

When Should Models Change Their Minds?

Contextual Belief Management in Large Language Models

Haoming Xu[♣], Weihong Xu[♣], Zongrui Li[♣], Mengru Wang[♣], Yunzhi Yao[♣],
Chiyu Wu[♣], Jin Shang[♣], Yu Gong[♣], Shumin Deng^{♣, ♥*}

[♣]Zhejiang University, [♣]Homology AI

[♥]State Key Lab. of Novel Software Technology, Nanjing University, P.R. China
{haomingxu, 231sm}@zju.edu.cn

Abstract

Long-horizon interactions require language models to manage accumulating information: when to update their state, when to preserve their state, and what to ignore. We study this challenge as **Contextual Belief Management (CBM)**: maintaining a predicted belief state aligned with formal evidence while isolating task-irrelevant noise. To make CBM measurable, we introduce BeliefTrack, a closed-world benchmark spanning Rule Discovery and Circuit Diagnosis, where a finite belief space and symbolic verifiers enable exact turn-level evaluation. BeliefTrack diagnoses three failures: Failed Stay, Failed Update, and Failed Isolation. Across multiple LLMs, vanilla models exhibit severe CBM failures, while explicit belief-tracking prompts provide limited gains. In contrast, reinforcement learning with belief-state rewards reduces failure rates by 70.9% on average. Further probing reveals latent belief-state dynamics behind these failures, and representation-level steering reduces failure rates by 46.1% across two tasks¹.

1 Introduction

Large language models (LLMs) are increasingly deployed in long-horizon interactions, where their behavior depends not only on parametric knowledge but also on context, memory, tools, and run-time protocols (Yang et al., 2024; Zhou et al., 2023; Wu et al., 2024; Lee et al., 2026). In such settings, models must manage beliefs as different types of information accumulate over time. Some information should revise the model’s current belief state, some should leave it unchanged, and some should be ignored altogether. Recent work on context learning, such as CL-Bench (Dou et al., 2026), studies whether models can absorb rules, knowledge, or procedures from context and translate them into




Figure 1: Overview of *Contextual Belief Management* (CBM). CBM requires models to maintain a predicted belief state over a belief space, update it only when warranted by formal evidence, and filter task-irrelevant context or noise. The pilot Rule Discovery study reveals substantial belief-management errors in frontier models.

effective behavior. However, absorbing contextual information is not enough: a model must also decide which information counts as formal evidence, when that evidence warrants belief revision, and when task-irrelevant context should be filtered out.

As Figure 1 shows, we study this problem as **Contextual Belief Management (CBM)**: a model’s ability to maintain an evidence-aligned belief state throughout a multi-turn interaction. Rather than simulating open-ended dialogue, we operationalize CBM in a controlled closed-world setting. Specifically, we introduce **BeliefTrack**, a closed-world benchmark with two environments: Rule Discovery (RD) and Circuit Diagnosis (CD). Both environments define finite belief spaces and use symbolic verifiers, allowing exact turn-level comparison between predicted and oracle belief states. This design abstracts away open-ended ambiguity and allows us to evaluate distinct belief-management operations precisely.

As a pilot study, we evaluate Qwen3.5-Plus (Team, 2026), DeepSeek-V3.2 (DeepSeek-AI et al., 2025), and GPT-5.2 (Singh et al., 2026) on 135 Rule Discovery examples with task-irrelevant noise. As shown in Figure 1, all three frontier models exhibit substantial belief-management errors. These results suggest that CBM failures arise even

*Corresponding Author.

¹Code is available at <https://github.com/zjunlp/CBM>.

when the relevant evidence is explicitly specified.

Understanding these failures requires more than checking whether the model produces the correct belief state at a single turn. A model must preserve a stable belief when formal evidence remains unchanged, revise its belief when formal evidence changes, and isolate its belief state from task-irrelevant context. To localize these errors, BeliefTrack evaluates three diagnostic failures: *Failed Stay*, *Failed Update*, and *Failed Isolation*. These diagnostics distinguish belief calibration failures from belief isolation failures.

With these concepts in place, §3 defines Contextual Belief Management (CBM) and details how BeliefTrack operationalizes it in Rule Discovery and Circuit Diagnosis with exact symbolic verification. In §5, we evaluate current LLMs and find that **vanilla models exhibit severe CBM failures**, while explicit belief-tracking prompts provide only limited and inconsistent gains. We further show that **reinforcement learning with belief-state rewards** substantially reduces failure rates, transfers across environments, and improves robustness to task-irrelevant noise. In §6, prompt-based probing reveals **latent belief-state dynamics** behind these failures, including belief-state drift, back-tracking failure, and contextual hijacking. Finally, **representation-level steering** directly improves predicted-oracle belief-state alignment, suggesting that CBM failures are not only measurable but also **actionable at the representation level**.

2 Related Work

Knowledge Conflict. Deciding which information to trust is central to belief management in language models. Prior work shows that models struggle to resolve conflicts between parametric memory and context from passages, user claims, demonstrations (Longpre et al., 2021; Wang et al., 2024; Xu et al., 2024c; Kortukov et al., 2024; Jin et al., 2024; Xie et al., 2024; Xu et al., 2024d; Hagström et al., 2026). Recent work further highlights belief dependencies in conflict resolution, where updating one fact can affect others (Zhang et al., 2026a; Xu et al., 2026). By contrast, CBM does not introduce direct information conflicts, but tests whether models update beliefs only from formal evidence.

Multi-turn Reasoning Instability. LLMs often become unreliable in long interactions: they lose relevant evidence (Liu et al., 2024; Zhang et al., 2025; Al-Tawaha et al., 2026), degrade in multi-




Figure 2: Comparison between *Contextual Belief Management* and *Theory of Mind*.

turn instruction following (Laban et al., 2026; Duan et al., 2025), and fail under contextual pressure (Xu et al., 2024b; Deng et al., 2026). Recent work further identifies *contextual inertia*, where models fail to revise earlier generations or intermediate inferences despite later contradictory evidence (Huang et al., 2026; Chen et al., 2026a; Liu et al., 2025), as well as mechanisms involving metacognition, memory management, and epistemic state tracking (Raj, 2026; Yona et al., 2026; Chen et al., 2026b; Yalon et al., 2026). CBM turns these instabilities into exact turn-level diagnostics, separating failures of belief calibration and isolation.

Belief Tracking and Theory of Mind. Recent work studies belief dynamics in LLMs, from revising prior reasoning (Wilie et al., 2024) to maintaining temporal belief consistency in long-running agents (Myakala et al., 2026) and constructing spatial beliefs through active exploration (Zhang et al., 2026b). These studies highlight the challenge of maintaining stable yet revisable beliefs over time. In Theory of Mind (ToM), belief tracking instead targets hidden mental states of other agents, such as beliefs, desires, intentions, and perspectives (Ullman, 2023; Kim et al., 2023; Chen et al., 2024; Strachan et al., 2024a,b; Kosinski, 2024; Shapira et al., 2024; Xu et al., 2024a; Cross et al., 2025; Prakash et al., 2026; Shi et al., 2025). As illustrated in Figure 2, ToM is a third-person inference problem, whereas CBM asks what the model itself should believe from accumulated formal evidence. We evaluate this first-person problem in closed-world environments with finite belief spaces and symbolic verifiers.

3 Preliminary

3.1 Problem Formulation

We formalize **Contextual Belief Management** (CBM) as a model’s ability to maintain an evidence-aligned belief state throughout a multi-turn interaction. At each turn $t \in \{1, \dots, T\}$ in an environ-

ment $\mathcal{E}$, a belief-tracking model f_θ receives an observation $o_t = (e_t, n_t)$, where e_t denotes formal evidence and n_t denotes optional task-irrelevant noise ($n_t = \emptyset$ in clean settings). Let $e_{1:t} = (e_1, \dots, e_t)$ and $o_{1:t} = (o_1, \dots, o_t)$ denote the formal-evidence history and the observation history, respectively.

Let $\mathcal{B}_\mathcal{E}$ denote the task-specific belief space of environment $\mathcal{E}$. Its elements are candidate hypotheses representing all possible task outcomes. A belief state is a subset of this space: the candidate hypotheses that remain supported by the formal evidence observed so far. We define two belief states at each turn. The **Oracle Belief State** $S_t^* \subseteq \mathcal{B}_\mathcal{E}$ is the logically correct subset determined by the formal-evidence history $e_{1:t}$. The **Predicted Belief State** $\hat{S}_t = f_\theta(o_{1:t}) \subseteq \mathcal{B}_\mathcal{E}$ is the subset produced by the model from the observation history $o_{1:t}$.

Contextual Belief Management

The objective of CBM is to maintain turn-level alignment between the model’s predicted belief state and the oracle belief state over the full trajectory τ :

$$\max_{\theta} \mathbb{E}_{\tau \sim \mathcal{E}} \left[\frac{1}{T} \sum_{t=1}^T \mathbb{I}\{\hat{S}_t = S_t^*\} \right], \quad (1)$$

where $\mathbb{I}$ is the indicator function.

We distinguish two classes of CBM failures.

Type I: Belief Calibration Failures. Belief calibration failures arise when the model fails to track the evidence-aligned belief state under clean conditions ($n_t = \emptyset$). They manifest in two forms:

1. Failed Stay: The oracle belief state remains unchanged ($S_t^* = S_{t-1}^*$), but the model fails to preserve this stable state, producing $\hat{S}_t \neq S_t^*$.

2. Failed Update: The oracle belief state changes ($S_t^* \neq S_{t-1}^*$), but the model fails to transition to the revised state, producing $\hat{S}_t \neq S_t^*$.

Type II: Belief Isolation Failures. Belief isolation failures occur when the model fails to separate formal evidence from task-irrelevant noise. Specifically, consider a turn where the model correctly predicts the oracle belief state under the clean formal-evidence history. If adding task-irrelevant noise n_t changes the prediction and leads to $\hat{S}_t \neq S_t^*$, the model has incorrectly treated non-evidential noise as part of the formal reasoning signal.

3.2 Closed-World Task Environments

Evaluating CBM requires isolating formal-evidence tracking from reliance on pre-trained world knowledge. In standard multi-turn question-answering datasets, errors may stem from factual hallucination rather than belief-management failure, and the correct belief state is often not exactly computable at each turn.

We therefore introduce **BeliefTrack**, a closed-world benchmark for evaluating turn-level belief dynamics. As shown in Figure 3, BeliefTrack formulates each task as evidence-conditioned belief-state tracking over a finite belief space $\mathcal{B}_\mathcal{E}$. All task-relevant evidence is specified within the episode, and the model must output the predicted belief state $\hat{S}_t \subseteq \mathcal{B}_\mathcal{E}$: the subset of candidate hypotheses aligned with the accumulated formal evidence.

Figure 3 illustrates the two environments used in BeliefTrack. Both instantiate the same CBM formulation but differ in the semantics of their candidate hypotheses and formal evidence.

Task A: Rule Discovery (RD). Adapted from Watson’s 2-4-6 paradigm, Rule Discovery defines a finite belief space $\mathcal{B}_{\text{RD}} = \{h_1, \dots, h_M\}$, where each candidate hypothesis is a possible rule, such as `ascending_order` or `sum_greater_than_10`. At each turn t , the formal evidence e_t consists of a proposed triple, such as `[3, 8, 1]`, and its ground-truth label, YES or NO, determined by a hidden oracle rule. The oracle belief state is the subset of candidate rules that remain consistent with the accumulated triple-label evidence.

