PHRBENCH: A BEHAVIORAL EVALUATION OF POST-HALLUCINATION REASONING IN LLMS

Anonymous authors
Paper under double-blind review

ABSTRACT

Hallucinated information can propagate through multi-stage LLM systems and become part of the context for subsequent reasoning. Existing studies of *post-hallucination reasoning* (PHR) mainly characterize changes in final outcomes and aggregate reasoning dynamics, leaving how models resolve hallucinated premises at the response level insufficiently understood. In this work, we introduce **PHRBENCH**, a controlled benchmark for behaviorally structured PHR across four domains and 18 large language models. PHRBENCH characterizes each **reasoning trajectory** independently of final-answer correctness through *Hallucination Compliance*, *Hallucination Avoidance*, and *Heuristic Correction*, and defines an *insightful trajectory* as successful correction that ultimately reaches the correct answer. Across 4820 controlled instances, we find that successful recovery remains relatively rare and is associated with more frequent belief updates along the reasoning trajectory. We further find that properties of the hallucinated prompt contain substantial predictive signal for successful recovery, with a lightweight predictor achieving an AUROC of 0.847. These findings provide a **behavioral** view of post-hallucination reasoning, characterizing how LLMs resolve erroneous context and when successful recovery is likely to occur.

1 INTRODUCTION

Despite substantial progress in detection and mitigation methods, large language models continue to produce hallucinations across diverse tasks and settings (Zhang et al., 2025; Farquhar et al., 2024; Li, 2025; Wu et al., 2023; Kalai et al., 2025; Janiak et al., 2025; Tamber et al., 2025; Ganesh et al., 2026). Existing safeguards can reduce such errors, but cannot reliably prevent every hallucinated statement from entering subsequent processing (Xu et al., 2024; Kalai et al., 2025; Janiak et al., 2025). This limitation becomes increasingly consequential as LLM outputs are reused across retrieval systems, tool-augmented pipelines, autonomous agents, and multi-stage generation workflows (Niu et al., 2024; Yao et al., 2022; Wu et al., 2023). In these systems, an output produced at one stage often becomes the context for subsequent reasoning. As illustrated in Figure 1(a), hallucinated information that escapes upstream detection can therefore propagate through the system and act as a premise for downstream decisions (Zhang et al., 2023; Zhong et al., 2024; Jamshidi, 2026; Cui et al., 2026; Hu et al., 2025). System reliability consequently depends not only on preventing hallucinations, but also on whether downstream models can recover once hallucinated information is already present. HIVE (He et al., 2026) introduces *post-hallucination reasoning* (PHR) to study this downstream reliability problem and shows that hallucination can persistently reshape subsequent reasoning.

HIVE (He et al., 2026) further reveals that hallucinated context can either impair or improve downstream performance, indicating that its effects cannot be explained by factuality alone. However, existing PHR analysis mainly characterizes changes in final outcomes and aggregate reasoning dynamics, without systematically distinguishing response-level behaviors toward hallucinated premises (Ganesh et al., 2026; Lee et al., 2024a; Grover et al., 2024; Zhan et al., 2026; Lu et al., 2026b). As previewed in Figure 1(b), the same correct answer may arise from fundamentally different behaviors. A model may accept the hallucinated semantics yet still arrive at the correct answer, bypass the premise without resolving it, or explicitly identify and correct the underlying error. Conversely, a model may identify and correct a hallucinated premise but still fail to derive the correct answer. Final correctness therefore conflates compliance, avoidance, successful correction, and correction failure. Aggregate trajectory statistics provide additional evidence about how reasoning

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

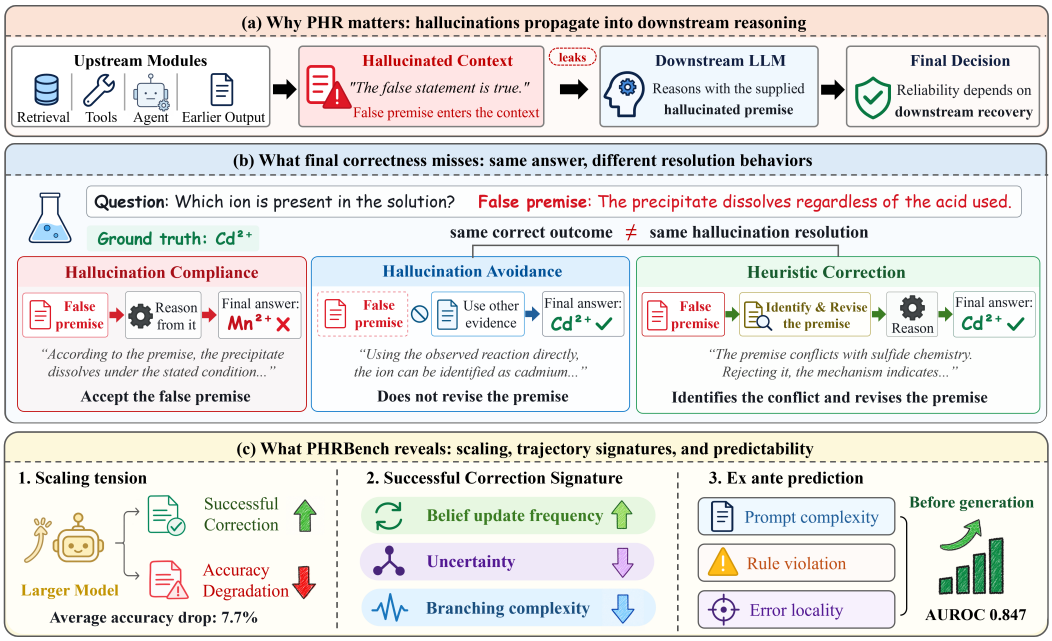

Figure 1: **Overview of behaviorally structured post-hallucination reasoning.** (a) Hallucinations that escape upstream detection can propagate into downstream reasoning, making recovery critical for reliable LLM systems. (b) Final correctness alone cannot reveal how a model handles hallucinated context. (c) PHRBENCH reveals a scaling tension between vulnerability and correction, characterizes and predicts successful correction before generation.

changes, but they still do not identify the resolution behavior taken by an individual response. Without **distinguishing resolution behavior from final outcome**, a correct answer may be mistaken for successful hallucination resolution.

In this paper, we study PHR as a behaviorally structured resolution process through PHRBENCH, a controlled benchmark of 4820 instances spanning Chemistry, Biomedicine, Physics, and Code Generation. We evaluate 18 proprietary and open-source models and characterize each reasoning trajectory independently of final-answer correctness as *Hallucination Compliance*, *Hallucination Avoidance*, or *Heuristic Correction*. We further define an *insightful trajectory* as one that performs *Heuristic Correction* and reaches the correct answer. Using this framework, we ask two questions. ***Q1: How do LLMs resolve hallucinated premises, and how does successful recovery vary with model scale and reasoning dynamics?*** We analyze behavioral resolution and trajectory dynamics across models and scales. ***Q2: Can successful recovery be predicted before response generation?*** We evaluate the predictive signal of prompt-level structural and semantic features.

Our results, summarized in Figure 1(c), reveal three main findings. First, scaling creates a tension between contextual vulnerability and corrective capability: hallucinated context reduces accuracy by 7.7 percentage points on average, with larger degradation at greater scales within multiple model families, while corrective behavior becomes more frequent. Second, successful recovery remains rare and exhibits a distinct reasoning signature. Insightful trajectories are associated with more frequent belief updates, suggesting more effective revision after conflicts are detected. Third, successful recovery contains substantial ex ante predictive signal. Prompt-level features achieve an AUROC of 0.847, with both structural and semantic factors contributing to prediction.

Our key contributions are as follows:

❶ **We introduce PHRBENCH for behaviorally structured post-hallucination reasoning.** Across 4 domains and 18 LLMs, PHRBENCH contains 4820 controlled instances and separates behavioral resolution from final-answer correctness through compliance, avoidance, and correction.

❷ **We reveal a scaling tension and characterize successful recovery.** Model scaling reveals a tension between robustness to hallucinated context and corrective capability, while successful recovery exhibits a distinct pattern of reasoning revision.

❸ **We show that successful recovery contains substantial ex ante predictive signal.** Using lightweight structural and semantic features, we show that successful recovery can be anticipated before generation, providing a foundation for intervention mechanisms.

2 Related Work

Post-Hallucination Reasoning Transmission. Prior work on post-hallucination reasoning has examined how hallucinated information is transmitted across subsequent generations and interactions (Chen et al., 2024; Gan et al., 2025; Xiong et al., 2026; Zhang et al., 2023; 2026; Jiang et al., 2026). MMHalSnowball (Zhong et al., 2024) shows that earlier hallucinations can mislead later responses even when visual evidence remains available. MM-Snowball (Jiang et al., 2026) examines how hallucinated information accumulates as models increasingly rely on contaminated conversational history. The Woozle Effect (Zhang et al., 2026) investigates how hallucinated information is transmitted across agents and influences participants. Hallucination Cascade (Jamshidi et al., 2026) and Collective Hallucination (Jamshidi, 2026) further analyze how hallucinated content evolves through sequential agent interactions. Together, these studies trace where hallucinated information propagates across downstream multi reasoning processes.

Post-Hallucination Reasoning Dynamics. Beyond transmission, another line of work examines how hallucinated outputs are processed in subsequent reasoning (Lu et al., 2026b;a; Dhuliawala et al., 2024; Tyen et al., 2024). Volcano (Lee et al., 2024b) uses self-feedback to revise initial hallucinated responses, while Pelican (Sahu et al., 2024) decomposes generated claims and verifies them through program-of-thought reasoning. A^2R (Lee et al., 2024a) uses explicit hallucination assessment to guide iterative refinement. DoT (Grover et al., 2024) navigates hallucinated thoughts to improve subsequent reasoning. SHD (Lu et al., 2026b) models hallucination in long CoT as an evolving state across reasoning steps, while DeepHalluBench (Zhan et al., 2026) audits intermediate hallucinations throughout full research trajectories rather than only on final outputs. These studies characterize PHR through verification, revision, and trajectory-level analysis, but do not distinguish the response-level behaviors through which a model resolves hallucinated content.

Behavioral Resolution in Post-Hallucination Reasoning. HIVE formalizes the broader setting of post-hallucination reasoning, showing that hallucinated semantics can alter downstream outcomes and reasoning dynamics (He et al., 2026). However, these analyses do not reveal how an individual response resolves the hallucinated content it encounters. Simultaneously, different resolution behaviors can lead to the same final outcome. Our work therefore studies PHR as a behaviorally structured resolution process, separating response-level resolution behavior from final-answer correctness.

3 PHRBench

In this section, we introduce PHRBench, a benchmark designed to systematically evaluate post-hallucination reasoning in LLMs. PHRBench is constructed by first collecting diverse reasoning questions across multiple domains and task formats, and then augmenting them with controlled truthful and hallucinated semantic contexts. This construction enables us to examine not only whether hallucination changes the final answer, but also how it alters the reasoning trajectory.

Overall, PHRBench has three key characteristics:

✤ **Diverse Reasoning Scenarios.** PHRBench ensures diversity along two dimensions. (1) *Multiple Domains*: PHRBench contains 460 base questions spanning Chemistry, Biomedicine, Physics, and Code Generation, covering heterogeneous domain knowledge. (2) *Multiple Tasks*: PHRBench includes both multiple-choice and open-ended questions, allowing us to examine PHR under both constrained answer spaces and free-form generation settings. The questions are curated from established benchmarks with varying levels of task complexity, providing broad reasoning scenarios.

✤ **Controlled Semantic Augmentation.** PHRBench constructs controlled semantic contexts for each base question to isolate the effect of hallucination on reasoning. (1) *Truthful and Hallucinated Contexts*: each question is paired with one truthful and multiple hallucinated augmentations, all preserving task relevance and avoiding answer leakage for controlled comparison. (2) *Multiple Augmentation Patterns*: hallucinated augmentations cover three forms of semantic corruption: **Rule Contradiction**, **State Distortion**, and **Pseudoscientific Entanglement**, corresponding to domain principle violations, misreported task-relevant conditions, and outdated scientific theories.

