeo. 2024. [Language Models as Compilers: Simulating Pseudocode Execution Improves Algorithmic Reasoning in Language Models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22471–22502, Miami, Florida, USA. Association for Computational Linguistics.

Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. Pal: Program-aided language models. In *International Conference on Machine Learning*, pages 10764–10799. PMLR.

Alan Garnham. 1981. Mental models as representations of text. *Memory & Cognition*, 9(6):560–565.

Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady,

David Peng, Yujie Qiao, Luke Benson, Lucy Sun, Alexander Wardle-Solano, Hannah Szabó, Ekaterina Zubova, Matthew Burtell, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, Rui Zhang, Alexander Fabbri, Wojciech Maciej Kryscinski, Semih Yavuz, Ye Liu, Xi Victoria Lin, Shafiq Joty, Yingbo Zhou, Caiming Xiong, Rex Ying, Arman Cohan, and Dragomir Radev. 2024. [FOLIO: Natural language reasoning with first-order logic](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 22017–22031, Miami, Florida, USA. Association for Computational Linguistics.

Chenxu Hu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Zhao, and Hang Zhao. 2023. Chatdb: Augmenting llms with databases as their symbolic memory. *arXiv preprint arXiv:2306.03901*.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.

Walter Kintsch. 1995. How readers construct situation models for stories: The role of syntactic cues and causal inferences. *Coherence in spontaneous text*, 1995:139–160.

Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. 2023a. [Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2947–2962.

Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. 2023b. [Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2947–2962.

Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show your work: Scratchpads for intermediate computation with language models. In *Deep Learning for Code Workshop*.

Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. 2023. [Logic-LM: Empowering large language models with symbolic solvers for faithful logical reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3806–3824, Singapore. Association for Computational Linguistics.

Mihir Parmar, Nisarg Patel, Neeraj Varshney, Mutsumi Nakamura, Man Luo, Santosh Mashetty, Arindam Mitra, and Chitta Baral. 2024. [LogicBench: Towards Systematic Evaluation of Logical Reasoning Ability of Large Language Models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages

634	13679–13707, Bangkok, Thailand. Association for Computational Linguistics.	692
635		693
636	Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, and Niki Van Stein. 2025. Multi-step reasoning with large language models, a survey. <i>ACM Computing Surveys</i> .	694
637		695
638		696
639		
640	Abulhair Saparov and He He. 2023. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In <i>The Eleventh International Conference on Learning Representations</i> .	697
641		698
642		699
643		700
644	Prithviraj Sen, Breno WSR de Carvalho, Ryan Riegel, and Alexander Gray. 2022. Neuro-symbolic inductive logic programming with logical neural networks. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , volume 36, pages 8212–8219.	701
645		702
646		703
647		
648		
649	Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. 2019. CLUTRR: A Diagnostic Benchmark for Inductive Reasoning from Text. In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 4506–4515, Hong Kong, China. Association for Computational Linguistics.	704
650		705
651		706
652		707
653		
654		
655		
656		
657		
658	Hongda Sun, Weikai Xu, Wei Liu, Jian Luan, Bin Wang, Shuo Shang, Ji-Rong Wen, and Rui Yan. 2024. De-termLR: Augmenting LLM-based Logical Reasoning from Indeterminacy to Determinacy. In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 9828–9862, Bangkok, Thailand. Association for Computational Linguistics.	708
659		709
660		710
661		711
662		
663		
664		
665		
666	Oyvind Tafjord, Bhavana Dalvi, and Peter Clark. 2021. ProofWriter: Generating implications, proofs, and abductive statements over natural language. In <i>Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021</i> , pages 3621–3634, Online. Association for Computational Linguistics.	
667		
668		
669		
670		
671		
672	Siyuan Wang, Zhongyu Wei, Yejin Choi, and Xiang Ren. 2024a. Can LLMs reason with rules? logic scaffolding for stress-testing and improving LLMs. In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 7523–7543, Bangkok, Thailand. Association for Computational Linguistics.	
673		
674		
675		
676		
677		
678		
679	Siyuan Wang, Zhongyu Wei, Yejin Choi, and Xiang Ren. 2024b. Symbolic working memory enhances language models for complex rule application. In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> , pages 17583–17604, Miami, Florida, USA. Association for Computational Linguistics.	
680		
681		
682		
683		
684		
685		
686	Siyuan Wang, Zhongyu Wei, Yejin Choi, and Xiang Ren. 2024c. Symbolic working memory enhances language models for complex rule application. In <i>Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing</i> , pages 17583–17604.	
687		
688		
689		
690		
691		
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	
	Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. 2024. Faithful Logical Reasoning via Symbolic Chain-of-Thought. In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 13326–13365, Bangkok, Thailand. Association for Computational Linguistics.	
	Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. 2024. Satlm: Satisfiability-aided language models using declarative prompting. <i>Advances in Neural Information Processing Systems</i> , 36.	
	Rolf A. Zwaan and Gabriel A. Radvansky. 1998. Situation models in language comprehension and memory. <i>Psychological Bulletin</i> , 123(2):162–185. Place: US Publisher: American Psychological Association.	