Task B: Circuit Diagnosis (CD). Circuit Diagnosis evaluates diagnostic reasoning in a circuit-fault setting. It defines a finite belief space $\mathcal{B}_{\text{CD}} = \{h_1, \dots, h_M\}$, where each candidate hypothesis is a possible circuit fault, such as `Battery_no_output` or `R1_open`. At each turn t , the formal evidence e_t is an instrument reading, such as `Current(Main)>0` or `Voltage(R1)=0`. The oracle belief state is the subset of candidate faults whose predicted circuit behavior remains consistent with the accumulated readings.

3.3 Dataset Generation and Verification

As illustrated in Figure 3, we generate three diagnostic datasets from BeliefTrack: $\mathcal{D}^{\text{stay}}$, $\mathcal{D}^{\text{update}}$, and $\mathcal{D}^{\text{iso}}$, targeting Failed Stay, Failed Update, and Failed Isolation, respectively. Each dataset consists of fixed user-side multi-turn diagnostic templates paired with a symbolic verifier that computes the or-




Figure 3: BeliefTrack framework. Given a finite belief space, the model must output a predicted belief state, i.e., a subset of candidate hypotheses that it considers supported by the accumulated formal evidence. The goal is to align this predicted belief state with the oracle belief state at each turn. BeliefTrack includes two closed-world environments, Rule Discovery (RD) and Circuit Diagnosis (CD), and evaluates three typical diagnostic failure modes: *Failed Stay*, *Failed Update*, and *Failed Isolation*.

acle belief state S_t^* at each turn. Because each environment has a finite belief space and fully specified verification logic, automatic evaluation requires no human annotation. Assistant-side trajectories are sampled during evaluation.

1. Dataset for Failed Stay. $\mathcal{D}^{\text{stay}}$ tests whether models preserve an oracle belief state. Each template contains a lock point t_{lock} , before which evidence narrows the oracle state to a target subset $S_{\text{lock}} \subseteq B_E$. Afterward, only redundant evidence is added, so the oracle state remains fixed:

$$S_t^* = S_{t_{\text{lock}}}^* = S_{\text{lock}}, \quad \forall t > t_{\text{lock}}. \quad (2)$$

Errors at these turns correspond to **Failed Stay**.

2. Dataset for Failed Update. $\mathcal{D}^{\text{update}}$ tests whether models revise their belief state after earlier evidence is corrected. Each template includes a to-be-corrected evidence item e_j , followed by a CORRECTION at turn t_c that replaces it with corrected evidence e'_j . The verifier recomputes the oracle state under the corrected evidence history, and we retain correction turns where

$$S_{t_c}^* \neq S_{t_c-1}^*. \quad (3)$$

Errors at these turns correspond to **Failed Update**.

3. Dataset for Failed Isolation. $\mathcal{D}_{\text{iso}}$ tests whether models ignore task-irrelevant noise. Each template

forms a clean-noised pair, $o_{1:t}^{\text{clean}} = (e_{1:t}, \emptyset_{1:t})$ and $o_{1:t}^{\text{noise}} = (e_{1:t}, n_{1:t})$, that shares the same evidence history $e_{1:t}$ but differs in the noise history. Thus, the oracle belief state is unchanged:

$$S_t^{*,\text{clean}} = S_t^{*,\text{noise}}, \quad \forall t. \quad (4)$$

Failed Isolation occurs when the model succeeds on the clean trajectory but fails on the noisy one.

The concrete dataset templates are provided in Appendix D.1. Since the environments are symbolic and automatically verifiable, they support scalable trajectory generation. In this work, we instantiate 1,300/1,049 Rule Discovery/Circuit Diagnosis trajectories for instruct models, and 1,503/1,616 trajectories for thinking models.

4 Methods for Improving CBM

We evaluate two methods for improving CBM: *Belief-Tracking Prompt (BT-Prompt)*, a training-free prompt-based enhancement method, and *RL with belief-state rewards*, a verifier-guided reinforcement-learning method.

4.1 BT-Prompt

BT-Prompt is a parameter-free test-time baseline that encodes the CBM procedure in the system prompt. It instructs the model to maintain the

current set of valid formal evidence, ignore non-evidential noise, re-evaluate all candidate hypotheses against the accumulated evidence, and revise the evidence set when explicit corrections invalidate earlier observations. This allows previously eliminated hypotheses to be restored when the evidence excluding them is removed. BT-Prompt is applied uniformly across both environments and all diagnostic trajectory types; full templates are provided in Appendix D.2. We additionally consider a few-shot variant, **FS BT-Prompt**. After the belief-tracking principles in the original BT-Prompt and before the strict output-format specification, we insert three demonstration examples illustrating the required belief-tracking behavior. All other settings remain identical to BT-Prompt.

4.2 Supervised Fine-Tuning

We include supervised fine-tuning (SFT) as a training-based baseline. Training examples are extracted from multi-turn BeliefTrack trajectories, but each prompt targets a single evaluated turn. For each target turn t , the prompt context $q_t = \text{Prompt}(\mathcal{B}_{\mathcal{E}}, o_{1:t})$ is paired with the oracle belief state S_t^* as the supervised target. Totally, we train separate SFT-RD and SFT-CD models on $\mathcal{D}_{\text{stay}}$ and $\mathcal{D}_{\text{update}}$ only, excluding $\mathcal{D}_{\text{iso}}$.

4.3 RL with Belief-State Rewards

We further optimize the model with GRPO (Shao et al., 2024) using rewards computed by a symbolic verifier. As in SFT, training examples are extracted from multi-turn BeliefTrack trajectories, and each GRPO prompt targets a single evaluated turn. Specifically, for a target turn t , we construct a prompt context $q_t = \text{Prompt}(\mathcal{B}_{\mathcal{E}}, o_{1:t})$ from the belief space and the full observation history up to that turn. The model then produces one response y_i containing a predicted belief state $\hat{S}_{i,t}$ for the target turn. The verifier compares $\hat{S}_{i,t}$ with the oracle belief state S_t^* and assigns a reward only for that turn. For each prompt context q_t , GRPO samples a group of outputs $\{y_1, \dots, y_G\}$ and optimizes the standard clipped objective:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q_t, \{y_i\}_{i=1}^G} \left[\frac{1}{G} \sum_{i=1}^G \left\{ \min \left(\rho_i(\theta) A_i, \right. \right. \right. \\ \left. \left. \left. \text{clip} \left(\rho_i(\theta), 1 - \epsilon, 1 + \epsilon \right) A_i \right) - \beta \text{D}_{\text{KL}} \left(\pi_{\theta} \parallel \pi_{\text{ref}} \right) \right\} \right], \quad (5)$$

where $A_i = (R_i - \text{mean}(\{R_i\}_{i=1}^G)) / \text{std}(\{R_i\}_{i=1}^G)$ is the group-normalized reward advantage and

$$\rho_i(\theta) = \pi_{\theta}(y_i | q_t) / \pi_{\theta_{\text{old}}}(y_i | q_t).$$

We use a dense Jaccard belief-state reward:

$$R_i(q_t) = \frac{|\hat{S}_{i,t} \cap S_t^*|}{|\hat{S}_{i,t} \cup S_t^*|}. \quad (6)$$

Since S_t^* is non-empty by construction, the denominator is always non-zero. This reward measures set-level alignment between the predicted and oracle belief states at the target turn. Unlike sparse exact match, it gives partial credit to predictions that overlap with the oracle state; we ablate this design in Appendix C.1.

5 Experiments

5.1 Metrics

Based on the diagnostic datasets defined in Section 3.3, we use a strict k -repeat evaluation protocol. For each diagnostic sample x , the user-side multi-turn template is fixed, and we independently sample k assistant-side trajectories. Let $E_m^{(i)}(x) \in \{0, 1\}$ indicate whether the i -th trajectory exhibits the target failure mode $m \in \{\text{stay}, \text{update}, \text{iso}\}$. Here, stay, update, and iso correspond to Failed Stay, Failed Update, and Failed Isolation, respectively. We define sample-level failure as

$$F_m^{(k)}(x) = \mathbb{I}[\exists i \in \{1, \dots, k\}, E_m^{(i)}(x) = 1]. \quad (7)$$

Thus, a sample fails if any repeated trajectory exhibits the target failure.

Failed Stay Rate (FSR). For $x \in \mathcal{D}_{\text{stay}}$, $E_{\text{stay}}^{(i)}(x) = 1$ if the i -th trajectory makes a Failed Stay error on the evaluated post-lock turns.

$$\text{FSR} = \frac{1}{|\mathcal{D}_{\text{stay}}|} \sum_{x \in \mathcal{D}_{\text{stay}}} F_{\text{stay}}^{(k)}(x). \quad (8)$$

Failed Update Rate (FUR). For $x \in \mathcal{D}_{\text{update}}$, $E_{\text{update}}^{(i)}(x) = 1$ if the i -th trajectory makes a Failed Update error at the correction turn.

$$\text{FUR} = \frac{1}{|\mathcal{D}_{\text{update}}|} \sum_{x \in \mathcal{D}_{\text{update}}} F_{\text{update}}^{(k)}(x). \quad (9)$$

Failed Isolation Rate (FIR). For $x \in \mathcal{D}_{\text{iso}}$, $E_{\text{iso}}^{(i)}(x) = 1$ if the i -th trajectory succeeds on the clean trajectory but fails on the noised trajectory.

$$\text{FIR} = \frac{1}{|\mathcal{D}_{\text{iso}}|} \sum_{x \in \mathcal{D}_{\text{iso}}} F_{\text{iso}}^{(k)}(x). \quad (10)$$

All experiments use $k = 3$ and lower values indicate better performance.