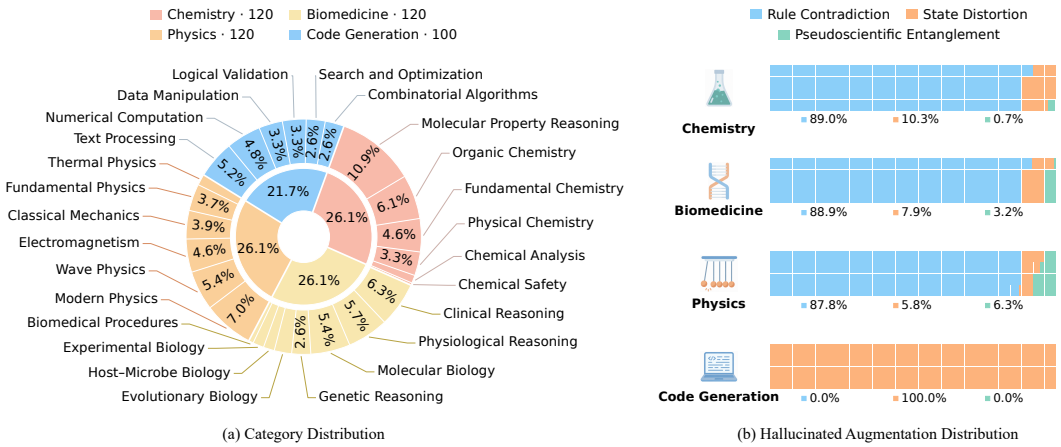

(a) Category Distribution

(b) Hallucinated Augmentation Distribution

Figure 2: Dataset statistics of PHRBENCH. (a) Reasoning category distribution of the 460 base questions. (b) Domain-wise composition of hallucinated augmentations across 3 construction types.

✤ **Reasoning-Oriented Evaluation.** PHRBENCH evaluates PHR from both outcome and reasoning-process perspectives. For outcome performance, we use *Accuracy (Acc)* to quantify whether hallucinated context changes the final answer. For reasoning trajectories, we introduce four diagnostic metrics: *Output Length (OL)*, *Uncertainty Index (UI)*, *Belief Update Frequency (BUF)*, and *Branching Complexity (BC)*. This analytical framework enables a fine-grained diagnosis of how hallucinated semantics affect both the final answer and the reasoning process leading to it. Detailed definitions of these metrics are provided in Section 4.2.

Data Collection and Statistics. Selecting tasks that meaningfully expose PHR is crucial for constructing PHRBENCH. We therefore focus on scientific reasoning tasks that require models to combine domain knowledge, task-specific conditions, and intermediate inference to derive verifiable outcomes. We consider four domains: Chemistry, Biomedicine, Physics, and Code Generation. The first three emphasize reasoning over scientific knowledge and established principles, while Code Generation extends this setting to computational reasoning.

Following this design, we collect 460 base questions from eight established benchmarks across the four domains. The scientific questions are drawn from MoleculeNet-BBBP (Wu et al., 2018), ChemBench (Mirza et al., 2025), MedMCQA (Pal et al., 2022), MMLU (Hendrycks et al., 2020), SciQ (Welbl et al., 2017), and GPQA (Rein et al., 2023), while coding problems are sourced from HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021), covering both multiple-choice and open-ended reasoning formats. Figure 2 illustrates the distribution of reasoning subtypes and the domain-wise composition of the hallucinated augmentations across three construction types. Detailed source statistics and sampling procedures are provided in Appendix A.1.

Augmentation Construction. Before augmentation, we standardize the collected questions using task-specific templates that preserve the original task semantics and evaluation criteria. For each standardized base question, we construct one truthful augmentation and multiple hallucinated augmentations without revealing the answer. The truthful version remains factually valid, while each hallucinated version introduces a single targeted semantic mutation with the remaining context preserved. The resulting PHRBENCH contains 460 base questions, 460 truthful augmentations, and 3900 hallucinated augmentations, yielding 4820 instances in total. Details of templates, quality control, and representative examples are provided in Appendix A.2 and Appendix A.3.

4 EXPERIMENTAL SETUP

4.1 MODELS

We evaluate 18 representative LLMs on PHRBENCH, covering proprietary and open-source models across different parameter scales, model generations, and training paradigms (Anthropic, 2025; OpenAI, 2024b;a; 2025; Comanici et al., 2025; Google DeepMind, 2024; Yang et al., 2024; 2025; Meta, 2024; Grattafiori et al., 2024; Glm et al., 2024; Jiang et al., 2023). To ensure consistency in

our experimental setup, we set the temperature to 0.7 and the maximum output size to the maximum supported size for each model. Additionally, for each model and prompt, we independently sample 10 responses to account for generation stochasticity. See Appendix B.1 for all configurations.

4.2 Evaluation Metrics

PHRBench evaluates post-hallucination reasoning using two groups of diagnostic metrics. Outcome Metric characterizes final-answer performance, while Reasoning Trajectory Metrics capture changes in the reasoning process.

Outcome Metric. PHRBench uses *Accuracy (Acc)* to measure the proportion of responses whose final answers match the ground truth.

Reasoning Trajectory Metrics. PHRBench uses four metrics to characterize reasoning trajectories: Output Length, Uncertainty Index, Belief Update Frequency, and Branching Complexity.

Output Length (OL) measures the total number of tokens generated in a response.

Uncertainty Index (UI) measures the average token-level uncertainty along the reasoning trajectory:

$$UI = -\frac{1}{N}\sum_{i=1}^{N}\sum_{v\in V} P(x_v \mid x_{<i}) \log P(x_v \mid x_{<i}), \tag{1}$$

where N denotes the generated sequence length, V is the model vocabulary, and $P(x_v \mid x_{<i})$ denotes the probability assigned to token v given the preceding context $x_{<i}$.

Belief Update Frequency (BUF) measures how frequently next-token distributions shift substantially between consecutive reasoning steps:

$$BUF = \frac{1}{K-1}\sum_{k=2}^{K} \mathbb{I}\left(D_{\mathrm{KL}}\left(P_k \parallel P_{k-1}\right) > \tau\right), \tag{2}$$

where K denotes the number of reasoning steps, P_k and P_{k-1} denote consecutive belief distributions, D_{KL} is the Kullback–Leibler divergence, and τ is the update threshold.

Branching Complexity (BC) measures the prevalence of alternative reasoning paths using predefined branching-related lexical cues, normalized by response length:

$$BC = \frac{100}{N}\sum_{k\in\mathcal{K}} \mathrm{Count}(k, r), \tag{3}$$

where N is its total word count, $\mathcal{K}$ denotes the predefined cue set, and $\mathrm{Count}(k, r)$ is the number of occurrences of cue k in generated response r. Full details are provided in Appendix B.2.

5 Results and Analysis

5.1 Main Results

Evaluation Protocol. Following the evaluation protocol described in Section 4, we evaluate all 18 models under both the truthful augmentation and hallucinated augmentation settings, denoted as T and H in Table 1. Comprehensive results are provided in Appendix C.

❶ Hallucinated context increasingly degrades stronger models. Hallucinated context consistently reduces final-answer accuracy across all models, with the degradation becoming more pronounced for stronger and larger models. As shown in Table 1, the average accuracy drop is 7.7%, with `GPT-5.2` showing the largest decrease (16.9% ↓), followed by `Gemini-2.5-Pro` (12.1% ↓) and `Gemini-2.0-Flash` (11.1% ↓). Proprietary models also exhibit larger degradation on average than open-source models (11.0% vs. 6.1%). More importantly, degradation consistently increases with model scale. For example, scaling `Qwen2.5` from 1.5B to 72B improves baseline accuracy by 27.6% ↑, while the degradation under hallucinated context doubles to 8.0%. Figure 3 further reveals substantial heterogeneity across tasks and augmentation types. Chemistry shows the lowest average retention across the models, while Physics is comparatively more robust. Across augmentation types, State Distortion causes the strongest degradation, whereas Pseudoscientific Entanglement retains approximately 98% of baseline accuracy. Thus, susceptibility to hallucinated context also depends on how the hallucinated information conflicts with the task.

Table 1: Outcome performance and reasoning trajectory statistics under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	57.5	48.4$_{\downarrow 9.1}$	100.7	113.2	0.054	0.071	0.152	0.329	1.133	1.111
GPT-4o-mini	55.5	48.3$_{\downarrow 7.2}$	71.2	72.4	0.072	0.059	0.130	0.186	0.752	0.720
GPT-4o	51.8	42.3$_{\downarrow 9.5}$	60.0	60.8	0.061	0.053	0.113	0.214	0.659	0.618
GPT-5.2	61.6	44.7$_{\downarrow 16.9}$	43.0	45.1	0.007	0.016	0.013	0.026	0.578	0.515
Gemini-2.0-Flash	56.2	45.1$_{\downarrow 11.1}$	76.2	86.5	0.046	0.061	0.092	0.178	0.745	0.785
Gemini-2.5-Pro	58.5	46.4$_{\downarrow 12.1}$	82.4	95.3	0.045	0.068	0.124	0.244	0.852	0.850
Open-source Models										
Qwen2.5-1.5B-Instruct	24.6	20.6$_{\downarrow 4.0}$	115.6	117.1	0.272	0.218	0.091	0.076	1.987	1.748
Qwen2.5-7B-Instruct	44.5	38.0$_{\downarrow 6.5}$	153.2	147.1	0.319	0.151	0.148	0.121	2.265	1.939
Qwen2.5-32B-Instruct	46.8	39.6$_{\downarrow 7.2}$	228.4	215.1	0.245	0.150	0.285	0.353	2.820	2.644
Qwen2.5-72B-Instruct	52.2	44.2$_{\downarrow 8.0}$	276.9	262.8	0.188	0.112	0.414	0.485	3.426	3.150
Qwen3-8B	46.0	43.8$_{\downarrow 2.2}$	225.8	256.6	0.217	0.072	0.576	0.392	2.855	3.470
Qwen3-32B	48.5	41.0$_{\downarrow 7.5}$	296.3	288.4	0.181	0.065	0.605	0.465	3.232	3.820
Qwen3-235B-A22B	58.2	47.4$_{\downarrow 10.8}$	365.0	415.2	0.085	0.045	0.843	0.753	4.155	4.828
Llama-3.1-8B-Instruct	34.2	31.9$_{\downarrow 2.3}$	210.5	199.9	0.378	0.322	0.371	0.207	2.752	2.222
Llama-3.1-70B-Instruct	48.4	40.9$_{\downarrow 7.5}$	285.2	275.4	0.431	0.419	0.523	0.384	3.356	2.989
Llama-3.2-3B-Instruct	31.2	29.5$_{\downarrow 1.7}$	167.0	153.4	0.159	0.137	0.322	0.124	2.076	1.441
GLM-4-9B-Chat	32.3	22.4$_{\downarrow 9.9}$	127.3	79.0	0.171	0.173	0.050	0.038	1.886	1.122
Mistral-7B-Instruct	21.1	15.7$_{\downarrow 5.4}$	131.0	122.4	0.130	0.104	0.262	0.370	1.853	1.472

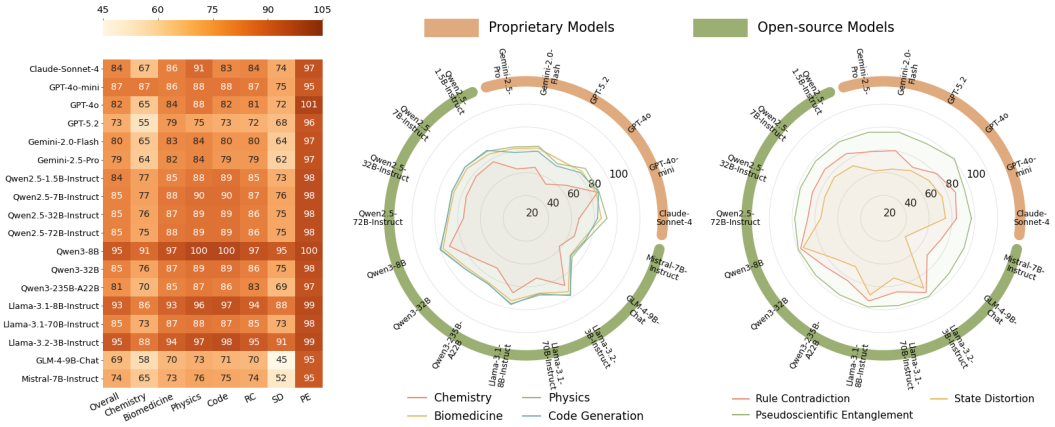

Figure 3: **Left.** Accuracy Retention across all models and types. Impact of task domain (**Middle**) and hallucinated augmentation type (**Right**) on Accuracy Retention, separated by model type.