A Experimental Setup: Datasets

To evaluate the validity of the situation model in actual reasoning processes, we conducted experiments on three logical reasoning datasets:

PrOntoQA (Saparov and He, 2023) is a synthetic dataset designed to evaluate LLMs’ deductive reasoning abilities, focusing on first-order logic (FOL) with modus ponens. Following prior work (Pan et al., 2023), we use the hardest fictional characters version, specifically the 5-hop subset, which requires multi-step inference. Each instance involves verifying the truth value of a new fact in a question-answering format.

ProofWriter (Tafford et al., 2021) is a benchmark for deductive reasoning under the open-world assumption (OWA), where each (Context, Question) pair is labeled as True, False, Unknown. We use the depth-5 subset, requiring multi-step inference, and follow in sampling 600 balanced instances for evaluation. Additionally, we select instances from the OWA subset, categorized by reasoning depth (0–5 steps), to compare performance at different levels of reasoning depth.

FOLIO (Han et al., 2024) is an open-domain dataset designed to evaluate complex first-order logic (FOL) reasoning in natural language. The problems align with real-world knowledge and feature highly natural linguistic expressions. For evaluation, we use the full test set of 204 examples, which requires intricate logical inference.

LogiQA 2.0 (NLI) (Liu et al., 2023b) recasts the original LogiQA passages into a Natural-Language-Inference setting to probe complex logical reasoning. Given a premise drawn from a civil-service-exam passage and a hypothesis sentence, the task is to determine whether the hypothesis is *Entailed* or *Not Entailed*. A single premise can be paired with several alternative hypotheses, yielding a total of 3,840 premise–hypothesis pairs.

B Confusable Pair Identification Task

Examples of confusable pair identification questions from (Garnham, 1981).

Questions

In each set of sentences, a and b are the confusable pair and c and d are the nonconfusable pair.

- (a) The girl was given a complete pedicure at the chiropodist’s.

(b) The girl was given a complete pedicure by the chiropodist.

(c) The girl had her handbag stolen at the chiropodist’s.

(d) The girl had her handbag stolen by the chiropodist.
- (a) The clerk was going to have the wood cut up for him at the joiner’s.

(b) The clerk was going to have the wood cut up for him by the joiner.

(c) The clerk was informed about the rail strike at the joiner’s.

(d) The clerk was informed about the rail strike by the joiner.
- (a) The judge got his contact lenses from the optician.

(b) The judge got his contact lenses in the optician’s.

(c) The judge answered a telephone call from the optician.

(d) The judge answered a telephone call in the optician’s.
- (a) The leader purchased three loaves from the baker.

(b) The leader purchased three loaves in the baker’s.

(c) The leader got a punch on the nose from the baker.

(d) The leader got a punch on the nose in the baker’s.

5. (a) The salesman was given two pounds for the coat at the pawnbroker's.
(b) The salesman was given two pounds for the coat by the pawnbroker.
(c) The salesman was bumped into on the step at the pawnbroker's.
(d) The salesman was bumped into on the step by the pawnbroker.
6. (a) The teacher wanted her clock repaired at the watchmaker's.
(b) The teacher wanted her clock repaired by the watchmaker.
(c) The teacher was saved from falling down at the watchmaker's.
(d) The teacher was saved from falling down by the watchmaker.
7. (a) The politician bought his wife a ring from the jeweller.
(b) The politician bought his wife a ring in the jeweller's.
(c) The politician received congratulations from the jeweller.
(d) The politician received congratulations in the jeweller's.
8. (a) The bishop bought a new pipe from the tobacconist.
(b) The bishop bought a new pipe in the tobacconist's.
(c) The bishop heard an account of the state of the nation from the tobacconist.
(d) The bishop heard an account of the state of the nation in the tobacconist's.
9. (a) The servant had a tooth removed at the dentist's.
(b) The servant had a tooth removed by the dentist.
(c) The servant was injured on the knee at the dentist's.
(d) The servant was injured on the knee by the dentist.

10. (a) The student bought some spring cabbage from the greengrocer.
(b) The student bought some spring cabbage in the greengrocer's.
(c) The student heard about the concert from the greengrocer.
(d) The student heard about the concert in the greengrocer's.
11. (a) The writer collected his manuscript from the publisher.
(b) The writer collected his manuscript in the publisher's.
(c) The writer got an invitation to a wedding from the publisher.
(d) The writer got an invitation to a wedding in the publisher's.
12. (a) The matron got three pork chops from the butcher.
(b) The matron got three pork chops in the butcher's.
(c) The matron heard an account of the cricket match from the butcher.
(d) The matron heard an account of the cricket match in the butcher's.
13. (a) The worker had her prescription made up at the chemist's.
(b) The worker had her prescription made up by the chemist.
(c) The worker was made to look foolish at the chemist's.
(d) The worker was made to look foolish by the chemist.
14. (a) The chairman bought a cup for the prize from the silversmith.
(b) The chairman bought a cup for the prize in the silversmith's.
(c) The chairman heard about the epidemic from the silversmith.
(d) The chairman heard about the epidemic in the silversmith's.