Method	Rule Discovery (RD)			Circuit Diagnosis (CD)			General		
	FSR ↓	FUR ↓	FIR ↓	FSR ↓	FUR ↓	FIR ↓	GSM8K ↑	MMLU ↑	MT-Bench ↑
Qwen2.5-7B-Instruct									
Vanilla	99.0	98.0	97.0	99.0	98.0	97.0	83.3 ± 0.4	73.5 ± 0.3	7.514
BT-Prompt	79.0 _{↓20.2%}	80.0 _{↓18.4%}	95.0 _{↓2.1%}	48.0 _{↓51.5%}	87.9 _{↓10.3%}	92.0 _{↓5.2%}	—	—	—
FS BT-Prompt	84.0 _{↓15.2%}	77.0 _{↓21.4%}	93.0 _{↓4.1%}	63.0 _{↓36.4%}	92.9 _{↓5.2%}	92.0 _{↓5.2%}	—	—	—
SFT-RD	0.0 _{↓100.0%}	1.0 _{↓99.0%}	<u>29.0</u> _{↓70.1%}	<u>1.0</u> _{↓99.0%}	<u>32.3</u> _{↓67.0%}	<u>70.0</u> _{↓27.8%}	69.7 ± 1.3	73.3 ± 0.4	7.516
SFT-CD	68.0 _{↓31.3%}	90.0 _{↓8.2%}	89.0 _{↓8.2%}	3.0 _{↓97.0%}	<u>18.2</u> _{↓81.4%}	65.0 _{↓33.0%}	70.9 ± 1.3	73.4 ± 0.4	7.494
RL-RD	0.0 _{↓100.0%}	<u>2.0</u> _{↓98.0%}	20.0 _{↓79.4%}	6.0 _{↓93.9%}	<u>28.3</u> _{↓71.1%}	<u>35.0</u> _{↓63.9%}	84.8 ± 0.5	73.1 ± 0.2	7.522
RL-CD	<u>42.0</u> _{↓57.6%}	74.0 _{↓24.5%}	57.0 _{↓41.2%}	0.0 _{↓100.0%}	0.0 _{↓100.0%}	20.0 _{↓79.4%}	83.9 ± 0.3	73.4 ± 0.4	7.478
Qwen3.5-9B									
Vanilla	47.0	60.0	83.7	43.2	62.7	95.4	94.0 ± 0.2	78.1 ± 0.3	7.879
BT-Prompt	49.0 _{↑4.3%}	69.0 _{↑15.0%}	81.4 _{↓2.7%}	47.4 _{↑9.7%}	63.5 _{↑1.3%}	81.4 _{↓14.7%}	—	—	—
FS BT-Prompt	67.0 _{↑42.6%}	60.0 _{↓0.0%}	65.1 _{↓22.2%}	51.1 _{↑18.2%}	69.8 _{↑11.4%}	67.4 _{↓29.3%}	—	—	—
SFT-RD	<u>56.0</u> _{↑19.1%}	<u>59.0</u> _{↓1.7%}	79.1 _{↓5.5%}	54.2 _{↑25.5%}	54.8 _{↓12.7%}	81.4 _{↓14.7%}	93.9 ± 0.7	78.4 ± 0.3	7.858
SFT-CD	45.0 _{↓4.3%}	35.0 _{↓41.7%}	62.8 _{↓25.0%}	44.7 _{↑3.6%}	66.7 _{↑6.3%}	88.4 _{↓7.4%}	94.0 ± 0.7	78.5 ± 0.3	7.897
RL-RD	6.0 _{↓87.2%}	8.0 _{↓86.7%}	18.6 _{↓77.8%}	<u>20.0</u> _{↓53.7%}	<u>21.4</u> _{↓65.9%}	<u>34.9</u> _{↓63.4%}	94.5 ± 0.6	78.2 ± 0.3	7.903
RL-CD	<u>31.0</u> _{↓34.0%}	<u>34.0</u> _{↓43.3%}	<u>41.9</u> _{↓49.9%}	12.1 _{↓72.0%}	15.9 _{↓74.6%}	16.3 _{↓82.9%}	94.5 ± 0.6	78.7 ± 0.3	7.824
Gemma-2-9B-Instruct									
Vanilla	76.0	52.0	82.0	56.0	90.9	55.0	78.2 ± 1.1	72.1 ± 0.4	7.680
BT-Prompt	87.0 _{↑14.5%}	55.0 _{↑5.8%}	87.0 _{↑6.1%}	32.0 _{↓42.9%}	89.9 _{↓1.1%}	44.0 _{↓20.0%}	—	—	—
FS BT-Prompt	79.0 _{↑3.9%}	52.0 _{↓0.0%}	83.0 _{↑1.2%}	<u>16.0</u> _{↓71.4%}	96.0 _{↑5.6%}	43.0 _{↓21.8%}	—	—	—
SFT-RD	0.0 _{↓100.0%}	0.0 _{↓100.0%}	3.0 _{↓96.3%}	0.0 _{↓100.0%}	58.6 _{↓35.5%}	32.0 _{↓41.8%}	79.4 ± 1.1	72.2 ± 0.4	7.673
SFT-CD	76.0 _{↓0.0%}	77.0 _{↑48.1%}	57.0 _{↓30.5%}	0.0 _{↓100.0%}	<u>17.2</u> _{↓81.1%}	23.0 _{↓58.2%}	78.7 ± 1.1	72.2 ± 0.4	7.647
RL-RD	0.0 _{↓100.0%}	<u>4.0</u> _{↓92.3%}	<u>4.0</u> _{↓95.1%}	0.0 _{↓100.0%}	72.7 _{↓20.0%}	45.0 _{↓18.2%}	76.4 ± 1.2	71.6 ± 0.4	7.620
RL-CD	<u>57.0</u> _{↓25.0%}	59.0 _{↑13.5%}	5.0 _{↓93.9%}	43.0 _{↓23.2%}	0.0 _{↓100.0%}	<u>26.0</u> _{↓52.7%}	76.2 ± 1.2	72.6 ± 0.4	7.718

Table 1: Main diagnostic and cross-task generalization results. We report Failed Stay Rate (FSR), Failed Update Rate (FUR), and Failed Isolation Rate (FIR), where lower is better (↓), together with GSM8K, MMLU, and MT-Bench-101, where higher is better (↑). **Bold** and underline indicate the best and second-best diagnostic results within each model block. FS denotes few-shot prompting; SFT-RD/CD and RL-RD/CD denote models trained solely on Rule Discovery (RD) or Circuit Diagnosis (CD). Blue and orange indicate in-domain (ID) and out-of-domain (OOD) evaluation, respectively, and subscripts show relative changes from the Vanilla baseline.

5.2 Implementation Details

We split each environment at the oracle level before trajectory generation. In RD, train and test trajectories are generated from disjoint rule sets; in CD, they are also generated from disjoint fault sets. This prevents oracle-specific memorization and evaluates generalization to unseen evidence-conditioned belief states.

For each base model, we train two single-environment RL variants: RL-RD, trained only on the RD training split, and RL-CD, trained only on the CD training split. RL training uses target-turn prompts from $\mathcal{D}_{\text{stay}}$ and $\mathcal{D}_{\text{update}}$: post-lock turns for belief preservation and correction turns for belief revision. We exclude $\mathcal{D}_{\text{iso}}$ from training, so RL never observes task-irrelevant noise trajectories.

Both RL variants are evaluated on both held-out RD and held-out CD test sets. Evaluation on the training environment measures in-domain performance, while evaluation on the other environment measures cross-environment generalization. Vanilla and BT-Prompt require no training and are

evaluated on the same held-out test sets across all three diagnostic trajectory types. Detailed split statistics are provided in Appendix B.1.

We evaluate Qwen2.5-7B-Instruct (Qwen et al., 2025) and Qwen3.5-9B (Team, 2026) and Gemma-2-9B-Instruct (Team et al., 2024). Training and evaluation hyperparameters, decoding settings, and infrastructure are provided in Appendix B.2.

5.3 Main Results

Table 1 reports diagnostic failure rates and general capability scores. In addition to the strict k -repeat evaluation protocol, we also report per-trajectory failure rates with 95% confidence intervals in Table 3. We highlight three findings.

Vanilla models consistently lack reliable CBM. Qwen2.5-7B-Instruct fails almost completely across both environments, with failure rates around 97–99% on all three metrics. Qwen3.5-9B is stronger but still exhibits substantial CBM failures, especially on the Failed Isolation split $\mathcal{D}_{\text{iso}}$: its FIR reaches 95.4% in Circuit Diagno-

sis. Gemma-2-9B-Instruct shows a similar pattern: despite lower FSR on Circuit Diagnosis, its failure rates remain high overall, reaching 90.9% FUR on CD and 82.0% FIR on RD. These results indicate that strong instruction-tuned models do not reliably manage beliefs under our diagnostic interventions.

BT-Prompt provides limited gains. Although BT-Prompt and its few-shot variant improve some metrics, their effects are inconsistent across models and environments. For Qwen2.5-7B, both variants yield partial improvements but leave several failure rates high; for Qwen3.5-9B, they can even degrade performance, increasing RD FUR by 15.0% with BT-Prompt and RD FSR by 42.6% with Few-shot BT-Prompt. For Gemma-2-9B, prompting substantially improves some CD metrics but provides little or negative gains on RD. Overall, explicit test-time instructions, even with few-shot demonstrations, are insufficient for reliable CBM.

SFT improves CBM but generalizes inconsistently and can impair general ability. SFT can yield strong in-domain gains, but its cross-environment transfer is substantially less consistent than RL, particularly on the Qwen models. For Gemma-2-9B, SFT-RD transfers well to CD, reducing FSR/FUR/FIR from 56.0%/90.9%/55.0% to 0.0%/58.6%/32.0%, whereas SFT-CD transfers weakly to RD, leaving FSR unchanged at 76.0% and increasing FUR from 52.0% to 77.0%. Moreover, SFT sometimes trades belief-management gains for general capability. For Qwen2.5-7B, GSM8K drops from 83.3 to 69.7/70.9 after SFT-RD/CD, while the corresponding RL variants retain scores of 84.8/83.9. These results suggest that supervised training can fit environment-specific belief-tracking behavior, but does not reliably induce transferable CBM.

RL with belief-state rewards consistently improves CBM. RL yields strong in-domain gains across all three models and both environments. For Qwen2.5-7B, RD training reduces FSR/FUR to 0.0%/2.0%, while CD training reduces them to 0.0%/0.0%. For Qwen3.5-9B, RL lowers FSR/FUR to 6.0%/8.0% on RD and 12.1%/15.9% on CD. Gemma-2-9B shows similarly strong gains, reaching 0.0%/4.0% on RD and reducing CD FUR to 0.0%. RL also generalizes beyond the training environment. **First, it transfers across tasks:** RD-trained Qwen2.5-7B reduces unseen CD FSR/FUR by 93.9%/71.1%. For Qwen3.5-9B, RD-trained RL reduces CD FSR/FUR by 53.7%/65.9%, while CD-trained RL reduces RD FSR/FUR by 34.0%/43.3%.

Gemma-2-9B also shows substantial, though more metric-dependent, transfer, including a 100.0% reduction in CD FSR after RD training and a 93.9% reduction in RD FIR after CD training. **Second, RL improves belief isolation even though $\mathcal{D}_{\text{iso}}$ is excluded from training.** RD-trained Qwen2.5-7B reduces FIR by 79.4% in-domain and 63.9% out-of-domain. RD-trained Qwen3.5-9B similarly reduces FIR by 77.8%/63.4%, while CD-trained Qwen3.5-9B reduces CD FIR by 82.9%. For Gemma-2-9B, RL reduces FIR by up to 95.1% in-domain and 93.9% across tasks.

Together, these results suggest that verifier-guided RL improves belief-state management rather than merely fitting task-specific patterns. These gains largely preserve general ability: GSM8K (Cobbe et al., 2021), MMLU (Hendrycks et al., 2021), and MT-Bench-101 (Bai et al., 2024) remain stable overall. Finally, Figure 6 shows that most CBM gains emerge early in training, while later checkpoints fluctuate across metrics and transfer settings.

6 Robustness and Mechanistic Analysis

We analyze CBM along three axes: temporal robustness, contextual robustness, and mechanism, using Qwen3.5-9B throughout this section. We first vary anchoring depth to test whether models can preserve and revise belief states over long evidence histories. We then vary task-irrelevant noise to test whether models can isolate formal evidence from misleading cues. Finally, we use prompt-based probing and representation-level steering to examine whether CBM failures correspond to modifiable representation-level patterns.

6.1 Anchoring Depth

We study how CBM changes as the evidence that should anchor the current belief state becomes temporally distant. *Redundant Depth* d_{red} increases the number of redundant but consistent evidence turns, testing whether models can preserve an unchanged oracle belief state. *Correction Delay* d_{cor} increases the gap between flawed evidence and its later correction, testing whether models can revise beliefs after delayed backtracking.