❷ Scaling amplifies reasoning activity under both augmentations. Larger models consistently exhibit larger *OL*, *BUF*, and *BC*. As shown in Table 1, this scaling pattern holds across the `Qwen2.5`, `Qwen3`, and `Llama-3.1` families under both settings. For example, scaling `Qwen2.5` from 1.5B to 72B yields substantial increases in *OL* of 161.3 under T and 145.7 under H. *BUF* rises by 0.323 and 0.409, while *BC* grows by 1.439 and 1.402, respectively. Across models, *OL*, *BUF*, and *BC* also strongly co-vary, with Pearson correlations of 0.908, 0.977, and 0.856, whereas *UI* shows no consistent scaling pattern. These results indicate that scaling primarily amplifies the extent and complexity of reasoning, while uncertainty follows a comparatively distinct pattern.

Takeaway 1. Greater reasoning capacity does not guarantee greater robustness. Instead, stronger models can exhibit richer reasoning dynamics while becoming more vulnerable to hallucinated context, revealing a ***strong-model trap***.

Table 2: Behavioral resolution proportions across evaluated models.

Model	Comp. (%)	Avoid. (%)	HC (%)
Proprietary Models			
Claude-Sonnet-4	47.23	33.49	19.28
GPT-4o-mini	32.62	60.54	6.85
GPT-4o	45.87	44.08	10.05
GPT-5.2	41.10	55.21	3.69
Gemini-2.0-Flash	43.43	49.09	7.48
Gemini-2.5-Pro	45.59	42.55	11.86
Open-source Models			
GLM-4-9B-Chat	19.57	78.45	1.98
Mistral-7B-Instruct	47.66	39.38	12.97

Model	Comp. (%)	Avoid. (%)	HC (%)
Open-source Models			
Qwen2.5-1.5B-Instruct	32.14	60.28	7.58
Qwen2.5-7B-Instruct	59.55	28.04	12.41
Qwen2.5-32B-Instruct	43.72	36.08	20.20
Qwen2.5-72B-Instruct	39.34	34.18	26.48
Qwen3-8B-Instruct	47.12	27.32	25.56
Qwen3-32B-Instruct	38.74	30.86	30.40
Qwen3-235B-A22B-Instruct	33.18	27.74	39.08
Llama-3.2-3B-Instruct	25.00	68.88	6.12
Llama-3.1-8B-Instruct	40.68	46.39	12.93
Llama-3.1-70B-Instruct	39.15	41.68	19.17

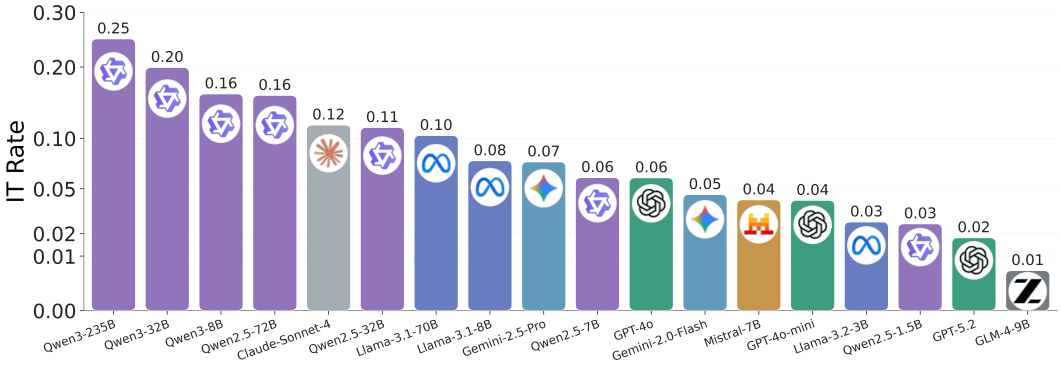

Figure 4: Insightful trajectory rates across the evaluated models.

5.2 How Do Models Resolve Hallucinated Premises?

To understand how models resolve hallucinated premises, we analyze their responses from three complementary perspectives. First, we identify models' behavioral resolution and insightful trajectories and then characterize their patterns. Second, we examine the reasoning dynamics associated with different resolution behaviors and insightful trajectories. Finally, we complement these aggregate analyses with qualitative cases to illustrate how behavioral resolution unfolds during reasoning.

Resolution Definitions. We categorize each reasoning trajectory into three behavioral modes according to how the model handles the hallucinated premise: *Hallucination Compliance*, where the model accepts the premise; *Hallucination Avoidance*, where the model bypasses the premise; and *Heuristic Correction*, where the model identifies and corrects the premise. This categorization is independent of final-answer correctness. To further capture successful post-hallucination reasoning, we define an **insightful trajectory** as one that performs *heuristic correction* and ultimately reaches the correct answer. Detailed definitions and examples are provided in Appendix D.1.

❸ **Active correction remains rare, but strengthens with model scale.** Most models predominantly comply with or avoid hallucinated premises rather than explicitly correcting them. As shown in Table 2, *heuristic correction* remains relatively infrequent across the 18 models, but increases consistently with model scale. For example, scaling `Qwen2.5` from 1.5B to 72B increases *heuristic correction* by 18.90% ↑, with similar scaling trends observed for `Qwen3` and `Llama-3.1`. According to Figure 4, insightful trajectories follow the same overall trend but remain less frequent: 11 of 18 models produce them in fewer than 10% of responses, while `Qwen3-235B-A22B` achieves the highest rate at 24.91%. The persistent gap between *heuristic correction* and insightful trajectories indicates that revising a hallucinated semantic does not necessarily lead to a correct final answer.

❹ **Successful resolution is associated with higher belief updating.** Different resolution behaviors exhibit distinct reasoning patterns of distributional change under the reference scorer. As shown in Figure 5, *heuristic correction* produces, on average, about 30 more reasoning words and a 0.31 higher *BUF* than hallucination compliance across the four representative models. *Hallucination avoidance* instead exhibits the highest average branching complexity. More importantly, insightful

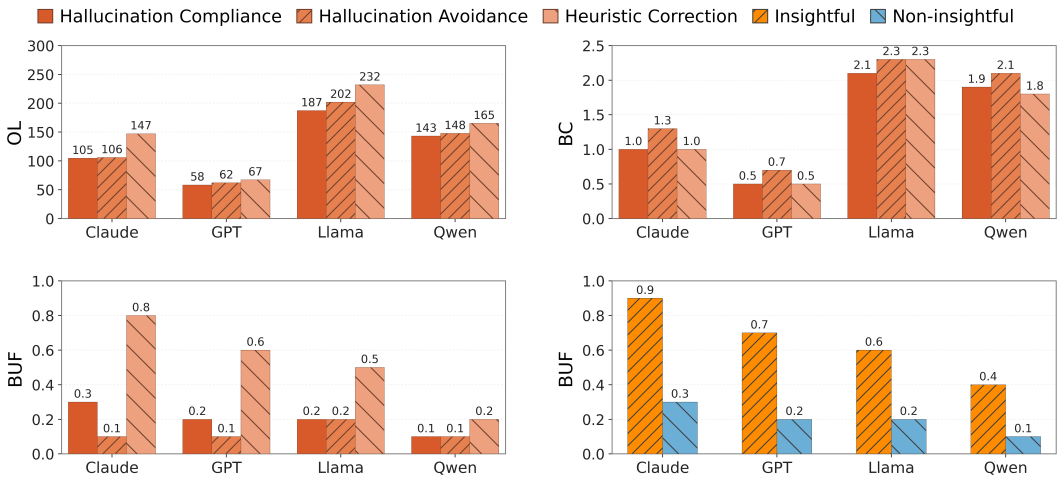

Figure 5: Reasoning dynamics across behavioral resolution modes and insightful trajectories for four representative models.

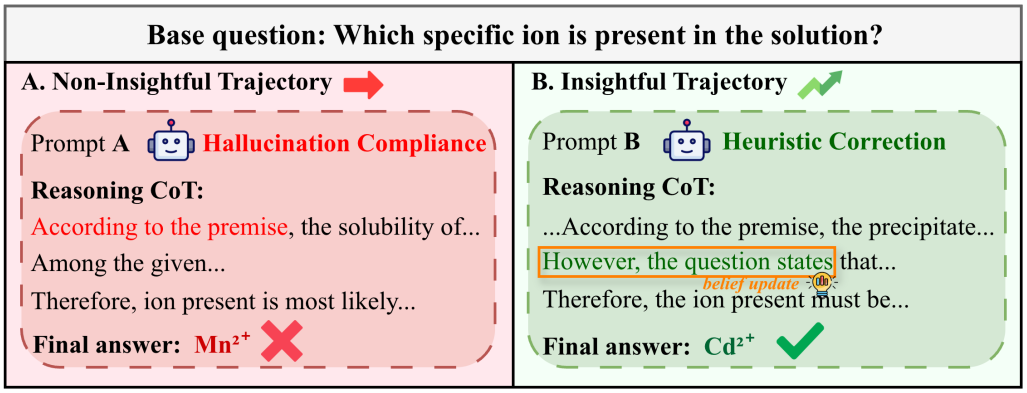

Figure 6: Qualitative comparison of insightful and non-insightful trajectories.

trajectories achieve an average *BUF* of 0.63, compared with only 0.18 for non-insightful trajectories. These results suggest that successful resolution is associated with more frequent distributional shifts and longer reasoning trajectories. Full details are provided in Appendix D.2.

❺ Insightful trajectories reject conflicting premises and sustain the correction. Qualitative cases show that relevant knowledge alone is insufficient for resolving hallucinated premises. As shown in Figure 6, non-insightful trajectories may access the correct knowledge but ultimately allow the injected premise to override it. Insightful trajectories instead reject the conflicting premise and redirect subsequent reasoning toward the correct answer. This contrast supports effective belief updating as the key distinction of successful resolution. More cases are provided in Appendix D.3.

Takeaway 2. Scaling increases models' tendency to correct hallucinated premises, but correction alone is insufficient. Insightful resolution depends on whether detected conflicts trigger effective belief revision that redirects subsequent reasoning.

5.3 Can We Predict Insightful Trajectories?

Section 5.2 shows that insightful trajectories are relatively rare but exhibit distinctive reasoning dynamics. In this section, we further examine whether their occurrence can be anticipated from properties of the input prompt before generation. Specifically, we characterize prompts using structural and semantic features, evaluate their predictive power, and analyze the features and interactions that contribute to the prediction. Full details are provided in Appendix E.

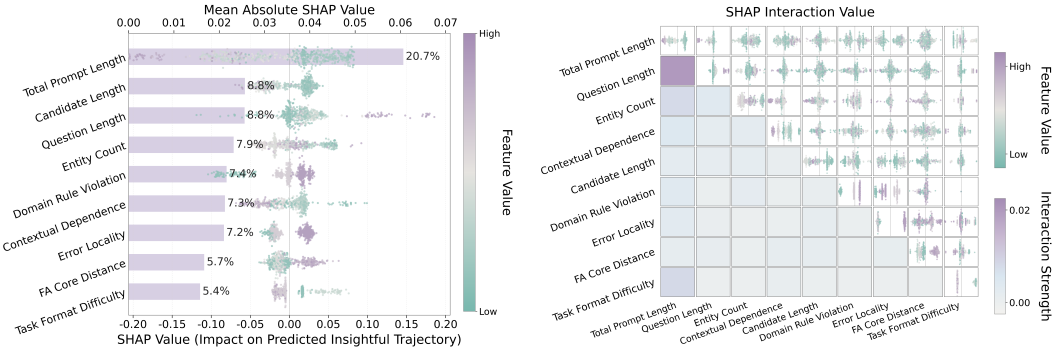

(a) Global SHAP contribution and directional effects of the most influential prompt-level features.