15. (a) The vicar went to have a suit fitted at the tailor's.
(b) The vicar went to have a suit fitted by the tailor.
(c) The vicar was given an old lampshade at the tailor's.
(d) The vicar was given an old lampshade by the tailor.
16. (a) The actress was given a shampoo and set at the hairdresser's.
(b) The actress was given a shampoo and set by the hairdresser.
(c) The actress's brother was discovered at the hairdresser's.
(d) The actress's brother was discovered by the hairdresser.
17. (a) The woman was given a thorough examination at the doctor's.
(b) The woman was given a thorough examination by the doctor.
(c) The woman was told about the new traffic scheme at the doctor's.
(d) The woman was told about the new traffic scheme by the doctor.
18. (a) The secretary bought some curtains from the draper.
(b) The secretary bought some curtains in the draper's.
(c) The secretary got an invitation to a party from the draper.
(d) The secretary got an invitation to a party in the draper's.
19. (a) The husband bought some sugar from the grocer.
(b) The husband bought some sugar in the grocer's.
(c) The husband received a message from the grocer.
(d) The husband received a message in the grocer's.

20. (a) The policeman had a bouquet made up at the florist's.
(b) The policeman had a bouquet made up by the florist.
(c) The policeman was attacked from behind at the florist's.
(d) The policeman was attacked from behind by the florist.
21. (a) The hostess bought a mink coat from the furrier.
(b) The hostess bought a mink coat in the furrier's.
(c) The hostess received a telegram from the furrier.
(d) The hostess received a telegram in the furrier's.
22. (a) The farmer bought his copy of "Farmer's Weekly" from the newsagent.
(b) The farmer bought his copy of "Farmer's Weekly" in the newsagent's.
(c) The farmer borrowed a pound from the newsagent.
(d) The farmer borrowed a pound in the newsagent's.
23. (a) The academic took a book to be bound at the bookbinder's.
(b) The academic took a book to be bound by the bookbinder.
(c) The academic was taken to an exhibition at the bookbinder's.
(d) The academic was taken to an exhibition by the bookbinder.
24. (a) The detective was advised not to go and inspect the house at the estate-agent's.
(b) The detective was advised not to go and inspect the house by the estate-agent.
(c) The detective was tipped off at the estate-agent's.
(d) The detective was tipped off by the estate-agent.

C Human Evaluation

We conducted the human evaluation after obtaining informed consent from all participants. The consent form explained that participation was voluntary and that we would collect and analyze only anonymized evaluation data (i.e., no personally identifiable information was stored or used). Participants were recruited on a voluntary basis and were not monetarily compensated. The evaluation was then carried out using the interface shown in Figure 8.

Confusable Pair Human Evaluation

From the four options, select only two options that can be considered similar in meaning.
You may translate the options into any language you are comfortable with before answering.

Question 1 / 24
Question ID: 1

- (a) The girl was given a complete pedicure by the chiropodist.
- (b) The girl had her handbag stolen at the chiropodist's.
- (c) The girl was given a complete pedicure at the chiropodist's.
- (d) The girl had her handbag stolen by the chiropodist.

Next

Figure 8: Interface for human evaluation

D Additional Prompts

Prompts Used in the Memory Step

You are reading a logical context sentence-by-sentence. Return ONLY a valid JSON object. No extra text.

Given:

- sentence: the next sentence to read
- memory_statements: list of normalized propositions accumulated so far
- memory_protagonists: list of protagonists accumulated so far

Extract NEW information from the sentence:

- 1) time: explicit time expressions in the sentence; if none, ["none"]
- 2) space: explicit location expressions; if none, ["none"]
- 3) intention: explicit goal/purpose; if none, ["none"]
- 4) proposition: a short normalized proposition for the sentence
 - keep negation
 - keep quantifiers roughly (every/each/are/is)
- 5) protagonists_in_sentence: entities/kinds mentioned in the sentence (can include ones already in memory)

Output JSON schema:

```
{
  "time": [...],
  "space": [...],
  "intention": [...],
  "proposition": "...",
  "protagonists_in_sentence": [...],
  "confidence": 0.0-1.0
}
```

sentence:
[[SENTENCE]]

memory_statements:
[[MEM_STATEMENTS]]

memory_protagonists:
[[MEM_PROTAGONISTS]]

Prompts for Final Reasoning from Memory

You are a careful reasoner.
Use the MEMORY (accumulated statements + reading log) to determine whether the QUESTION's conclusion is:

- True: entailed by the MEMORY
- False: contradicted by the MEMORY
- Uncertain: neither entailed nor contradicted (insufficient information)

You must output exactly ONE of:
True, False, Uncertain.

Then end with exactly:
Final Answer: <True/False/Uncertain>

MEMORY_STATEMENTS:
[[STATEMENTS]]

READING_LOG:
[[READING]]

QUESTION (conclusion to evaluate):
[[QUESTION]]

OPTIONS:
[[OPTIONS]]