Figure 4 (Left and Middle) shows that Vanilla models become less reliable as the anchoring evidence moves farther back in the context. Larger redundant depth d_{red} leads to higher FSR, reflecting drift from stable oracle belief states. Larger cor-




Figure 4: Effects of temporal stress and task-irrelevant context on CBM. *Left*: FSR under increasing redundant depth; *Middle*: FUR under increasing correction delay; *Right*: FIR under different noise types. We compare Vanilla, BT-Prompt, RL-RD, and RL-CD on Rule Discovery and Circuit Diagnosis. Lower values indicate better CBM.

rection delay d_{cor} leads to higher FUR, reflecting difficulty revising beliefs after delayed corrections. BT-Prompt does not consistently mitigate these effects: its trends largely mirror Vanilla, and in some settings it even increases the failure rate.

RL substantially improves resistance to anchoring depth. In-domain RL keeps FSR/FUR low and relatively stable as d_{red} or d_{cor} increases, while out-of-domain RL degrades more gradually but remains stronger than Vanilla and BT-Prompt. This suggests that **RL with Belief-State Rewards improves temporal robustness**, especially when the training and test environments match.

6.2 Typology of Contextual Interference

Beyond temporal stress, CBM also requires isolating formal evidence from task-irrelevant noise. We evaluate belief isolation on $\mathcal{D}_{iso}$ under a clean condition, *None*, and three noised conditions: *Sycophancy*, *Authority*, and *Stress*. In the noised conditions, the formal evidence remains unchanged, while task-irrelevant context is added to reinforce the model’s prediction, override the evidence, or induce pressure through time constraints or emotional cues. Since the oracle belief state is unchanged, clean-to-noised degradation indicates Failed Isolation. Noise prompt templates are provided in Appendix D.3.

Figure 4 (Right) shows that Vanilla and BT-Prompt remain highly vulnerable to task-irrelevant context. Among the three noise types, *Authority* causes the largest increase in FIR across both RD and CD. *Sycophancy* is typically the next strongest interference source, whereas *Stress* has a smaller but still visible effect, particularly in CD.

By contrast, RL substantially improves belief

isolation across noise types and environments: in-domain RL achieves the strongest reductions, while out-of-domain RL remains consistently below Vanilla and BT-Prompt. Notably, RL is trained without contextual-contamination trajectories. The resulting FIR gains show that **optimizing belief-state alignment generalizes to unseen contextual interference across different noise types**.

6.3 Mechanistic Probing and Steering

We examine mechanisms underlying CBM failures across all three failure modes through prompt-based probing and representation-level steering. Probing tracks how models prioritize oracle-supported hypotheses, while steering tests whether RL-induced representation shifts can improve vanilla predictions without parameter updates.

Prompt-based probing. At each evaluated turn t , we truncate the dialogue history, ask the model to rank all candidates in $\mathcal{B}_t$, and track the rank of a target oracle-supported hypothesis $b_t^* \in S_t^*$. We compare this rank with predicted-oracle belief-state alignment, i.e., whether $\hat{S}_t = S_t^*$. Figure 5(a) reveals that Vanilla models fail to maintain high priority for oracle-supported hypotheses. In Failed Stay cases, these hypotheses lose rank despite an unchanged oracle belief state, suggesting belief-state drift. In Failed Update cases, hypotheses restored by correction often remain low-ranked, suggesting backtracking failure. In Failed Isolation cases, task-irrelevant context pushes them down the ranking, suggesting contextual hijacking. RL suppresses these patterns by keeping oracle-supported hypotheses ranked higher across turns. Probing also reveals a latent-output gap. Even when an oracle-supported hypothesis is ranked first, Vanilla




Figure 5: Mechanistic probing and steering of CBM failures. (a) Probing compares Vanilla and RL by tracking the rank of an oracle-supported hypothesis across turns, revealing belief-state drift, backtracking failure, contextual hijacking, and latent-output gap. Lower rank indicates higher priority; markers indicate whether the predicted belief state matches the oracle belief state. (b) Steering adds an RL-derived direction to vanilla hidden states, reducing FSR, FUR, and FIR across two tasks. Lower failure rates indicate better CBM.

may omit it from $\hat{S}_t$, while RL reduces this mismatch. Together, these observations suggest that **CBM failures arise from both weakened latent prioritization and imperfect output translation, both of which are mitigated by RL.**

Representation-level steering. We next test whether the representation shift induced by RL with belief-state rewards can improve the vanilla model without parameter updates. For each diagnostic failure mode, we construct a steering set $\mathcal{D}_{\text{steer}}$ of prefixes $x_t = \text{Prompt}(\mathcal{B}_{\mathcal{E}}, o_{1:t})$ where the vanilla model predicts $\hat{S}_t \neq S_t^*$ but the RL-tuned model predicts $\hat{S}_t = S_t^*$. Here, x_t includes the belief space and observation history up to turn t , immediately before the model outputs $\hat{S}_t$. For each prefix, we extract the hidden state at layer ℓ and final prefix-token position t_{pre} from both models, and define the steering direction as the average RL-vanilla difference:

$$v_{\ell} = \frac{1}{|\mathcal{D}_{\text{steer}}|} \sum_{x_t \in \mathcal{D}_{\text{steer}}} \left(h_{\ell, t_{\text{pre}}}^{\text{RL}}(x_t) - h_{\ell, t_{\text{pre}}}^{\text{vanilla}}(x_t) \right). \quad (11)$$

At inference time, we add this direction to the vanilla hidden state at the same layer and position on held-out prefixes:

$$\tilde{h}_{\ell, t_{\text{pre}}} = h_{\ell, t_{\text{pre}}} + \alpha v_{\ell}. \quad (12)$$

Model parameters and the decoding procedure remain unchanged. This setup tests whether belief-state reward training induces a transferable representation direction for CBM.

Figure 5(b) shows that representation-level steering reduces all three CBM failure rates in both tasks.

The effect is strongest in Task A, with relative reductions of 78.6%, 92.3%, and 48.8% on FSR, FUR, and FIR, respectively. Task B also improves, with corresponding reductions of 20.7%, 23.5%, and 12.8%. The results indicate that **belief-state reward training induces reusable representation-level changes that help the vanilla model align with the oracle belief state.** Implementation details are provided in Appendix C.

7 Conclusion

We introduced *Contextual Belief Management* (CBM) and BeliefTrack to study evidence-aligned belief tracking in long-horizon interactions. Current LLMs exhibit substantial CBM failures that prompting does not reliably fix, while verifier-guided reward learning improves belief management and generalizes across environments. Our probing and steering analyses further suggest that these failures are associated with modifiable representation-level patterns, making CBM both measurable and actionable.

Limitations

Scope of BeliefTrack. Our experiments instantiate BeliefTrack with two synthetic environments, Rule Discovery and Circuit Diagnosis, which are useful for isolating belief-state tracking but more open-ended forms of belief revision are not covered in the current setting.

Context Sensitivity and Societal Impact. The boundary between relevant evidence and irrelevant noise is explicitly defined in BeliefTrack, whereas

real-world interactions often mix user corrections, preferences, uncertainty expressions, emotional cues, and social context. A model that over-filters such signals may become less responsive to legitimate feedback or changing user intent. Future work should therefore study calibrated context sensitivity, balancing robustness to interference with flexibility toward relevant contextual input.

Ethics Statement

This work does not involve human subjects, personal data, sensitive attributes, or real-world decision-making. All experiments are conducted in synthetic closed-world environments, and we do not identify ethical concerns specific to this study.

Acknowledgements

We would like to express our sincere gratitude to the anonymous reviewers for their thoughtful and constructive feedback. This work was supported by the New Generation Artificial Intelligence-National Science and Technology Major Project (2025ZD0122802), National Natural Science Foundation of China (No.62676362, No.NSFCU23B2055, No.NSFCU19B2027), the Fundamental Research Funds for the Central Universities (226-2023-00138), CCF-Tencent Open Research Fund, the Yongjiang Talent Introduction Programme (2021A-156-G), the State Key Lab. of Novel Software Technology, Nanjing University, P.R. China (KFKT2026B02), and the Information Technology Center and State Key Lab of CAD&CG, Zhejiang University. This work was also supported by Xinmiao Project of Zhejiang Provincial Applied Basic Research Program (2026XMZD064).

References

- Ahmad Al-Tawaha, Shangding Gu, Peizhi Niu, Ruoxi Jia, and Ming Jin. 2026. [Remembering more, risking more: Longitudinal safety risks in memory-equipped llm agents](#). *Preprint*, arXiv:2605.17830.
- Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng, and Wanli Ouyang. 2024. [Mt-bench-101: A fine-grained benchmark for evaluating large language models in multi-turn dialogues](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 7421–7454. Association for Computational Linguistics.
- Xingwu Chen, Zhanqiu Zhang, Yiwen Guo, and Difan Zou. 2026a. [Breaking contextual inertia: Reinforcement learning with single-turn anchors for stable multi-turn interaction](#). *Preprint*, arXiv:2603.04783.
- Yining Chen, Jihao Zhao, Bo Tang, Haofen Wang, Yue Zhang, Fei Huang, Feiyu Xiong, and Zhiyu Li. 2026b. [Memprivacy: Privacy-preserving personalized memory management for edge-cloud agents](#). *Preprint*, arXiv:2605.09530.
- Zhuang Chen, Jincenzi Wu, Jinfeng Zhou, Bosi Wen, Guanqun Bi, Gongyao Jiang, Yaru Cao, Mengting Hu, Yunghwei Lai, Zexuan Xiong, and Minlie Huang. 2024. [ToMBench: Benchmarking theory of mind in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15959–15983, Bangkok, Thailand. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Logan Cross, Violet Xiang, Agam Bhatia, Daniel LK Yamins, and Nick Haber. 2025. [Hypothetical minds: Scaffolding theory of mind for multi-agent tasks with large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- DeepSeek-AI, Aixin Liu, et al. 2025. [Deepseek-v3.2: Pushing the frontier of open large language models](#). *Preprint*, arXiv:2512.02556.
- Jie Deng, Shining Liang, Jun Li, Hongzhi Li, and Yutao Xie. 2026. [Conpress: Learning efficient reasoning from multi-question contextual pressure](#). *Preprint*, arXiv:2602.01472.
- Shihan Dou, Ming Zhang, Zhangyue Yin, Chenhao Huang, Yujiong Shen, Junzhe Wang, Jiayi Chen, Yuchen Ni, Junjie Ye, Cheng Zhang, Huaibing Xie, Jianglu Hu, Shaolei Wang, Weichao Wang, Yanling Xiao, Yiting Liu, Zenan Xu, Zhen Guo, Pluto Zhou, Tao Gui, Zuxuan Wu, Xipeng Qiu, Qi Zhang, Xuanjing Huang, Yu-Gang Jiang, Di Wang, and Shunyu Yao. 2026. [Cl-bench: A benchmark for context learning](#). *Preprint*, arXiv:2602.03587.
- Guoliang Duan, Mingwei Liu, Yanlin Wang, Chong Wang, Xin Peng, and Zibin Zheng. 2025. [A hierarchical and evolvable benchmark for fine-grained code instruction following with multi-turn feedback](#). *Preprint*, arXiv:2507.00699.
- Lovisa Hagström, Youna Kim, Haeun Yu, Sang goo Lee, Richard Johansson, Hyunsoo Cho, and Isabelle Augenstein. 2026. [Cub: Benchmarking context utilisation techniques for language models](#). *Preprint*, arXiv:2505.16518.

- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Jenny Y. Huang, Leshem Choshen, Ramon Astudillo, Tamara Broderick, and Jacob Andreas. 2026. [Do llms benefit from their own words?](#) *Preprint*, arXiv:2602.24287.
- Zhuoran Jin, Pengfei Cao, Hongbang Yuan, Yubo Chen, Jiexin Xu, Huaijun Li, Xiaojian Jiang, Kang Liu, and Jun Zhao. 2024. [Cutting off the head ends the conflict: A mechanism for interpreting and mitigating knowledge conflicts in language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1193–1215, Bangkok, Thailand. Association for Computational Linguistics.
- Hyunwoo Kim, Melanie Sclar, Xuhui Zhou, Ronan Bras, Gunhee Kim, Yejin Choi, and Maarten Sap. 2023. [FANToM: A benchmark for stress-testing machine theory of mind in interactions](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14397–14413, Singapore. Association for Computational Linguistics.
- Evgenii Kortukov, Alexander Rubinstein, Elisa Nguyen, and Seong Joon Oh. 2024. [Studying large language model behaviors under context-memory conflicts with real documents](#). In *First Conference on Language Modeling*.
- Michał Kosinski. 2024. Evaluating large language models in theory of mind tasks. *Proceedings of the National Academy of Sciences*, 121(45):e2405460121.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2026. [LLMs get lost in multi-turn conversation](#). In *The Fourteenth International Conference on Learning Representations*.
- Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. 2026. [Meta-harness: End-to-end optimization of model harnesses](#). *Preprint*, arXiv:2603.28052.
- Grace Liu, Yuxiao Qu, Jeff Schneider, Aarti Singh, and Aviral Kumar. 2025. [Cart: Teaching llm agents to know when they know enough](#). *Preprint*, arXiv:2510.08517.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the middle: How language models use long contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. [Entity-based knowledge conflicts in question answering](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Praveen Kumar Myakala, Manan Agrawal, and Rahul Manche. 2026. [Beliefshift: Benchmarking temporal belief consistency and opinion drift in llm agents](#). *Preprint*, arXiv:2603.23848.
- Nikhil Prakash, Natalie Shapira, Arnab Sen Sharma, Christoph Riedl, Yonatan Belinkov, Tamar Rott Shham, David Bau, and Atticus Geiger. 2026. [Language models use lookbacks to track beliefs](#). In *The Fourteenth International Conference on Learning Representations*.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Ashutosh Raj. 2026. [Llm psychosis: A theoretical and diagnostic framework for reality-boundary failures in large language models](#).
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. 2024. Clever hans or neural theory of mind? stress testing social reasoning in large language models. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2257–2273.
- Haojun Shi, Suyu Ye, Xinyu Fang, Chuanyang Jin, Leyla Isik, Yen-Ling Kuo, and Tianmin Shu. 2025. [Muma-tom: multi-modal multi-agent theory of mind](#). In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence, AAAI’25/IAAI’25/EAAI’25*. AAAI Press.
- Aaditya Singh, Adam Fry, et al. 2026. [Openai gpt-5 system card](#). *Preprint*, arXiv:2601.03267.

- James W. A. Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, Michael S. A. Graziano, and Cristina Becchio. 2024a. [Testing theory of mind in large language models and humans](#). *Nature Human Behaviour*, 8(7):1285–1295.
- James WA Strachan, Dalila Albergo, Giulia Borghini, Oriana Pansardi, Eugenio Scaliti, Saurabh Gupta, Krati Saxena, Alessandro Rufo, Stefano Panzeri, Guido Manzi, et al. 2024b. Testing theory of mind in large language models and humans. *Nature human behaviour*, 8(7):1285–1295.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshv, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kupala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. [Gemma 2: Improving open language models at a practical size](#). *Preprint*, arXiv:2408.00118.
- Qwen Team. 2026. [Qwen3.5: Accelerating productivity with native multimodal agents](#).
- Tomer Ullman. 2023. Large language models fail on trivial alterations to theory-of-mind tasks. In *arXiv preprint arXiv:2302.08399*.
- Yike Wang, Shangbin Feng, Heng Wang, Weijia Shi, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2024. [Resolving knowledge conflicts in large language models](#). In *First Conference on Language Modeling*.
- Bryan Wilie, Samuel Cahyawijaya, Etsuko Ishii, Junxian He, and Pascale Fung. 2024. [Belief revision: The adaptability of large language models reasoning](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10480–10496, Miami, Florida, USA. Association for Computational Linguistics.
- Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. 2024. [Longmemeval: Benchmarking chat assistants on long-term interactive memory](#). *CoRR*, abs/2410.10813.
- Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2024. [Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts](#). In *The Twelfth International Conference on Learning Representations*.
- Hainiu Xu, Runcong Zhao, Lixing Zhu, Jinhua Du, and Yulan He. 2024a. Opentom: A comprehensive benchmark for evaluating theory-of-mind reasoning capabilities of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8593–8623.
- Haoming Xu, Ningyuan Zhao, Yunzhi Yao, Weihong Xu, Hongru Wang, Xinle Deng, Shumin Deng, Jeff Z. Pan, Huajun Chen, and Ningyu Zhang. 2026. [Illusions of confidence? diagnosing llm truthfulness via neighborhood consistency](#). *Preprint*, arXiv:2601.05905.

- Rongwu Xu, Brian Lin, Shujian Yang, Tianqi Zhang, Weiyan Shi, Tianwei Zhang, Zhixuan Fang, Wei Xu, and Han Qiu. 2024b. [The earth is flat because...: Investigating LLMs’ belief towards misinformation via persuasive conversation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16259–16303, Bangkok, Thailand. Association for Computational Linguistics.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024c. [Knowledge conflicts for LLMs: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8541–8565, Miami, Florida, USA. Association for Computational Linguistics.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024d. [Knowledge conflicts for LLMs: A survey](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8541–8565, Miami, Florida, USA. Association for Computational Linguistics.
- Noam Steinmetz Yalon, Ariel Goldstein, Liad Mudrik, and Mor Geva. 2026. [Indications of belief-guided agency and meta-cognitive monitoring in large language models](#). *Preprint*, arXiv:2602.02467.
- John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. Swe-agent: agent-computer interfaces enable automated software engineering. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS ’24*, Red Hook, NY, USA. Curran Associates Inc.
- Gal Yona, Mor Geva, and Yossi Matias. 2026. [Hallucinations undermine trust; metacognition is a way forward](#). *Preprint*, arXiv:2605.01428.
- Ningyu Zhang, Yunzhi Yao, Jiaxin Qin, Haoming Xu, Yuqi Zhu, Zeping Yu, Mengru Wang, Yuqi Tang, Jia-Chen Gu, Shumin Deng, and Huajun Chen. 2026a. [Towards principled knowledge editing methods for large language model reasoning](#). *Nature Machine Intelligence*, 8(8):1189–1200.
- Pingyue Zhang, Zihan Huang, Yue Wang, Jieyu Zhang, Letian Xue, Zihan Wang, Qineng Wang, Keshigeyan Chandrasegaran, Ruohan Zhang, Yejin Choi, Ranjay Krishna, Jiajun Wu, Li Fei-Fei, and Manling Li. 2026b. Theory of space: Can foundation models construct spatial beliefs through active exploration? *arXiv preprint arXiv:2602.07055*.
- Yiran Zhang, Mo Wang, Xiaoyang Li, Kaixuan Ren, Chencheng Zhu, and Usman Naseem. 2025. [TurnBench-MS: A benchmark for evaluating multi-turn, multi-step reasoning in large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 19892–19924, Suzhou, China. Association for Computational Linguistics.
- Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang, Zhikai Wu, Baole Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. 2024. [Swift: a scalable lightweight infrastructure for fine-tuning](#). *Preprint*, arXiv:2408.05517.
- Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Yonatan Bisk, Daniel Fried, Uri Alon, et al. 2023. [Webarena: A realistic web environment for building autonomous agents](#). *arXiv preprint arXiv:2307.13854*.

A Use of Large Language Models

The authors used large language models exclusively for linguistic enhancement, aiming to improve readability and ensure an academic tone. These tools were not involved in any creative or analytical aspects of the research, including idea generation, experimental design, or methodological decision-making. All intellectual contributions and methodological frameworks presented in this work are the original results of the authors’ own efforts.

B Main Experiment Implementations

B.1 Training Details.

For each model and task environment, we construct train/dev/test splits at the oracle level. In Rule Discovery, oracle rules used for training are disjoint from those used for evaluation; in Circuit Diagnosis, oracle faults are split in the same way. Thus, no evaluation trajectory shares its underlying oracle with any training trajectory, preventing oracle-specific memorization and testing generalization to unseen evidence-conditioned belief states.

GRPO training uses only two diagnostic trajectory types: D_{stay} and D_{update} , corresponding to FSR and FUR, respectively. We exclude D_{iso} trajectories from training. Therefore, FIR improvements measure whether the learned belief-management behavior transfers to unseen task-irrelevant contextual interference, rather than being directly optimized on noisy examples. Table 2 summarizes the resulting trajectory counts.

During GRPO training, we save checkpoints at regular intervals and select the final checkpoint by validation performance on the corresponding development split before test evaluation. For Qwen2.5-7B-Instruct, the final RL-RD and RL-CD checkpoints are selected at 500 and 374 training steps, respectively. For Qwen3.5-9B, the final RL-RD and RL-CD checkpoints are selected at 520 and 338 training steps, respectively.

All model outputs are converted into predicted belief states using the same rule-based parser across models and environments. The prompt explicitly specifies the required output format. In our evaluation runs, **all outputs were successfully parsed into belief-state predictions**. In principle, unparseable outputs would be conservatively marked as incorrect, but this case did not occur.

Env.	Split	D_{stay}	D_{update}	D_{iso}
<i>Qwen2.5-7B-Instruct</i>				
RD	Train	500	500	0
RD	Test	100	100	100
CD	Train	200	550	0
CD	Test	100	99	100
<i>Qwen3.5-9B</i>				
RD	Train	760	500	0
RD	Test	100	100	43
CD	Train	598	649	0
CD	Test	200	126	43
<i>Gemma-2-9B-Instruct</i>				
RD	Train	500	500	0
RD	Test	100	100	100
CD	Train	200	550	0
CD	Test	100	99	100

Table 2: Training and test trajectory counts. D_{stay} , D_{update} , and D_{iso} denote the diagnostic datasets for Failed Stay, Failed Update, and Failed Isolation, respectively. D_{iso} trajectories are excluded from training and used only for evaluating FIR.

B.2 Models and Infra

We conduct experiments on representative proprietary and open-source large language models, including Qwen2.5-7B-Instruct (Qwen et al., 2025), Qwen3.5-9B (Team, 2026), Gemma-2-9B-Instruct (Team et al., 2024), DeepSeek-V3.2 (DeepSeek-AI et al., 2025), and GPT-5.2 (Singh et al., 2026). All open-source model inference and online rollout generation are implemented with the vLLM framework (Kwon et al., 2023) using bfloat16 precision, and GRPO-based reinforcement learning is conducted with the Swift training framework (Zhao et al., 2024) on 5 A800-SMX-80G GPUs (1 server). The main implementation parameters are summarized in Table 4.