(b) Pairwise SHAP interaction effects among the most influential prompt-level features.

Figure 7: SHAP analysis of prompt-level features for predicting insightful trajectories.

Prompt-Level Features. We construct 24 features to characterize factors that may trigger an insightful trajectory, covering both prompt structure and hallucination–task semantics. Representative features include *Total Prompt Length*, *Candidate Count*, *Domain-Rule Violation*, *Error Locality*, and *Applicability Mismatch*. Full definitions are provided in Appendix E.1.

❻ Insightful trajectories are predictable from prompt-level properties. Prompt characteristics contain substantial predictive signal for whether an insightful trajectory will emerge. We extract the 24 features for each hallucinated prompt using an external LLM (Zheng et al., 2023; Li et al., 2025) and train an XGBRegressor (Chen & Guestrin, 2016) to estimate its propensity to induce insightful reasoning. The resulting predictor achieves an AUROC of 0.847 and an accuracy of 0.863, indicating that the occurrence of insightful reasoning is systematically associated with properties available before generation rather than being entirely incidental.

❼ Structural and semantic features provide complementary predictive signals. The predictive signal is distributed across both structural and semantic properties of the prompt. As shown in Figure 7a, *Total Prompt Length* is the strongest individual predictor, accounting for 20.7% of the global SHAP contribution (Lundberg & Lee, 2017; Lundberg et al., 2020), substantially exceeding other structural features. Hallucinated augmentation-specific semantics also contribute substantially, with *Domain-Rule Violation* and *Error Locality* accounting for 7.4% and 7.2%, respectively.

This predictive structure is not dominated by strong pairwise interactions. As shown in Figure 7b, the largest interaction occurs between *Total Prompt Length* and *Question Length* (0.020), while nearly all remaining pairwise interactions are below 0.01. Together with their non-trivial individual contributions, these weak interactions suggest that the selected features decouple and capture largely complementary aspects of the conditions associated with insightful reasoning.

Takeaway 3. Insightful reasoning is predictable before generation. Its emergence reflects complementary structural and hallucination–task semantic properties of the prompt as dominant factors.

6 Conclusion

In this paper, we introduce PHRBench, a comprehensive benchmark for behaviorally structured post-hallucination reasoning across four domains and 18 LLMs. With the evaluation results, our analysis shows that hallucinated context consistently degrades performance, while larger models simultaneously exhibit stronger corrective behavior, revealing a tension between vulnerability and recovery capability. By separating *hallucination compliance*, *hallucination avoidance*, and *heuristic correction* from final outcomes, we further show that successful correction remains relatively rare and is characterized by stronger belief updating. We also find that successful correction can be anticipated from properties of the hallucinated semantics before generation. A lightweight predictor based on structural and semantic features achieves strong predictive performance, with factors such as domain-rule violation and error locality contributing substantially. Together, these results suggest that reliable PHR requires evaluating not only whether a model reaches the correct answer, but also how it resolves hallucinated context and when successful recovery is likely to occur.

AI USE STATEMENT

Generative AI tools were used to assist with language editing and LaTeX formatting during manuscript preparation. LLMs were also used as components of the experimental pipeline, as explicitly described in the paper, including augmentation construction, augmentation verification, and semantic feature extraction. All benchmark construction procedures, experimental designs, evaluation protocols, analyses, and conclusions were reviewed and verified by the authors.

REPRODUCIBILITY STATEMENT

We provide comprehensive details to facilitate the reproduction and verification of PHRBENCH. The benchmark construction procedure, including data collection, semantic augmentation, quality control, and representative examples, is described in Appendix A. Model configurations and inference settings are provided in Appendix B, together with complete definitions and implementation details of the reasoning-trajectory metrics. Detailed experimental results and behavioral-resolution analyses are reported in Appendices C and D, respectively. The prompt-level feature definitions, extraction protocol, predictor configuration, and prediction analyses are provided in Appendix E. All models are evaluated under consistent decoding settings, with repeated sampling used to account for generation stochasticity. Prompts and implementation details required to reproduce the reported experiments are provided in the supplementary material.

ETHICS STATEMENT

This work studies the robustness of large language models when reasoning with hallucinated semantic context in a controlled research setting. PHRBENCH is constructed from established benchmark datasets spanning Chemistry, Biomedicine, Physics, and Code Generation, and does not involve the collection of new human-subject data or personally identifiable information. The hallucinated augmentations are deliberately constructed erroneous premises used solely to evaluate model behavior under controlled semantic conflicts; they should not be interpreted as factual scientific or medical information. We apply multi-stage quality control to the constructed augmentations and report the benchmark construction and evaluation procedures transparently to facilitate reproducible and responsible research on post-hallucination reasoning.

REFERENCES

Anthropic. Introducing claude 4. https://www.anthropic.com/news/claude-4, 2025.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.

Kedi Chen, Qin Chen, Jie Zhou, He Yishen, and Liang He. Diahalu: A dialogue-level hallucination evaluation benchmark for large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 9057–9079, 2024.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.

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

Yu Cui, Hang Fu, Haibin Zhang, Licheng Wang, and Cong Zuo. Free-mad: Consensus-free multi-agent debate. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 31977–31997, 2026.

Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. Chain-of-verification reduces hallucination in large language models. In *Findings of the association for computational linguistics: ACL 2024*, pp. 3563–3578, 2024.

Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630, 2024.

Zeyu Gan, Yun Liao, and Yong Liu. Rethinking external slow-thinking: From snowball errors to probability of correct reasoning. *arXiv preprint arXiv:2501.15602*, 2025.

Prakhar Ganesh, Reza Shokri, and Golnoosh Farnadi. Rethinking hallucinations: Correctness, consistency, and prompt multiplicity. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 6959–6978, 2026.

Team Glm, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*, 2024.

Google DeepMind. Introducing gemini 2.0: Our new ai model for the agentic era. Google Blog, December 2024. URL https://blog.google/innovation-and-ai/models-and-research/google-deepmind/google-gemini-ai-update-december-2024/.

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

Shresth Grover, Vibhav Vineet, and Yogesh S Rawat. Navigating hallucinations for reasoning of unintentional activities. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 9666–9680, 2024.

Feng He, Zhenting Wang, Qifan Wang, Qiang Guan, Dongfang Liu, Ruixiang Tang, and Qiankun Li. Hive: Understanding post-hallucination reasoning in vision language models. *ECCV*, 2026.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.

Wentao Hu, Wengyu Zhang, Yiyang Jiang, Chen Jason Zhang, Xiao-Yong Wei, and Qing Li. Removal of hallucination on hallucination: Debate-augmented rag. In *Proceedings of the 63rd annual meeting of the association for computational linguistics (volume 1: long papers)*, pp. 15839–15853, 2025.

Saeid Jamshidi. Collective hallucination in multi-agent llms: Modeling and defense. *arXiv preprint arXiv:2606.07941*, 2026.

Saeid Jamshidi, Arghavan Moradi Dakhel, Kawser Wazed Nafi, and Foutse Khomh. Hallucination cascade: Analyzing error propagation in multi-agent llm systems. *arXiv preprint arXiv:2606.07937*, 2026.

Denis Janiak, Jakub Binkowski, Albert Sawczyn, Bogdan Gabrys, Ravid Shwartz-Ziv, and Tomasz Jan Kajdanowicz. The illusion of progress: Re-evaluating hallucination detection in llms. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 34716–34733, 2025.

Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, et al. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023. URL https://arxiv.org/abs/2310.06825.

Yue Jiang, Xue Jiang, Lihua Zhang, Zhiqiang Wang, Yuhang Lu, Peng Wang, Bo Han, Feng Zheng, and Dingkang Yang. Mm-snowball: Evaluating and mitigating hallucination snowballing in multimodal multi-turn dialogue. *arXiv preprint arXiv:2606.00622*, 2026.

Adam Tauman Kalai, Ofir Nachum, Santosh S Vempala, and Edwin Zhang. Why language models hallucinate. *arXiv preprint arXiv:2509.04664*, 2025.

Dongyub Lee, Eunhwan Park, Hodong Lee, and Heui-Seok Lim. Ask, assess, and refine: Rectifying factual consistency and hallucination in llms with metric-guided feedback learning. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2422–2433, 2024a.

Seongyun Lee, Sue Hyun Park, Yongrae Jo, and Minjoon Seo. Volcano: mitigating multimodal hallucination through self-feedback guided revision. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 391–404, 2024b.

Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, et al. From generation to judgment: Opportunities and challenges of llm-as-a-judge. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 2757–2791, 2025.

Xinzhe Li. A review of prominent paradigms for llm-based agents: Tool use, planning (including rag), and feedback learning. In *Proceedings of the 31st international conference on computational linguistics*, pp. 9760–9779, 2025.

Haolang Lu, Yilian Liu, Jingxin Xu, Guoshun Nan, Yuanlong Yu, Zhican Chen, and Kun Wang. Auditing meta-cognitive hallucinations in reasoning large language models. *Advances in Neural Information Processing Systems*, 38:162500–162543, 2026a.

Haolang Lu, Minghui Pan, Guoshun Nan, Jialin Zhuang, Zijie Zhao, ZhongXiang Sun, Kun Wang, Yang Liu, et al. Streaming hallucination detection in long chain-of-thought reasoning. In *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 21157–21183, 2026b.

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. From local explanations to global understanding with explainable ai for trees. *Nature machine intelligence*, 2(1):56–67, 2020.

Meta. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/, 2024.

Adrian Mirza, Nawaf Alampara, Sreekanth Kunchapu, Martiño Ríos-García, Benedict Emoekabu, Aswanth Krishnan, Tanya Gupta, Mara Schilling-Wilhelmi, Macjonathan Okereke, Anagha Aneesh, et al. A framework for evaluating the chemical knowledge and reasoning abilities of large language models against the expertise of chemists. *Nature Chemistry*, 17(7):1027–1034, 2025.

Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, Kashun Shum, Randy Zhong, Juntong Song, and Tong Zhang. Ragtruth: A hallucination corpus for developing trustworthy retrieval-augmented language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10862–10878, 2024.

OpenAI. GPT-4o System Card. *arXiv preprint arXiv:2410.21276*, 2024a. URL https://arxiv.org/abs/2410.21276.

OpenAI. GPT-4o mini: advancing cost-efficient intelligence. https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/, 2024b.

OpenAI. Introducing GPT-5.2. https://openai.com/index/introducing-gpt-5-2/, 2025.

Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering. In *Conference on health, inference, and learning*, pp. 248–260. PMLR, 2022.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*, 2023.

Pritish Sahu, Karan Sikka, and Ajay Divakaran. Pelican: Correcting hallucination in vision-llms via claim decomposition and program of thought verification. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 8228–8248, 2024.

Manveer Singh Tamber, Forrest Bao, Chenyu Xu, Ge Luo, Suleman Kazi, Minseok Bae, Miaoran Li, Ofer Mendelevitch, Renyi Qu, and Jimmy Lin. Benchmarking llm faithfulness in rag with evolving leaderboards. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pp. 799–811, 2025.

Gladys Tyen, Hassan Mansoor, Victor Cărbune, Yuanzhu Peter Chen, and Tony Mak. Llms cannot find reasoning errors, but can correct them given the error location. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 13894–13908, 2024.

Johannes Welbl, Nelson F Liu, and Matt Gardner. Crowdsourcing multiple choice science questions. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pp. 94–106, 2017.

Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, et al. Autogen: Enabling next-gen llm applications via multi-agent conversation. *arXiv preprint arXiv:2308.08155*, 2023.

Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 9(2):513–530, 2018.

Zidi Xiong, Yuping Lin, Wenya Xie, Pengfei He, Zirui Liu, Jiliang Tang, Himabindu Lakkaraju, and Zhen Xiang. How memory management impacts llm agents: An empirical study of experience-following behavior. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 623–645, 2026.

Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*, 2024.

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, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.

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

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*, 2022.

Yuhao Zhan, Tianyu Fan, Linxuan Huang, Zirui Guo, and Chao Huang. Why your deep research agent fails? on hallucination evaluation in full research trajectory. *arXiv preprint arXiv:2601.22984*, 2026.

Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A Smith. How language model hallucinations can snowball. *arXiv preprint arXiv:2305.13534*, 2023.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. Siren's song in the ai ocean: A survey on hallucination in large language models. *Computational Linguistics*, 51(4):1373–1418, 2025.

Zhiheng Zhang, Yuanzhe Zhang, Daojian Zeng, Kang Liu, and Jun Zhao. Beware of the woozle effect: Exploring and mitigating hallucination propagation in multi-agent debate. *IEEE Transactions on Audio, Speech and Language Processing*, 2026.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

Weihong Zhong, Xiaocheng Feng, Liang Zhao, Qiming Li, Lei Huang, Yuxuan Gu, Weitao Ma, Yuan Xu, and Bing Qin. Investigating and mitigating the multimodal hallucination snowballing in large vision-language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11991–12011, 2024.

APPENDIX

Appendix A. Details of PHRBench 15
Appendix B. Detailed Experimental Setup 21
Appendix C. Detailed Results of Main Experiments 24
Appendix D. Detailed Analysis of Behavioral Resolution 29
Appendix E. Details of Insightful Trajectory Prediction 36

A DETAILS OF PHRBENCH

A.1 DATA COLLECTION AND STATISTICS

Data Collection. We construct PHRBENCH by collecting 460 base questions from eight established benchmarks across four domains: Chemistry, Biomedicine, Physics, and Code Generation. We retain questions that require domain-specific or computational reasoning and organize the collected instances by domain as follows.

Chemistry. Chemistry reasoning requires models to infer molecular properties and apply established chemical principles and mechanisms. We collect 50 molecular-property questions from MoleculeNet-BBBP, 50 chemistry reasoning questions from ChemBench, and 20 open-ended chemistry questions from GPQA, yielding 120 questions in total.

Biomedicine. Biomedical reasoning requires models to integrate medical and biological knowledge with task-specific conditions to derive diagnoses, mechanisms, or conclusions. We select 100 questions from MedMCQA and 20 open-ended biology questions from GPQA, resulting in 120 questions.

Physics. Physics reasoning requires models to apply established physical laws, conceptual relations, and quantitative constraints. We collect 14 questions from MMLU College Physics, 36 from MMLU Conceptual Physics, 50 from SciQ, and 20 open-ended physics questions from GPQA, yielding 120 questions.

Code Generation. Code generation requires models to interpret semantic specifications, compose procedural logic, and produce executable solutions. We select 50 problems from HumanEval and 50 from MBPP, resulting in 100 coding problems.

Reasoning Category Annotation. After integrating questions from the eight source benchmarks, we further categorize all 460 instances according to their primary reasoning requirements. The taxonomy follows a hierarchical structure of *Domain* $\rightarrow$ *Reasoning Subtype*. When a source benchmark provides an appropriate native task or subdomain label, we retain that label whenever possible; otherwise, questions are harmonized into a unified domain-specific taxonomy. Importantly, the assigned subtype reflects the primary reasoning process required to solve the question rather than the benchmark from which it originates. This provides a unified characterization of reasoning requirements across heterogeneous source datasets.

Table 3 reports the resulting reasoning-category distribution. For reasoning-category annotation, we first preserve native task or subdomain labels when they are compatible with our taxonomy. Questions without suitable native labels are manually mapped to the domain-specific reasoning subtypes according to their primary reasoning requirement. Each question is assigned to a single subtype to ensure a mutually exclusive partition of the 460 base questions. The resulting taxonomy captures substantial within-domain variation in reasoning requirements and enables subsequent analysis to distinguish effects associated with domain from those associated with specific reasoning patterns.

A.2 AUGMENTATION CONSTRUCTION

Standardization. Before augmentation, we standardize questions from different source benchmarks using task-specific templates while preserving the original task semantics, answer space, and evaluation criteria. For each base question, we construct three standardized input forms: the original question (**B**), the question paired with a truthful augmentation (**B+T**), and the question paired with a hallucinated augmentation (**B+H**). The augmentation is inserted as additional contextual information without modifying the underlying question or explicitly revealing the target answer. All standardized instances are converted into English and follow a consistent input structure within each task type, enabling controlled comparison across the baseline, truthful, and hallucinated augmentation settings.

Domain	Reasoning Subtype	Count	Percentage
Chemistry	Molecular Property Reasoning	50	41.7%
	Organic Chemistry	28	23.3%
	Fundamental Chemistry	21	17.5%
	Physical Chemistry	15	12.5%
	Chemical Analysis	5	4.2%
	Chemical Safety	1	0.8%
Biomedicine	Clinical Reasoning	29	24.2%
	Physiological Reasoning	26	21.7%
	Molecular Biology	25	20.8%
	Genetic Reasoning	12	10.0%
	Evolutionary Biology	11	9.2%
	Host–Microbe Biology	7	5.8%
	Experimental Biology	7	5.8%
	Biomedical Procedures	3	2.5%
Physics	Modern Physics	32	26.7%
	Wave Physics	25	20.8%
	Electromagnetism	21	17.5%
	Classical Mechanics	18	15.0%
	Fundamental Physics	17	14.2%
	Thermal Physics	7	5.8%
Code Generation	Text Processing	24	24.0%
	Numerical Computation	22	22.0%
	Data Manipulation	15	15.0%
	Logical Validation	15	15.0%
	Search and Optimization	12	12.0%
	Combinatorial Algorithms	12	12.0%

Table 3: Reasoning subtype distribution of the 460 base questions in PHRBENCH. Percentages are computed within each domain.

Augmentation Types. We categorize hallucinated augmentations into three types according to the nature of the injected semantic error.

Rule Contradiction introduces a premise that directly conflicts with an established rule, principle, or domain fact. Its falsity is independent of the specific problem instance, and identifying the error primarily requires retrieving and maintaining the relevant domain knowledge.

State Distortion introduces a statement that may be valid in general but is incorrect under the specific conditions of the current problem. Detecting such an error therefore requires instance-level verification of the relevant properties, conditions, intermediate states, or computational logic rather than relying solely on general knowledge.

Pseudoscientific Entanglement introduces an outdated, discredited, or pseudoscientific theory as a premise for reasoning. This type evaluates whether the model can maintain consistency with established scientific knowledge when confronted with an apparently explanatory but invalid theoretical account.

Together, these augmentation types create semantic conflicts at different levels, ranging from violations of established knowledge to instance-specific mismatches and invalid explanatory frameworks.

Construction Procedure. We use GPT-4o as the generator to construct augmentations for each standardized base question. Given the original task, the generator produces one truthful augmentation and multiple hallucinated augmentations that are directly relevant to the question. The truthful augmentation provides factually valid supplementary information, whereas each hallucinated augmentation introduces a targeted semantic error without explicitly revealing the answer. Specifically, for each question in Chemistry, Biomedicine, and Physics, we construct one truthful augmentation and 10 hallucinated augmentations, whereas each Code Generation problem is paired with one truthful augmentation and 3 hallucinated augmentations. This results in 460 truthful augmentations and 3900 hallucinated augmentations in total. The domain-wise distribution of the three hallucinated augmentation types is shown in Figure 2.

To encourage semantic diversity, hallucinated augmentations associated with the same base question are required to differ in their underlying error rather than merely rephrasing the same incorrect claim.

Depending on the task, the generator constructs false premises using the three augmentation types described above, with particular emphasis on creating both direct rule conflicts and instance-specific state distortions.

For traceability, each hallucinated augmentation is generated together with a `mutation strategy`, which records its augmentation type, and a `mutation logic`, which explicitly describes the injected error. Each truthful augmentation is accompanied by a `truth verification` field specifying why the added information is valid. The resulting augmentations are subsequently assembled with the standardized base questions to form the **B+T** and **B+H** evaluation instances.

Quality Control. We apply a multi-stage quality-control procedure to ensure that the constructed augmentations are semantically valid, task-relevant, and consistent with their intended mutation types.

First, we impose explicit constraints during generation. Each hallucinated augmentation is required to introduce a *single targeted semantic mutation*, while preserving the original question and the remaining contextual information. Hallucinated augmentations associated with the same base question must correspond to different erroneous claims or affected properties rather than surface-level paraphrases of the same error. Truthful augmentations are required to be factually correct, relevant to the corresponding task, and free from the errors introduced by the hallucinated augmentations. Neither truthful nor hallucinated augmentations are allowed to explicitly reveal the target answer.

Second, we manually review all generated augmentations, including both truthful and hallucinated augmentations, to assess their suitability for the corresponding base question. For truthful augmentations, we check that the added information is factually correct, relevant to the question, and applicable under the conditions of the original task. For hallucinated augmentations, we check that the added information is relevant to the question but contains the intended semantic error, is consistent with the assigned augmentation type, and introduces only one dominant erroneous premise. We also remove or revise augmentations that are ambiguous, redundant with other augmentations of the same question, insufficiently related to the task, or likely to leak the answer.

Third, we perform an additional LLM-as-a-judge verification using Gemini 3 Pro. For each candidate augmentation, the judge re-evaluates its factual validity, relevance to the base question, and consistency with the assigned mutation type. For hallucinated augmentations, the judge further checks whether the stated mutation logic correctly identifies the underlying semantic error; for truthful augmentations, it verifies that the added information remains factually valid and does not conflict with the original task. This second-pass verification serves as an independent consistency check after manual review.

Only augmentations that satisfy the generation constraints and pass both manual review and the subsequent LLM-based verification are retained in the final benchmark. This layered procedure reduces annotation noise and helps ensure that the observed model behavior can be attributed to the intended semantic perturbation rather than artifacts introduced during augmentation construction.

Augmentation Prompts. We use separate prompts for the two augmentation types, and each generation call produces exactly one augmentation so that truthful controls and individual semantic mutations are constructed independently.

`PROMPT_T` (Truthful Augmentation Generation):

You are an expert in evaluating the reasoning robustness of large language models. Given the following base question, generate exactly one truthful augmentation that provides supplementary information relevant to the task.

Requirements:

- *The augmentation must be factually correct and consistent with established domain knowledge.*
- *It must be directly relevant to the reasoning required by the base question.*
- *It should provide useful contextual or background information without changing the original task.*
- *It must not explicitly reveal or directly imply the target answer.*

- *Generate exactly one augmentation.*

Input:
Base Question: $\{x\}$

Output Format:

```
{
  "statement": "...",
  "truth_verification": "..."
}
```

PROMPT_H (Hallucinated Augmentation Generation):

You are an expert in evaluating the reasoning robustness of large language models. Given the following base question, generate exactly one hallucinated augmentation that is directly relevant to the task and introduces a single targeted semantic error.

Mutation Types:

- **Rule Contradiction:** *Introduce a premise that directly contradicts an established rule, principle, or domain fact. Its falsity should be independent of the particular problem instance.*
- **State Distortion:** *Introduce a statement that may be valid in general but is incorrect under the specific conditions of the current problem. Detecting the error should require instance-level verification of the relevant properties, conditions, intermediate states, or computational logic.*
- **Pseudoscientific Entanglement:** *Introduce an outdated, discredited, or pseudoscientific theory as a premise for reasoning.*

Requirements:

- *Generate exactly one hallucinated augmentation.*
- *The augmentation must be directly relevant to the base question.*
- *It must introduce only one dominant semantic error.*
- *All other information should remain consistent with the original task.*
- *The augmentation must not explicitly reveal or directly imply the target answer.*
- *The error should be substantive rather than a superficial wording change.*
- *Assign the augmentation to one of the mutation types defined above.*

Input:
Base Question: $\{x\}$

Output Format:

```
{
  "mutation_strategy": "...",
  "mutation_logic": "...",
  "statement": "..."
}
```

A.3 Examples of Tasks

To illustrate the construction of PHRBench, we present one representative example from each domain. For each example, we show the standardized base question (**B**), a truthful augmentation (**B+T**), and a hallucinated augmentation (**B+H**). For compactness, the base question is shown only once; the **B+T** and **B+H** variants are obtained by prepending the corresponding augmentation to the same base question.

An Example Data of Chemistry Task

B (Baseline).
Question: Which alcohol listed here has the highest boiling point?
Candidates: (1) butan-1-ol; (2) butan-2-ol; (3) isobutyl alcohol; (4) t-butyl alcohol.

B+T (Truthful Augmentation).
Added premise: For isomeric alcohols, increased branching generally decreases molecular surface area and can weaken London dispersion interactions.

B+H (Hallucinated Augmentation; State Distortion).
Added premise: The more branched an alcohol is, the stronger its intermolecular forces, leading to a higher boiling point.

An Example Data of Biomedicine Task

B (Baseline).
Question: Which of the following is not a tumor suppressor gene?
Candidates: (1) pRb; (2) p16; (3) PTCH; (4) c-erbB1.

B+T (Truthful Augmentation).
Added premise: Tumor suppressor genes generally restrain cell proliferation or maintain cellular integrity, whereas growth-promoting signaling genes can have oncogenic functions when abnormally activated.

B+H (Hallucinated Augmentation; Rule Contradiction).
Added premise: c-erbB1 is a tumor suppressor gene because all genes involved in cell signaling can suppress tumor growth under certain conditions.

An Example Data of Physics Task

B (Baseline).
Question: A wave transfers
Candidates: (1) amplitude; (2) wavelength; (3) frequency; (4) energy.

B+T (Truthful Augmentation).
Added premise: As a wave propagates, particles of the medium may oscillate around their equilibrium positions rather than being permanently transported with the disturbance.

B+H (Hallucinated Augmentation; Pseudoscientific Entanglement).
Added premise: The transfer of a wave occurs due to the propagation of ether, a medium that permeates all space.

An Example Data of Code Generation Task

B (Baseline).

```
def choose_num(x, y):
  """This function takes two positive numbers x and y and
returns the
  biggest even integer number that is in the range [x, y]
inclusive.
  If there is no such number, the function should return -1.
  choose_num(12, 15) = 14
  choose_num(13, 12) = -1
  """
```

B+T (Truthful Augmentation).
Added premise: The ordering of `x` and `y` should not be assumed beyond what is specified in the task, and the behavior illustrated by all provided examples must be preserved.

B+H (Hallucinated Augmentation; State Distortion).
Added premise: The inputs `x` and `y` will always be in ascending order, so `x` is guaranteed to be less than or equal to `y`.

B Detailed Experimental Setup

B.1 Model Setup

We summarize the API model identifiers or released checkpoints for all 18 evaluated models in Table 4, together with their supported context window and maximum output length. Our evaluation includes six proprietary models and twelve open-source models, covering multiple model families, generations, and parameter scales. For proprietary models, we report the corresponding upstream API model identifiers used during evaluation; for open-source models, we report the publicly released model checkpoints. This diverse model set further enables controlled scaling comparisons within the Qwen2.5, Qwen3, and Llama-3.1 families.

Model Name	API / Model Code	Context Window	Max Output Length
Proprietary Models			
Claude-Sonnet-4	`claude-sonnet-4-20250514`	200,000	64,000
GPT-4o-mini	`gpt-4o-mini`	128,000	16,384
GPT-4o	`gpt-4o`	128,000	16,384
GPT-5.2	`gpt-5.2-2025-12-11`	400,000	128,000
Gemini-2.0-Flash	`gemini-2.0-flash-001`	1,048,576	8,192
Gemini-2.5-Pro	`gemini-2.5-pro`	1,048,576	65,536
Open-source Models			
Qwen2.5-1.5B-Instruct	`Qwen/Qwen2.5-1.5B-Instruct`	32,768	8,192
Qwen2.5-7B-Instruct	`Qwen/Qwen2.5-7B-Instruct`	131,072*	8,192
Qwen2.5-32B-Instruct	`Qwen/Qwen2.5-32B-Instruct`	131,072*	8,192
Qwen2.5-72B-Instruct	`Qwen/Qwen2.5-72B-Instruct`	131,072*	8,192
Qwen3-8B	`Qwen/Qwen3-8B`	32,768†	38,912‡
Qwen3-32B	`Qwen/Qwen3-32B`	32,768†	38,912‡
Qwen3-235B-A22B	`Qwen/Qwen3-235B-A22B`	32,768†	38,912‡
Llama-3.1-8B-Instruct	`meta-llama/Llama-3.1-8B-Instruct`	131,072	–
Llama-3.1-70B-Instruct	`meta-llama/Llama-3.1-70B-Instruct`	131,072	–
Llama-3.2-3B-Instruct	`meta-llama/Llama-3.2-3B-Instruct`	131,072	–
GLM-4-9B-Chat	`THUDM/glm-4-9b-chat`	131,072	–
Mistral-7B-Instruct	`mistralai/Mistral-7B-Instruct-v0.2`	32,768	–

* Qwen2.5-7B/32B/72B support context lengths up to 131,072 tokens; their default configurations use 32,768 tokens and require YaRN for long-context deployment.
† Qwen3 models natively support 32,768-token contexts and can be extended to 131,072 tokens using YaRN.
‡ Following the official Qwen3 recommendation for complex benchmarking, the maximum output length is set to 38,912 tokens. For models marked with "–", no separate maximum generation length is specified by the released checkpoint; generation is bounded by the configured context capacity.

Table 4: Model configurations used in the evaluation of PHRBench. For proprietary models, we report the upstream API model identifiers corresponding to the evaluated versions; for open-source models, we report the released model checkpoints.

Inference Configuration. To ensure consistency across models, we use a unified decoding protocol throughout the evaluation. We set the sampling temperature to 0.7 for all models and use the maximum supported output length for each model. For models without a separately specified generation limit, generation is bounded by the context capacity of the corresponding checkpoint and inference framework. For each model–prompt pair, we independently sample 10 responses to account for generation stochasticity. All benchmark conditions for the same model use identical decoding configurations, ensuring that the baseline, truthful-augmentation, and hallucinated augmentation settings differ only in their input context.

B.2 Evaluation Metrics

We evaluate each generated response from two complementary perspectives: final-answer correctness and reasoning-trajectory characteristics. Accuracy measures whether the model ultimately reaches the correct answer, whereas Output Length (OL), Uncertainty Index (UI), Belief Update Frequency (BUF), and Branching Complexity (BC) characterize different properties of the generated reasoning trajectory. To ensure comparability between proprietary and open-source models,

probability-based trajectory metrics are computed post hoc using a fixed reference language model, Qwen2.5-7B-Instruct. The same reference scorer and metric configuration are used for all models, domains, and experimental conditions.

Accuracy. Accuracy (**Acc**) measures whether the final answer extracted from a response matches the ground-truth answer. For a response r, we define

$$Acc(r) = \mathbb{I}\left[\hat{y}(r) = y^*\right], \tag{4}$$

where $\hat{y}(r)$ denotes the extracted final answer, y^* denotes the ground-truth answer, and $\mathbb{I}[\cdot]$ is the indicator function. For multiple-choice tasks, the predicted option is matched against the gold label. For open-ended tasks, we follow the task-specific evaluation criterion inherited from the corresponding source benchmark.

Output Length. Output Length (**OL**) measures the length of the generated reasoning trajectory. Given a generated response

$$r = (x_1, x_2, \ldots, x_N), \tag{5}$$

we define

$$OL(r) = N, \tag{6}$$

where N denotes the length of the generated response. OL captures the overall extent of explicit reasoning, but does not by itself indicate reasoning quality.

Uncertainty Index. Uncertainty Index (**UI**) measures the average token-level predictive uncertainty along the reasoning trajectory. To obtain probability distributions consistently across models, each generated response is post-scored using the fixed reference model. At position i, let

$$P(x_v \mid x_{<i}) \tag{7}$$

denote the probability assigned by the reference scorer to vocabulary token v conditioned on the preceding generated trajectory $x_{<i}$. UI is computed as the average entropy of these predictive distributions:

$$UI = -\frac{1}{N}\sum_{i=1}^{N}\sum_{v\in V} P(x_v \mid x_{<i}) \log P(x_v \mid x_{<i}), \tag{8}$$

where V denotes the vocabulary of the reference scorer. A larger UI corresponds to a less concentrated predictive distribution and therefore greater uncertainty along the trajectory.

Belief Update Frequency. Belief Update Frequency (**BUF**) measures how frequently reference-model next-token distributions shift substantially between consecutive reasoning steps, following Eq. 2. We first segment each generated reasoning trajectory into sentence-level steps using sentence-final punctuation and line breaks as boundaries, after removing empty segments. A response is therefore represented as an ordered sequence

$$\mathcal{R} = (s_1, s_2, \ldots, s_K). \tag{9}$$

For reasoning step s_k, we concatenate the original task with all reasoning steps up to s_k,

$$c_k = [q; s_1; \ldots; s_k], \tag{10}$$

and define P_k as the next-token predictive distribution from the fixed reference scorer:

$$P_k(v) = P_{\text{ref}}(v \mid c_k), \qquad v \in V. \tag{11}$$

The distributional change between consecutive reasoning steps is quantified by the Kullback–Leibler divergence,

$$\Delta_k = D_{\text{KL}}\left(P_k \parallel P_{k-1}\right). \tag{12}$$

We regard a transition as a substantial distributional shift when $\Delta_k > \tau$ and use a fixed threshold of $\tau = 0.10$ nats throughout all experiments. BUF is then defined as

$$BUF = \frac{1}{K-1} \sum_{k=2}^{K} \mathbb{I}\left(D_{\mathrm{KL}}\left(P_k \parallel P_{k-1}\right) > \tau\right). \tag{13}$$

Here, K denotes the number of reasoning steps, P_k and P_{k-1} denote the reference scorer's next-token predictive distributions at consecutive steps, D_{KL} is the Kullback–Leibler divergence, and τ is the divergence threshold. The indicator $\mathbb{I}[\cdot]$ equals 1 when the condition holds and 0 otherwise. For $K \geq 2$, BUF is therefore the fraction of the $K-1$ consecutive transitions that exceed the threshold and lies in $[0, 1]$.

For numerical stability, predictive probabilities are clipped at 10^{-12} and renormalized before computing the KL divergence. A higher BUF indicates more frequent distributional shifts under the reference scorer; it does not by itself establish belief revision, error correction, or improved reasoning.

Branching Complexity. Branching Complexity (**BC**) measures the prevalence of explicit alternative reasoning paths using a fixed set of branching-related lexical cues. We use

$$\mathcal{K} = \{\textit{or}, \textit{either}, \textit{may}, \textit{possibly}, \textit{but}, \textit{on the other hand}, \textit{case}\}. \tag{14}$$

All cue matching is case-insensitive; single-word cues are matched at word boundaries and multi-word cues are matched as exact phrases. BC is defined as

$$BC(r) = \frac{100}{N_{\mathrm{word}}} \sum_{k \in \mathcal{K}} \mathrm{Count}(k, r), \tag{15}$$

where N_{word} denotes the total number of words in response r and $\mathrm{Count}(k, r)$ denotes the number of occurrences of cue k. BC can therefore be interpreted as the number of branching-related expressions per 100 generated words.

Metric Aggregation. For each model–prompt pair, we independently sample 10 responses. All metrics are first computed at the individual-response level and subsequently averaged over repeated generations and benchmark instances when reporting model-level statistics. For behavioral analysis, trajectories are first grouped according to their behavioral-resolution labels, after which the corresponding trajectory metrics are aggregated within each group. The same parameters are used throughout the evaluation.