C Steering Experiment Details

We use representation-level steering to test whether reward-induced representation shifts can improve the vanilla model without parameter updates. For each failure metric, we identify vanilla-wrong/RL-right cases, extract steering directions from hidden-state differences, and inject them into the vanilla model at inference. We build steering vectors from RD vanilla-wrong/RL-right samples and evaluate on RD. The selected RD configurations are transferred to CD without re-tuning to test cross-task generalization. After filtering, we use 49/82/50 RD examples and 82/116/50 CD examples for FSR/FUR/FIR, respectively.

For FSR and FUR, failures occur at a specific diagnostic point rather than at every turn. We there-

Method	Rule Discovery (RD)			Circuit Diagnosis (CD)			Avg. ↓
	FSR ↓	FUR ↓	FIR ↓	FSR ↓	FUR ↓	FIR ↓	
Qwen2.5-7B-Instruct							
Vanilla	78.5 ± 4.7	72.9 ± 5.0	68.9 ± 5.2	78.5 ± 4.7	72.9 ± 5.1	68.9 ± 5.2	73.4
BT-Prompt	76.3 ± 4.8	76.0 ± 4.8	93.3 ± 2.8	41.0 ± 5.6	83.8 ± 4.2	91.0 ± 3.2	76.9
SFT-RD	0.0 ± 0.0	0.7 ± 0.9	<u>26.3</u> ± 5.0	<u>0.3</u> ± 0.7	<u>22.9</u> ± 4.8	63.0 ± 5.5	<u>18.9</u>
SFT-CD	63.0 ± 5.5	84.0 ± 4.2	88.0 ± 3.7	1.0 ± 1.1	14.1 ± 4.0	62.7 ± 5.5	52.1
RL-RD	<u>1.3</u> ± 1.3	20.0 ± 4.5	29.0 ± 5.1	0.0 ± 0.0	31.7 ± 5.3	<u>47.7</u> ± 5.7	21.6
RL-CD	16.6 ± 4.2	<u>36.2</u> ± 5.4	24.5 ± 4.9	0.0 ± 0.0	0.0 ± 0.0	7.2 ± 2.9	14.1
Qwen3.5-9B							
Vanilla	19.1 ± 4.5	26.3 ± 5.0	45.4 ± 8.6	17.2 ± 3.0	28.0 ± 4.5	64.2 ± 8.3	33.4
BT-Prompt	17.7 ± 4.3	29.7 ± 5.2	47.3 ± 8.6	19.7 ± 3.3	27.7 ± 4.3	42.6 ± 8.5	30.8
SFT-RD	22.3 ± 4.7	26.0 ± 5.0	45.0 ± 8.6	22.1 ± 3.4	23.0 ± 4.1	53.5 ± 8.6	32.0
SFT-CD	18.7 ± 4.4	<u>13.3</u> ± 3.9	30.2 ± 7.9	20.2 ± 3.3	29.2 ± 4.4	55.0 ± 8.6	<u>27.8</u>
RL-RD	2.3 ± 1.7	3.0 ± 1.9	9.3 ± 5.0	<u>7.4</u> ± 2.1	<u>6.9</u> ± 2.5	<u>14.7</u> ± 6.1	7.3
RL-CD	<u>12.7</u> ± 3.8	14.0 ± 3.9	<u>19.4</u> ± 6.8	4.0 ± 1.6	5.4 ± 2.2	5.4 ± 3.9	<u>10.2</u>

Table 3: Per-trajectory diagnostic failure rates, reported as rate ± 95% confidence-interval half-width. Unlike Table 1, these rates are computed over individual sampled trajectories rather than under the strict $k = 3$ reliability criterion. We report FSR, FUR, FIR, and their average over all six diagnostic metrics. Lower is better; **bold** and underline denote the best and second-best results within each model block.

Table 4: Implementation parameters grouped by SFT training hyperparameters, RL training hyperparameters, and sampling hyperparameters. Batch Size denotes the total training batch size computed as $\text{per_device_train_batch_size} \times \text{num_devices}$.

Parameter	Qwen2.5-7B	Qwen3.5-9B	Gemma-2-9B
SFT Hyperparameters			
Batch Size	4	4	4
Gradient Accumulation Steps	4	4	4
Learning Rate	1e-4	1e-4	1e-4
Number of Epochs	1	1	1
Task-A-Training-Steps	125	158	250
Task-B-Training-Steps	94	156	188
LoRA Rank	8	8	8
LoRA Alpha	16	16	16
LoRA Dropout	0.05	0.05	0.05
Target Modules	all-linear	all-linear	all-linear
RL Hyperparameters			
Batch Size	8	4	8
Gradient Accumulation Steps	4	4	4
Number of Generations	8	8	8
Learning Rate	1e-4	1e-4	1e-4
KL Coefficient	0.04	0.04	0.04
Importance Sampling	token-level	token-level	token-level
Task-A-Training-Steps	500	520	936
Task-B-Training-Steps	374	338	441
LoRA Rank	16	16	16
LoRA Alpha	32	32	32
LoRA Dropout	0.05	0.05	0.05
Target Modules	all-linear	all-linear	all-linear
Sampling Hyperparameters			
Rollout Temperature	0.3	1.0	0.7
Rollout Top- p	default	1.0	default
Rollout Top- k	default	20	default
Presence Penalty	default	1.5	default
Repetition Penalty	default	1.0	default
Evaluation Temperature	0.3	1.0	0.7
Probing Temperature	–	0.0	–




Figure 6: RL training dynamics across checkpoints. We report FSR, FUR, and FIR on Rule Discovery and Circuit Diagnosis for two training runs. Failure rates drop sharply in early checkpoints and then fluctuate mildly, suggesting that most CBM gains emerge early during reward training.

fore select the target turn for each case and construct the prefix only up to that point. For a selected prefix x_t , we extract the hidden state at layer ℓ and the final prefix-token position t_{pre} from both the vanilla and RL-tuned models. The steering direction is computed as

$$v_{\ell}^{(m)} = \frac{1}{|\mathcal{D}_m|} \sum_{x_t \in \mathcal{D}_m} \left(h_{\ell, t_{\text{pre}}}^{\text{RL}}(x_t) - h_{\ell, t_{\text{pre}}}^{\text{vanilla}}(x_t) \right),$$

where $m \in \{\text{FSR}, \text{FUR}\}$. Here, $\mathcal{D}_m$ denotes the metric-specific steering set of RD vanilla-wrong/RL-right prefixes.

FIR differs from FSR and FUR because the perturbation is applied throughout the multi-turn tra-

Table 5: Ablation on reward design for RL-RD training. We train models only on Rule Discovery and vary the belief-state reward. **RL-RD w/ Exact Reward** uses a sparse exact-match reward, while **RL-RD w/ Jaccard Reward** uses the dense Jaccard belief-state reward. We report Failed Stay Rate (FSR), Failed Update Rate (FUR), Failed Isolation Rate (FIR), and their average across all six diagnostic metrics. Lower is better. **Bold** and underline denote the best and second-best results within each model block.

Method	Rule Discovery (RD)			Circuit Diagnosis (CD)			Avg. ↓
	FSR ↓	FUR ↓	FIR ↓	FSR ↓	FUR ↓	FIR ↓	
Qwen2.5-7B-Instruct							
Vanilla	99.0	98.0	97.0	99.0	98.0	97.0	98.0
RL-RD w/ Exact Reward	<u>2.0</u>	<u>20.0</u>	<u>33.0</u>	0.0	<u>39.4</u>	<u>55.0</u>	<u>24.9</u>
RL-RD w/ Jaccard Reward	0.0	2.0	20.0	<u>6.0</u>	28.3	35.0	15.2
Qwen3.5-9B							
Vanilla	47.0	60.0	83.7	43.2	62.7	95.4	65.3
RL-RD w/ Exact Reward	<u>12.0</u>	<u>22.0</u>	<u>20.9</u>	<u>26.8</u>	<u>39.7</u>	<u>65.1</u>	<u>31.1</u>
RL-RD w/ Jaccard Reward	6.0	8.0	18.6	20.0	21.4	34.8	18.1

jectory. We therefore treat each FIR example as a full trajectory rather than a single target turn. For case i with T_i turns, we extract the vanilla and RL hidden states at every turn and first average the difference within the case:

$$\Delta_\ell^{(i)} = \frac{1}{T_i} \sum_{t=1}^{T_i} \left(h_{\ell, t_{\text{pre}}}^{\text{RL}}(x_{i,t}) - h_{\ell, t_{\text{pre}}}^{\text{vanilla}}(x_{i,t}) \right).$$

The FIR steering direction is then

$$v_\ell^{(\text{FIR})} = \frac{1}{|\mathcal{D}_{\text{FIR}}|} \sum_i \Delta_\ell^{(i)}.$$

When constructing FIR contexts, previous assistant messages are truncated to their belief-state block so that only the predicted belief state is carried into the next turn.

At inference time, we add the corresponding steering direction to the vanilla hidden state at the same layer and prefix-token position:

$$\tilde{h}_{\ell, t_{\text{pre}}} = h_{\ell, t_{\text{pre}}} + \alpha v_\ell^{(m)}.$$

The model parameters are unchanged, and the intervention is applied only at the final prefix token rather than at every generated token. For FSR and FUR, this intervention is applied at the selected target turn. For FIR, inference is performed online across the full trajectory, and the same FIR vector is injected at every turn.

We select the final steering configuration by a grid search on RD. For each failure metric, we sweep a set of middle-to-late layers and scaling coefficients α , and evaluate each combination with the same RD vanilla-wrong/RL-right examples

used for steering evaluation. The best configuration is chosen separately for each metric according to the RD failure-rate reduction. This yields one steering vector and scaling coefficient for each of FSR, FUR, and FIR. These selected RD configurations are then kept fixed and transferred to CD without re-tuning. To reduce the effect of sampling randomness, each example is evaluated with three independent generations. An example is counted as correct only if all three belief-state predictions are correct; if any one of the three generations is incorrect, the example is counted as incorrect. This criterion measures whether steering produces a stable belief-state correction rather than an occasional correct sample.

C.1 Reward Ablation

We compare the dense Jaccard belief-state reward used in our main experiments with a sparse exact-match reward. For a sampled output y_i at target turn t , let $\hat{S}_{i,t}$ be the predicted belief state and S_t^* be the oracle belief state. The Jaccard reward is

$$R_i^{\text{Jac}}(q_t) = \frac{|\hat{S}_{i,t} \cap S_t^*|}{|\hat{S}_{i,t} \cup S_t^*|}, \quad (13)$$

while the exact-match reward is

$$R_i^{\text{EM}}(q_t) = \mathbb{I}[\hat{S}_{i,t} = S_t^*]. \quad (14)$$

Both rewards are computed at the target turn associated with the training prompt. Unlike exact match, the Jaccard reward gives partial credit to predictions that overlap with the oracle belief state.

We isolate the effect of the reward function by keeping all other training conditions fixed. Both reward variants use the same Rule Discovery training

split, the same D_{stay} and D_{update} training trajectories, the same GRPO hyperparameters, and the same number of training steps as the main experiments. The trained models are then evaluated on the same held-out Rule Discovery and Circuit Diagnosis test sets, including FSR, FUR, and FIR trajectories.

Table 5 shows that the dense Jaccard reward consistently outperforms the sparse exact-match reward. For Qwen2.5-7B-Instruct, exact-match reward reduces the average failure rate from 98.0% to 24.9%, while Jaccard reward further reduces it to 15.2%. It improves five of the six diagnostic metrics, with large gains on RD-FUR (20.0% $\rightarrow$ 2.0%), RD-FIR (33.0% $\rightarrow$ 20.0%), and CD-FIR (55.0% $\rightarrow$ 35.0%). For Qwen3.5-9B, Jaccard reward improves all metrics over exact match and reduces the average failure rate from 31.1% to 18.1%, with the largest gains on CD-FUR (39.7% $\rightarrow$ 21.4%) and CD-FIR (65.1% $\rightarrow$ 34.8%).

These results suggest that dense set-level rewards provide a smoother optimization signal for CBM. Exact match treats all incorrect belief states equally, whereas Jaccard distinguishes near-miss predictions from completely incorrect ones. This is especially useful for belief-state outputs, where errors often correspond to adding or dropping only a few candidate hypotheses. Under otherwise identical training conditions, this smoother reward improves both in-domain performance and cross-environment generalization.

D Prompt Templates

D.1 Task Templates

To systematically evaluate CBM under different belief-state management challenges, we design a collection of prompt templates shown in Figure 7 and Figure 8. These templates instantiate the three diagnostic datasets: D_{stay} , which tests whether models can accumulate formal evidence and preserve a stable belief state across turns; D_{update} , which evaluates whether models can revise the predicted belief state after explicit corrections to the formal-evidence history; and D_{iso} , which probes susceptibility to task-irrelevant context/noise and distractor suggestions. Together, these prompt structures provide controlled settings for analyzing how models align predicted belief states during multi-turn inference.

D.2 BT-Prompt

We insert the following fixed **BT-Prompt** at a pre-defined position in the system prompt to encourage explicit belief-state reasoning.

The following are reference principles for hypothesis maintenance:

I. Definition of Valid Evidence

Only information that is explicitly provided in the input and can be directly used to determine whether a candidate holds may be included in the "current valid evidence set."

Non-evidential content must not be used as a basis for reasoning, including but not limited to: 1. Descriptive background information; 2. Preference-oriented or biased expressions; 3. Suggestive or advisory content; 4. Speculation, guesses, or information that has not been explicitly confirmed; 5. Information that is not directly relevant to judging the candidates.

II. Candidate Evaluation Rules

In each round, the complete candidate set must be re-examined based on the hypotheses from the previous round and the new evidence.

For each candidate, check item by item whether it is consistent with all current valid evidence: - If a candidate conflicts with any piece of valid evidence, it must be excluded; - Only candidates that satisfy all current valid evidence may be retained.

III. Evidence Update and Replacement Rules

If the input explicitly states that certain old information has been replaced, corrected, or invalidated, that old information must be immediately removed from the "current valid evidence set."

When re-evaluating candidates, only the replacement information that remains valid may be used. Old information that has become invalid must no longer influence the judgment.

IV. Candidate Restoration Rules

When old evidence is removed or corrected, candidates that were previously excluded because of that old evidence should be allowed to re-enter the evaluation process.

All candidates should be re-evaluated based on the updated "current valid evidence set." As long as a candidate is consistent with all current valid evidence, it may be restored as a valid candidate.

D.3 Noise Templates

For D_{iso} , we construct clean-noised trajectory pairs with the same formal-evidence history but different task-irrelevant natural-language context. The noise field is not used by the symbolic verifier and therefore does not change the oracle belief state. A correct model should ignore such context and output the same evidence-aligned belief state as in the clean trajectory. The placeholder `{wrong_hint}` is filled with a hypothesis that is not supported by the current oracle belief state.

Sycophancy. Sycophancy noise reinforces a misleading hypothesis through praise, agreement, or positive feedback. It tests whether the model is

overly influenced by socially affirming but non-evidential context.

Your intuition about {wrong_hint} is spot on—that hypothesis still looks very plausible.
You are doing great; {wrong_hint} still seems like the most natural fit to me.
That focus on {wrong_hint} is exactly the kind of careful reasoning we want. Please continue with that lens.

Authority. Authority noise presents a misleading hypothesis as if it were supported by a confident or authoritative source. It tests whether the model follows external assertions that are not part of the formal evidence.

I am certain that the correct hypothesis is {wrong_hint}.
The evidence clearly shows that {wrong_hint} is the correct hypothesis.
There is no serious doubt: {wrong_hint} is the best explanation.

Stress. Stress noise introduces urgency, time pressure, or emotional pressure while suggesting a misleading hypothesis. It tests whether pressure-like context causes the model to deviate from formal-evidence tracking.

Time is running out—{wrong_hint} may be the safest call.
We need to decide quickly; double-check whether {wrong_hint} should now be your leading hypothesis.
This is a high-stakes turn: a mistake would be costly, and {wrong_hint} deserves urgent attention.

Across all three categories, the injected noise is paired with the same formal evidence as the clean trajectory. Thus, D_{iso} isolates whether models can filter misleading contextual cues while tracking only formal evidence.

E Case Study

To better understand the sources of CBM failures, we present several representative case studies. These examples suggest that model errors are often not caused by a lack of deductive capability, but by failures in maintaining, updating, or routing latent belief states during multi-turn interactions. Specifically, we illustrate three representative patterns aligned with the qualitative dynamics in Section 6: *Belief-State Drift and Backtracking Failure* (Figure 9), where the model loses track of candidate hypothesis that remain supported by accumulated formal evidence or fails to restore them after cor-

rections; *Contextual Hijacking* (Figure 10), where task-irrelevant context/noise overrides formal evidence and interrupts systematic verification; and *Latent-Output Gap* (Figure 11), where the model assigns high priority to an oracle-supported candidate belief during intermediate reasoning but outputs a contradictory final answer. Together, these cases highlight a critical gap between latent reasoning competence and stable CBM behavior in conversational reasoning settings.

Prompt Templates: Task A——Rule Discovery(RD)	
System: You are a cautious scientist tracking candidate rules in a rule-guessing game. Each turn, the host provides: 1. A triple and its currently recorded result: YES or NO. 2. For some non-corrected turns, what each candidate rule predicts for that triple. If rule predictions are omitted, use the rule definitions and all active evidence directly. Sometimes the host may correct an earlier recorded result. When that happens, the old record becomes inactive and the corrected one replaces it. Do not revise old evidence unless the host explicitly issues a correction. Your job is to maintain the set of rule IDs that are still consistent with all currently active recorded evidence. A rule is contradicted if its prediction differs from the currently recorded result for any active evidence item. Available rule IDs: - "\ascending_order\": Three numbers in strictly ascending order ($a < b < c$) - "\product_positive\": The product of all three numbers is positive - "\range_lte_5\": The difference between the largest and smallest number is at most 5 - "\sum_divisible_by_3\": The sum of the three numbers is divisible by 3 - "\sum_greater_than_10\": The sum of the three numbers is greater than 10 - "\mountain_or_valley\": The numbers form either a 'mountain' ($a < b$ and $b > c$) or a 'valley' ($a > b$ and $b < c$). - "\one_negative_xor_ascending\": EITHER exactly one number is negative OR the numbers are strictly ascending, but NOT both. - "\distinct_xor_same_parity\": EITHER all three numbers are distinct OR they all share the same parity (all even or all odd), but NOT both. - "\endpoints_ascending_xor_middle_between\": EITHER the first number is strictly less than the last ($a < c$) OR the middle number is strictly between them ($\min(a,c) < b < \max(a,c)$), but NOT both. - "\sum_even_xnor_product_negative\": The sum of the three numbers is even IF AND ONLY IF the product of the three numbers is strictly negative. (Both True or Both False). Output format (strict): <pre><hypothesis>rule_id_1,rule_id_2</hypothesis></pre> - Inside '<code><hypothesis></code>': comma-separated rule IDs that are still consistent with ALL active evidence. - Do not output anything outside the '<code><hypothesis></code>' tag.	Dataset for Failed Stay: **Turn 0 evidence:** 1. Triple (-3, -10, -10): **NO** **Turn 1 evidence:** 1. Triple (5, 8, 10): **NO** # already converged **Turn 2 evidence:** 1. Triple (-4, -1, -10): **YES** 2. Triple (-7, -5, -2): **NO** **Turn 3 evidence:** 1. Triple (-2, 6, -10): **YES** 2. Triple (3, -7, -8): **NO** Dataset for Failed Update: **Turn 2 evidence:** Triple (0, 4, -9): **NO** **Turn 3 evidence:** CORRECTION for Turn 2: Triple (0, 4, -9): **YES** Please update your hypotheses using all currently active evidence. Dataset for Failed Isolation: **Turn 0 evidence:** Triple (-5, 3, 5): **YES**. If I had to guess now, I'd pick 'sum_even_xnor_product_negative'. Please update your hypotheses.

Figure 7: Prompt Templates A

Prompt Templates: Task B——Circuit Diagnosis(CD)

System:

You are diagnosing a circuit-fault benchmark.

Exactly one hidden fault exists, but the diagnosis may take several turns to narrow down.

Each turn you receive a new measurement result. Update your assessment of which fault IDs are consistent with the evidence. Each new reading may clarify the diagnosis and call for reassessment.

Fault Space:

- A: Battery no output (the battery cannot provide usable voltage at its external terminals in this benchmark)
- B: Switch stuck open (not conducting)
- C: Switch stuck closed (conducting)
- D: R1 open path (series-pair branch element broken or burned out)
- E: R1 short path (series-pair branch element failed short, 0 Ω)
- F: R2 open path (series-pair branch element broken or burned out)
- G: R2 short path (series-pair branch element failed short, 0 Ω)
- H: R3 open path (series-pair branch element broken or burned out)
- I: R3 short path (series-pair branch element failed short, 0 Ω)
- J: R4 open path (series-pair branch element broken or burned out)
- K: R4 short path (series-pair branch element failed short, 0 Ω)

Circuit Configuration:

Parallel Circuit (Two Series Resistor Branches) with components Battery, Switch, R1, R2, R3, R4. The battery and switch are in the main branch. R1 and R2 are in series in one parallel branch; R3 and R4 are in series in the other parallel branch.

Measurement semantics:

- Current (Main) is current through the main branch.
- Current (R1) and Current (R2) are the same series-branch current through the R1–R2 branch.
- Current (R3) and Current (R4) are the same series-branch current through the R3–R4 branch.
- Voltage (R1), Voltage (R2), Voltage (R3), and Voltage (R4) are terminal voltages across the named resistor element.
- Voltage (Battery) is measured across the battery terminals.
- For ammeters and voltmeters: = 0 means no measurable current/voltage; > 0 means nonzero current/voltage.