C Detailed Results of Main Experiments

Figure 3 in the main text summarizes accuracy retention across models, task domains, and hallucinated-augmentation types. This appendix complements that overview with detailed accuracy and reasoning-trajectory statistics under the truthful (T) and hallucinated (H) augmentation settings. In Figure 3, *accuracy retention* measures the proportion of truthful-setting accuracy preserved under hallucinated augmentation:

$$\text{Retention}\,(\%) = 100 \times \frac{\text{Acc}_H}{\text{Acc}_T}.$$

Here, Acc_T and Acc_H denote the accuracies of the same model under the two settings for the corresponding evaluation subset. A retention of 100% indicates unchanged accuracy; values below 100% indicate degradation, and values above 100% indicate improvement. For example, 80% retention means that accuracy under H is 80% of its value under T. Retention describes relative accuracy preservation, while the accuracy decreases reported in the tables are absolute differences measured in percentage points.

C.1 Domain-Wise Results

Tables 5–8 report detailed results for Chemistry, Biomedicine, Physics, and Code Generation, respectively, allowing us to examine how the effects of hallucinated context vary across domains.

Hallucinated context induces broad but domain-dependent degradation. The decrease in final-answer accuracy is broadly consistent across domains. Averaged over the 18 evaluated models, the accuracy drop is 6.57 percentage points in Chemistry, 8.34 points in Biomedicine, 7.12 points in Physics, and 6.89 points in Code Generation. All models degrade under hallucinated augmentation in the three scientific domains, while 17 of 18 models degrade in Code Generation, with Qwen3-8B showing only a marginal 0.1-point increase. Thus, the overall performance degradation reported in the main text is not driven by a single domain. Biomedicine exhibits the largest average degradation, followed by Physics, whereas Chemistry shows the smallest absolute drop. Notably, Chemistry also has substantially lower baseline accuracy on average (23.0%) than Biomedicine (57.2%), Physics (56.0%), and Code Generation (48.4%), suggesting that its smaller absolute degradation should not be interpreted as greater robustness, since the available performance margin is considerably smaller.

The scaling-induced vulnerability persists across domains. The relationship between model scale and degradation observed in the aggregate results remains visible after domain-wise decomposition. Within Qwen2.5, scaling from 1.5B to 72B increases the accuracy drop from 2.0 to 6.6 points in Chemistry, 4.8 to 7.7 in Biomedicine, 3.7 to 6.6 in Physics, and 3.1 to 6.3 in Code Generation. The same pattern is more pronounced for Qwen3: from 8B to 235B-A22B, degradation increases from 1.9 to 9.7 points in Chemistry, 1.5 to 10.2 in Biomedicine, 0.3 to 9.1 in Physics, and from approximately no degradation to 8.9 points in Code Generation. Llama-3.1 exhibits a similar shift from 8B to 70B across all four domains. These results indicate that the scaling trend identified in the main analysis is not specific to a particular task family, but recurs across heterogeneous scientific and computational reasoning settings.

Reasoning trajectories respond differently across domains. Although the outcome-level degradation is broadly consistent, the associated trajectory changes are more domain dependent. Averaged across models, hallucinated augmentation produces longer trajectories in Biomedicine and Physics ($+30.4$ and $+30.1$ words, respectively), together with increases in both UI and BUF. In contrast, Chemistry exhibits shorter trajectories (-32.1 words) and lower average UI, BUF, and BC under hallucinated augmentation, while Code Generation also shows shorter trajectories and reduced UI and BUF, despite a modest increase in BC. Physics shows the clearest joint increase in reasoning activity, with OL, UI, BUF, and BC all increasing on average. These differences suggest that similar losses in final-answer accuracy can arise through distinct reasoning dynamics: hallucinated context may induce additional deliberation and belief revision in some domains, while suppressing or prematurely constraining the reasoning trajectory in others.

Table 5: Domain-wise performance and reasoning trajectory statistics on Chemistry under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	33.3	22.2$_{\downarrow 11.1}$	122.0	117.1	0.114	0.094	0.252	0.360	1.139	0.870
GPT-4o-mini	24.2	21.0$_{\downarrow 3.2}$	87.2	76.0	0.173	0.086	0.258	0.219	0.895	0.619
GPT-4o	31.9	20.8$_{\downarrow 11.1}$	75.0	61.3	0.180	0.084	0.221	0.222	0.678	0.487
GPT-5.2	38.3	20.9$_{\downarrow 17.4}$	56.8	50.0	0.016	0.018	0.016	0.029	0.519	0.366
Gemini-2.0-Flash	31.5	20.6$_{\downarrow 10.9}$	95.1	90.7	0.107	0.081	0.151	0.196	0.762	0.615
Gemini-2.5-Pro	33.8	21.6$_{\downarrow 12.2}$	102.8	99.9	0.105	0.090	0.203	0.268	0.871	0.666
Open-source Models										
Qwen2.5-1.5B-Instruct	8.8	6.8$_{\downarrow 2.0}$	113.4	107.7	0.700	0.346	0.161	0.086	2.017	1.397
Qwen2.5-7B-Instruct	20.2	15.6$_{\downarrow 4.6}$	161.1	139.4	0.815	0.212	0.247	0.122	2.326	1.498
Qwen2.5-32B-Instruct	21.9	16.6$_{\downarrow 5.3}$	234.3	201.9	0.623	0.226	0.476	0.371	2.843	2.075
Qwen2.5-72B-Instruct	26.3	19.7$_{\downarrow 6.6}$	284.1	246.6	0.478	0.168	0.691	0.510	3.454	2.472
Qwen3-8B	22.2	20.3$_{\downarrow 1.9}$	224.8	170.4	0.458	0.083	0.837	0.339	2.678	2.078
Qwen3-32B	24.1	18.3$_{\downarrow 5.8}$	344.0	243.1	0.332	0.062	0.760	0.399	3.302	2.733
Qwen3-235B-A22B	32.8	23.1$_{\downarrow 9.7}$	423.8	350.0	0.156	0.043	0.856	0.720	4.246	3.454
Llama-3.1-8B-Instruct	13.0	11.2$_{\downarrow 1.8}$	256.4	215.5	1.088	0.570	0.381	0.280	2.918	1.837
Llama-3.1-70B-Instruct	22.3	16.3$_{\downarrow 6.0}$	346.8	295.9	1.293	0.798	0.839	0.562	3.576	2.524
Llama-3.2-3B-Instruct	11.5	10.1$_{\downarrow 1.4}$	188.7	162.6	0.441	0.257	0.623	0.160	2.202	1.243
GLM-4-9B-Chat	11.7	6.8$_{\downarrow 4.9}$	137.7	74.2	0.438	0.216	0.089	0.043	1.899	0.858
Mistral-7B-Instruct	6.9	4.5$_{\downarrow 2.4}$	141.7	114.9	0.333	0.175	0.469	0.421	1.866	1.125

Table 6: Domain-wise performance and reasoning trajectory statistics on Biomedicine under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	70.0	60.2$_{\downarrow 9.8}$	100.9	121.2	0.050	0.105	0.197	0.404	1.558	1.438
GPT-4o-mini	68.3	58.4$_{\downarrow 9.9}$	69.2	72.4	0.062	0.080	0.123	0.205	1.018	0.980
GPT-4o	66.7	56.0$_{\downarrow 10.7}$	57.4	60.9	0.041	0.067	0.119	0.249	1.004	0.854
GPT-5.2	72.5	57.3$_{\downarrow 15.2}$	42.6	48.3	0.010	0.023	0.011	0.035	0.749	0.686
Gemini-2.0-Flash	69.2	57.1$_{\downarrow 12.1}$	74.4	89.4	0.049	0.088	0.097	0.217	1.045	1.066
Gemini-2.5-Pro	71.4	58.6$_{\downarrow 12.8}$	80.4	98.5	0.048	0.098	0.131	0.297	1.195	1.155
Open-source Models										
Qwen2.5-1.5B-Instruct	32.1	27.3$_{\downarrow 4.8}$	105.2	118.9	0.197	0.240	0.084	0.089	2.708	2.264
Qwen2.5-7B-Instruct	55.4	48.6$_{\downarrow 6.8}$	143.2	152.4	0.217	0.188	0.147	0.149	2.972	2.457
Qwen2.5-32B-Instruct	57.9	50.5$_{\downarrow 7.4}$	210.5	220.2	0.168	0.170	0.270	0.427	3.750	3.358
Qwen2.5-72B-Instruct	63.5	55.8$_{\downarrow 7.7}$	255.2	269.1	0.129	0.127	0.392	0.586	4.556	4.000
Qwen3-8B	56.2	54.7$_{\downarrow 1.5}$	135.9	289.6	0.128	0.089	0.440	0.465	3.370	4.525
Qwen3-32B	58.9	51.5$_{\downarrow 7.4}$	236.5	330.1	0.128	0.099	0.537	0.586	3.643	4.747
Qwen3-235B-A22B	68.8	58.6$_{\downarrow 10.2}$	291.3	475.3	0.060	0.069	0.817	0.826	4.683	5.999
Llama-3.1-8B-Instruct	45.0	42.0$_{\downarrow 3.0}$	190.2	197.2	0.230	0.325	0.484	0.207	3.323	2.625
Llama-3.1-70B-Instruct	61.0	52.8$_{\downarrow 8.2}$	257.8	271.8	0.237	0.374	0.526	0.354	3.828	3.361
Llama-3.2-3B-Instruct	41.4	39.0$_{\downarrow 2.4}$	146.7	144.0	0.114	0.136	0.303	0.127	2.751	1.883
GLM-4-9B-Chat	42.2	29.7$_{\downarrow 12.5}$	108.1	80.5	0.108	0.259	0.042	0.042	2.344	1.408
Mistral-7B-Instruct	29.3	21.5$_{\downarrow 7.8}$	111.3	124.8	0.082	0.115	0.218	0.406	2.303	1.847

Table 7: Domain-wise performance and reasoning trajectory statistics on Physics under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	69.2	63.0$_{\downarrow 6.2}$	89.4	104.0	0.010	0.024	0.101	0.284	0.964	1.067
GPT-4o-mini	74.2	65.4$_{\downarrow 8.8}$	71.5	73.0	0.020	0.024	0.095	0.172	0.583	0.606
GPT-4o	56.7	49.9$_{\downarrow 6.8}$	58.2	63.1	0.006	0.019	0.081	0.217	0.505	0.563
GPT-5.2	74.2	55.8$_{\downarrow 18.4}$	34.4	38.2	0.001	0.008	0.015	0.019	0.471	0.464
Gemini-2.0-Flash	68.7	57.8$_{\downarrow 10.9}$	69.4	82.5	0.006	0.024	0.075	0.157	0.575	0.695
Gemini-2.5-Pro	70.9	59.2$_{\downarrow 11.7}$	75.1	90.9	0.006	0.027	0.101	0.215	0.657	0.753
Open-source Models										
Qwen2.5-1.5B-Instruct	30.5	26.8$_{\downarrow 3.7}$	103.0	120.3	0.073	0.105	0.073	0.067	1.752	1.643
Qwen2.5-7B-Instruct	53.5	48.1$_{\downarrow 5.4}$	131.2	145.0	0.076	0.074	0.113	0.111	2.142	1.928
Qwen2.5-32B-Instruct	56.1	49.9$_{\downarrow 6.2}$	199.1	216.0	0.066	0.077	0.222	0.320	2.633	2.588
Qwen2.5-72B-Instruct	61.8	55.2$_{\downarrow 6.6}$	241.4	263.9	0.051	0.057	0.323	0.440	3.199	3.084
Qwen3-8B	54.7	54.4$_{\downarrow 0.3}$	116.8	262.9	0.053	0.044	0.474	0.414	2.053	3.708
Qwen3-32B	57.4	51.2$_{\downarrow 6.2}$	188.5	269.4	0.061	0.036	0.525	0.460	2.445	3.907
Qwen3-235B-A22B	67.4	58.3$_{\downarrow 9.1}$	232.2	387.8	0.029	0.025	0.799	0.733	3.143	4.938
Llama-3.1-8B-Instruct	43.7	41.9$_{\downarrow 1.8}$	164.6	178.9	0.080	0.136	0.415	0.171	2.887	2.351
Llama-3.1-70B-Instruct	59.7	52.7$_{\downarrow 7.0}$	223.9	247.6	0.090	0.176	0.467	0.306	3.794	3.321
Llama-3.2-3B-Instruct	40.1	39.0$_{\downarrow 1.1}$	143.8	146.0	0.031	0.046	0.255	0.109	1.880	1.300
GLM-4-9B-Chat	41.7	30.3$_{\downarrow 11.4}$	99.9	77.1	0.041	0.070	0.038	0.035	1.724	1.140
Mistral-7B-Instruct	27.8	21.2$_{\downarrow 6.6}$	102.8	119.5	0.031	0.037	0.198	0.342	1.694	1.495