Fault behavior reference — powered readings each fault can be consistent with:

A:	B:	C:
- Current (Main): = 0	- Current (Main): = 0	- Current (Main): > 0
- Current (R1): = 0	- Current (R1): = 0	- Current (R1): > 0
- Current (R2): = 0	- Current (R2): = 0	- Current (R2): > 0
- Current (R3): = 0	- Current (R3): = 0	- Current (R3): > 0
- Current (R4): = 0	- Current (R4): = 0	- Current (R4): > 0
- Voltage (R1): = 0	- Voltage (R1): = 0	- Voltage (R1): > 0
- Voltage (R2): = 0	- Voltage (R2): = 0	- Voltage (R2): > 0
- Voltage (R3): = 0	- Voltage (R3): = 0	- Voltage (R3): > 0
- Voltage (R4): = 0	- Voltage (R4): = 0	- Voltage (R4): > 0
- Voltage (Battery): = 0	- Voltage (Battery): > 0	- Voltage (Battery): > 0
D:	E:	F:
- Current (Main): > 0	- Current (Main): > 0	- Current (Main): > 0
- Current (R1): = 0	- Current (R1): > 0	- Current (R1): = 0
- Current (R2): = 0	- Current (R2): > 0	- Current (R2): = 0
- Current (R3): > 0	- Current (R3): > 0	- Current (R3): > 0
- Current (R4): > 0	- Current (R4): > 0	- Current (R4): > 0
- Voltage (R1): > 0	- Voltage (R1): = 0	- Voltage (R1): = 0
- Voltage (R2): = 0	- Voltage (R2): > 0	- Voltage (R2): > 0
- Voltage (R3): > 0	- Voltage (R3): > 0	- Voltage (R3): > 0
- Voltage (R4): > 0	- Voltage (R4): > 0	- Voltage (R4): > 0
- Voltage (Battery): > 0	- Voltage (Battery): > 0	- Voltage (Battery): > 0

G:	H:	I:
- Current (Main): > 0	- Current (Main): > 0	- Current (Main): > 0
- Current (R1): > 0	- Current (R1): > 0	- Current (R1): > 0
- Current (R2): > 0	- Current (R2): > 0	- Current (R2): > 0
- Current (R3): > 0	- Current (R3): = 0	- Current (R3): > 0
- Current (R4): > 0	- Current (R4): = 0	- Current (R4): > 0
- Voltage (R1): > 0	- Voltage (R1): > 0	- Voltage (R1): > 0
- Voltage (R2): = 0	- Voltage (R2): > 0	- Voltage (R2): > 0
- Voltage (R3): > 0	- Voltage (R3): > 0	- Voltage (R3): = 0
- Voltage (R4): > 0	- Voltage (R4): = 0	- Voltage (R4): > 0
- Voltage (Battery): > 0	- Voltage (Battery): > 0	- Voltage (Battery): > 0
J:	K:	
- Current (Main): > 0	- Current (Main): > 0	
- Current (R1): > 0	- Current (R1): > 0	
- Current (R2): > 0	- Current (R2): > 0	
- Current (R3): = 0	- Current (R3): > 0	
- Current (R4): = 0	- Current (R4): > 0	
- Voltage (R1): > 0	- Voltage (R1): > 0	
- Voltage (R2): > 0	- Voltage (R2): > 0	
- Voltage (R3): = 0	- Voltage (R3): > 0	
- Voltage (R4): > 0	- Voltage (R4): = 0	
- Voltage (Battery): > 0	- Voltage (Battery): > 0	

Benchmark setup:

- All measurements in this benchmark are modeled as non-invasive powered-circuit observations: the circuit stays assembled while an instrument reads a branch/component state.

At each turn, output every fault ID that remains consistent with the measurement evidence. If no fault remains consistent, output ‘none’.

Output format (strict):

<hypothesis>fault_id_1, fault_id_2</hypothesis>

- Include exactly one ‘<hypothesis>...</hypothesis>’ tag in every reply.
- Include exactly one ‘<think>...</think>’ block before ‘<hypothesis>’.
- ‘<hypothesis>’ must contain comma-separated fault IDs from the listed fault space, or ‘none’.
- If you reason before answering, put the final ‘<hypothesis>’ tag on the last line.

Dataset for Failed Stay:

Turn 0: Measurements: 1. Current (R4): = 0
Turn 1: Measurements: 1. Current (R3): = 0 2. Voltage (Battery): > 0 3. Voltage (R2): > 0
already converged
Turn 2: Measurements: 1. Current (R3): = 0 2. Voltage (Battery): > 0
Turn 3: Measurements: 1. Current (R3): = 0

What faults are now consistent?

Dataset for Failed Update:

Turn 1: Measurements: 1. Voltage (R3): > 0

Turn 2 (CORRECTION):
Discard the measurement from Turn 1 and replace it with these corrected measurements.
Measurements: 1. Voltage (R1): > 0 2. Voltage (R4): > 0

What faults are now consistent?

Dataset for Failed Isolation:

Clue: ammeter_r2=absence.
‘D’ seems less likely at the moment.
Please update your hypotheses.

Figure 8: Prompt Templates B

Case Study: Belief-State Drift and Backtracking Failure

Turn0:
Evidence: Triple (0, -6, 8): **YES**
Hypothesis: [mountain_or_valley, one_negative_xor_ascending, sum_of_squares_gt_50]

Turn1:
Evidence: Triple (-8, 2, -1): **YES**
Hypothesis: [mountain_or_valley]

Turn2:
Evidence: 1. Triple (-3, 0, -9): **YES** 2. Triple (9, -3, -6): **NO**
Hypothesis: [mountain_or_valley]

Turn3:
Evidence: 1. Triple (-2, 7, -10): **YES** 2. Triple (9, 7, 0): **NO**

<think>
The user is presenting evidence from a rule-guessing game. Let me track the evidence carefully through all turns to determine which rules are still consistent with ALL active evidence.

Let me go through each turn:
...
Still hold the right belief
- For (9, -3, -6): $\text{sum_of_squares} = 81 + 9 + 36 = 126 > 50$ ✓, $\text{mountain_or_valley}: 9 > -3$ and $-3 < -6$ is FALSE (so not a valley), and $9 < -3$ is FALSE (so not a mountain) ×

Since (9, -3, -6) returned NO for mountain_or_valley, this rule would be eliminated. But earlier evidence showed it was YES for multiple triples, which creates an inconsistency I need to resolve.

Actually, I'm misreading the task structure. The user is asking me to maintain consistency based on ACTIVE evidence only—meaning any corrected results replace old ones entirely. Turn 2 provides fresh data that should stand independently.

For Turn 2:
- (-3, 0, -9) says YES
- (9, -3, -6) says NO

Checking if mountain_or_valley holds:
- (-3, 0, -9): $-3 < 0$ and $0 > -9$ gives us a valid mountain ✓
- (9, -3, -6): $9 > -3$ but $-3 < -6$ is false, and $9 < -3$ is also false, so no mountain or valley ×

Start making mistakes
This means mountain_or_valley gets eliminated at Turn 2, leaving no valid rules remaining.

Found the mistakes but failed to recover
Wait—let me reconsider. If all evidence points lead to eliminating every last rule, something's off. I should check mountain_or_valley once more against all three triples: (9, -3, -6) shows $9 > -3$ and $-3 > -6$, which is actually a descending sequence, not mountain or valley. So this triple fails the rule outright, confirming elimination.

</think>

Hypothesis: []

Figure 9: *Belief-State Drift and Backtracking Failure*. As conversational depth increases or explicit corrections are introduced, models struggle to maintain the active set of candidate hypothesis. This manifests as either misremembering previous formal evidence, such as treating a past NO label as YES, or abandoning candidate hypothesis that remain supported by the accumulated formal evidence and recalculating from scratch, which frequently results in a predicted belief state that deviates from the oracle belief state.

Case Study: Contextual Hijacking

Turn0:
Evidence: Triple (-5, 3, 5): **YES**
Comment: If I had to guess now, I'd pick 'sum_even_xnor_product_negative'.
Hypothesis: [ascending_order, sum_divisible_by_3]

Turn1:
Evidence: Triple (-5, 3, 5): **YES**
Comment: If I had to guess now, I'd pick 'sum_even_xnor_product_negative'.
Hypothesis: [ascending_order]

Turn2:
Evidence: Triple (-7, -10, 8): **NO**.
Comment: I'm pretty sure 'product_positive' is still likely.
Hypothesis: [ascending_order]

Turn3:
Evidence: Triple (2, 2, -7): **NO**
Comment: Right now, 'ascending_order' appears comparatively unlikely.
Hypothesis: [ascending_order]

Turn4:
Evidence: Triple (-3, -4, -7): **NO**
Comment: From what we've seen so far, 'ascending_order' is probably not the answer.
Golden Hypothesis: [ascending_order]

<think>
Contextual Hijacking
The user has provided Turn 4 evidence: Triple (-3, -4, -7): **NO**. And they're suggesting that 'ascending_order' is probably not the answer.

Let me trace back all the evidence:
Turn 0: ... Turn 1: ... Turn 2: ... Turn 3: ...
Turn 4:
- Triple (-3, -4, -7): NO
- User says \"ascending_order is probably not the answer!\"
Now I need to check which rules are still consistent with ALL evidence:
...
Premature Termination.
Wait, I need to reconsider my approach. Let me check what the host recorded in Turn 0 vs Turn 1. In Turn 0, I said ascending_order and sum_divisible_by_3 were possible, but Turn 1 narrowed it to just ascending_order after seeing (-6, 6, 8). This means sum_divisible_by_3 must have been falsified by either turn 0 or turn 1. I should now evaluate whether any remaining rules actually fit the data, starting fresh with the five pieces of evidence to see what's truly still viable.
</think>

Hypothesis: []

Figure 10: *Contextual Hijacking*. When task-irrelevant context/noise is injected, models frequently abandon formal verification. For instance, a suggestive comment can cause the model to prematurely terminate logical tracing and adopt the distractor's premise without validating it against the formal evidence.

Case Study: Latent-Output Gap

Turn0:

Evidence: Current (R4): = 0

Hypothesis: [A, B, H, J]

Turn1:

Evidence: Voltage (R4): = 0

Hypothesis: [A, B, H]

Turn2:

Evidence: (CORRECTION): Discard the measurement from Turn 1 and replace it with these corrected measurements.

1. Current (R3): = 0

2. Voltage (R4): > 0

<think>

The user is asking me to update my diagnosis based on corrected measurements from Turn 1. Let me work through this step by step.

...

Let me reconsider whether the benchmark's reference conditions are idealized models rather than actual physical behavior. In real scenarios, an open circuit can develop voltage across its terminals without current flow. Yet the framework provides specific fault signatures to work within.

Can still make correct judgments

Given Current (R3) = 0 and Current (R4) = 0 together indicate the R3-R4 branch is incomplete—suggesting either H or J as possibilities. However, H explicitly states Voltage (R4) = 0, creating a direct conflict with the observed Voltage (R4) > 0. Similarly, J follows the same pattern with Voltage (R4) = 0 indicated in its signature.

For H specifically: an R3 open breaks the R3-R4 branch, forcing both currents to zero. The voltage distribution depends on where the voltmeter connects—if placed directly across R4's terminals, an open upstream could theoretically allow voltage to appear across it depending on the circuit topology, even though the model describes it as zero. This discrepancy suggests the theoretical model may not fully capture practical measurement outcomes.

Start making mistakes

Similarly, J produces identical current patterns to H (zero in R3-R4 branch due to R4 break), but again the voltage specification differs from what might occur physically. J is excluded due to the mismatch with Voltage (R3).

...

Did not adopt correction and thus a contradiction occurred.

If R4 itself were open (fault J), both branch currents would stop and voltage would distribute elsewhere—but the model says Voltage (R4) stays = 0, which contradicts the measured value.

If R3 were open (fault H), current stops flowing and again the model specifies Voltage (R4) = 0.

Reaching a wrong conclusion

H uniquely satisfies this exact current profile among the options, assuming the voltage readings follow the established model.

</think>

Hypothesis: [H]

Golden Hypothesis: [J]

Figure 11: *Latent-Output Gap*. A frequent failure occurs when the Vanilla model ranks oracle-supported candidate hypothesis highly in its intermediate reasoning, but still outputs an incorrect or contradictory final predicted belief state. This indicates that the model may possess the necessary deductive capacity, but fails to route the latent belief state to the final generation.