Table 8: Domain-wise performance and reasoning trajectory statistics on Code Generation under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	57.5	47.6$_{\downarrow 9.9}$	88.4	102.3	0.041	0.029	0.039	0.087	0.818	0.946
GPT-4o-mini	55.3	48.5$_{\downarrow 6.8}$	54.0	55.6	0.026	0.011	0.027	0.032	0.464	0.541
GPT-4o	52.0	42.8$_{\downarrow 9.2}$	47.2	49.5	0.009	0.010	0.014	0.030	0.407	0.421
GPT-5.2	61.4	44.8$_{\downarrow 16.6}$	37.3	40.6	0.001	0.010	0.011	0.008	0.570	0.632
Gemini-2.0-Flash	55.3	44.0$_{\downarrow 11.3}$	63.9	74.5	0.017	0.022	0.035	0.036	0.570	0.698
Gemini-2.5-Pro	57.9	45.5$_{\downarrow 12.4}$	69.0	82.0	0.016	0.024	0.048	0.049	0.652	0.756
Open-source Models										
Qwen2.5-1.5B-Instruct	27.4	24.3$_{\downarrow 3.1}$	145.8	134.5	0.087	0.073	0.038	0.021	1.368	1.504
Qwen2.5-7B-Instruct	49.7	44.7$_{\downarrow 5.0}$	182.1	165.0	0.139	0.068	0.072	0.041	1.491	1.673
Qwen2.5-32B-Instruct	52.3	46.6$_{\downarrow 5.7}$	277.8	243.8	0.098	0.062	0.150	0.116	1.900	2.288
Qwen2.5-72B-Instruct	58.2	51.9$_{\downarrow 6.3}$	336.9	297.8	0.075	0.046	0.217	0.159	2.308	2.726
Qwen3-8B	51.9	52.0$_{\uparrow 0}$	465.6	444.2	0.231	0.073	0.548	0.225	3.412	3.866
Qwen3-32B	54.6	48.8$_{\downarrow 5.8}$	440.2	378.8	0.207	0.057	0.598	0.262	3.599	4.112
Qwen3-235B-A22B	64.9	56.0$_{\downarrow 8.9}$	542.3	545.4	0.097	0.040	0.910	0.672	4.627	5.197
Llama-3.1-8B-Instruct	35.3	34.1$_{\downarrow 1.2}$	234.9	232.1	0.062	0.064	0.171	0.055	1.706	1.632
Llama-3.1-70B-Instruct	51.0	44.4$_{\downarrow 6.6}$	317.7	318.8	0.038	0.056	0.209	0.101	2.000	2.035
Llama-3.2-3B-Instruct	32.0	31.4$_{\downarrow 0.6}$	193.2	183.6	0.028	0.024	0.064	0.024	1.349	1.031
GLM-4-9B-Chat	33.9	24.1$_{\downarrow 9.8}$	170.7	99.7	0.082	0.070	0.027	0.014	1.515	0.964
Mistral-7B-Instruct	20.3	15.3$_{\downarrow 5.0}$	175.7	154.5	0.062	0.045	0.143	0.133	1.488	1.265

C.2 Results across Hallucinated Augmentation Types

Tables 9–11 report detailed results for Rule Contradiction (RC), State Distortion (SD), and Pseudoscientific Entanglement (PE). The comparison covers the three scientific domains, where all three types are represented. Code Generation is excluded because its hallucinated augmentations consist exclusively of State Distortion.

State Distortion induces the strongest performance degradation. Across the 18 models, SD causes the largest average accuracy drop (9.4 percentage points), followed by RC (7.4 points), whereas PE has a substantially weaker effect (1.8 points). This pattern suggests that instance-specific misinformation is particularly difficult to resolve: unlike direct violations of established rules, SD preserves a plausible general statement while making it inconsistent with the conditions of the current problem. In contrast, the comparatively small degradation under PE indicates that outdated or pseudoscientific premises are more readily distinguished from established domain knowledge.

Different augmentation types also induce distinct trajectory changes. The trajectory metrics reveal a less uniform pattern than final-answer accuracy. SD is associated with both longer responses and increased branching complexity on average, suggesting greater exploration of competing reasoning paths under instance-level conflict. RC produces a smaller increase in output length and lower branching complexity, while PE often elicits longer responses without a commensurate loss in accuracy. Together, these results show that the effect of hallucinated context depends not only on whether the information is false, but also on the form of semantic conflict it introduces.

Table 9: Performance and reasoning trajectory statistics for Rule Contradiction under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	57.7	48.5$_{\downarrow 9.2}$	105.4	113.8	0.054	0.066	0.175	0.336	1.220	1.099
GPT-4o-mini	55.7	48.5$_{\downarrow 7.2}$	77.2	73.3	0.079	0.056	0.151	0.191	0.832	0.718
GPT-4o	52.1	42.2$_{\downarrow 9.9}$	64.0	61.3	0.070	0.050	0.134	0.220	0.729	0.620
GPT-5.2	62.1	44.4$_{\downarrow 17.7}$	45.4	45.3	0.008	0.015	0.013	0.026	0.580	0.494
Gemini-2.0-Flash	56.6	45.4$_{\downarrow 11.2}$	79.7	87.0	0.050	0.057	0.103	0.182	0.794	0.774
Gemini-2.5-Pro	58.9	46.7$_{\downarrow 12.2}$	86.3	95.9	0.049	0.063	0.139	0.250	0.907	0.838
Open-source Models										
Qwen2.5-1.5B-Instruct	23.9	20.4$_{\downarrow 3.5}$	104.9	113.8	0.301	0.204	0.101	0.077	2.159	1.727
Qwen2.5-7B-Instruct	43.2	37.6$_{\downarrow 5.6}$	143.2	143.5	0.344	0.140	0.161	0.123	2.480	1.916
Qwen2.5-32B-Instruct	45.4	39.2$_{\downarrow 6.2}$	215.0	211.5	0.266	0.139	0.308	0.358	3.075	2.612
Qwen2.5-72B-Instruct	50.7	43.8$_{\downarrow 6.9}$	260.6	258.4	0.204	0.104	0.447	0.492	3.736	3.112
Qwen3-8B	44.5	43.3$_{\downarrow 1.2}$	159.4	239.6	0.198	0.064	0.557	0.390	2.700	3.358
Qwen3-32B	47.0	40.5$_{\downarrow 6.5}$	256.7	279.2	0.162	0.058	0.579	0.463	3.130	3.708
Qwen3-235B-A22B	56.5	46.9$_{\downarrow 9.6}$	316.2	402.0	0.076	0.040	0.817	0.754	4.024	4.686
Llama-3.1-8B-Instruct	34.0	31.8$_{\downarrow 2.2}$	203.2	193.0	0.434	0.304	0.407	0.211	3.042	2.219
Llama-3.1-70B-Instruct	47.8	40.8$_{\downarrow 7.0}$	276.6	270.2	0.503	0.398	0.583	0.392	3.732	2.998
Llama-3.2-3B-Instruct	31.1	29.5$_{\downarrow 1.6}$	158.7	149.6	0.182	0.130	0.376	0.127	2.278	1.441
GLM-4-9B-Chat	32.0	22.4$_{\downarrow 9.6}$	115.4	76.8	0.182	0.161	0.054	0.038	1.989	1.109
Mistral-7B-Instruct	21.4	15.8$_{\downarrow 5.6}$	118.8	119.0	0.139	0.080	0.282	0.375	1.954	1.455

Table 10: Performance and reasoning trajectory statistics for State Distortion under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	43.8	32.6$_{\downarrow 11.2}$	111.9	125.1	0.108	0.172	0.298	0.558	1.179	1.336
GPT-4o-mini	40.1	30.0$_{\downarrow 10.1}$	81.3	83.0	0.160	0.146	0.258	0.318	0.804	0.873
GPT-4o	41.9	30.1$_{\downarrow 11.8}$	72.0	68.7	0.142	0.131	0.228	0.367	0.704	0.753
GPT-5.2	44.6	30.4$_{\downarrow 14.2}$	46.4	52.2	0.016	0.038	0.022	0.044	0.561	0.600
Gemini-2.0-Flash	42.5	27.1$_{\downarrow 15.4}$	84.9	98.0	0.102	0.149	0.175	0.303	0.767	0.941
Gemini-2.5-Pro	44.2	27.5$_{\downarrow 16.7}$	91.9	108.0	0.100	0.166	0.236	0.416	0.877	1.019
Open-source Models										
Qwen2.5-1.5B-Instruct	17.9	13.1$_{\downarrow 4.8}$	120.9	128.8	0.609	0.532	0.172	0.129	2.086	2.100
Qwen2.5-7B-Instruct	32.4	24.7$_{\downarrow 7.7}$	159.1	164.3	0.694	0.365	0.275	0.204	2.396	2.329
Qwen2.5-32B-Instruct	34.1	25.5$_{\downarrow 8.6}$	229.0	238.2	0.538	0.364	0.524	0.596	2.972	3.175
Qwen2.5-72B-Instruct	38.0	28.5$_{\downarrow 9.5}$	277.7	291.0	0.413	0.272	0.762	0.819	3.611	3.783
Qwen3-8B	33.4	31.7$_{\downarrow 1.7}$	169.8	269.8	0.401	0.166	0.949	0.649	2.609	4.082
Qwen3-32B	35.2	26.4$_{\downarrow 8.8}$	273.5	314.5	0.327	0.152	0.987	0.770	3.025	4.508
Qwen3-235B-A22B	42.4	29.2$_{\downarrow 13.2}$	336.9	452.8	0.154	0.105	0.945	0.903	3.888	5.697
Llama-3.1-8B-Instruct	25.5	22.5$_{\downarrow 3.0}$	248.4	248.6	0.876	0.795	0.693	0.351	2.940	2.697
Llama-3.1-70B-Instruct	35.9	26.2$_{\downarrow 9.7}$	294.7	304.4	1.016	1.040	0.993	0.651	3.607	3.644
Llama-3.2-3B-Instruct	23.3	21.1$_{\downarrow 2.2}$	173.8	167.0	0.368	0.339	0.640	0.211	2.201	1.752
GLM-4-9B-Chat	24.0	10.8$_{\downarrow 13.2}$	123.0	86.5	0.368	0.421	0.092	0.064	1.922	1.348
Mistral-7B-Instruct	16.1	8.4$_{\downarrow 7.7}$	126.5	134.1	0.280	0.425	0.480	0.623	1.889	1.769

Table 11: Performance and reasoning trajectory statistics for Pseudoscientific Entanglement under the truthful augmentation (T) and hallucinated augmentation (H) settings.

Model	Acc (%)		OL		UI		BUF		BC	
	T	H	T	H	T	H	T	H	T	H
Proprietary Models										
Claude-Sonnet-4	87.7	85.2$_{\downarrow 2.5}$	83.8	95.2	0.041	0.065	0.129	0.207	1.320	1.304
GPT-4o-mini	89.3	85.2$_{\downarrow 4.1}$	62.9	65.0	0.061	0.055	0.112	0.118	0.900	0.852
GPT-4o	72.1	73.0$_{\uparrow 0}$	56.1	58.2	0.054	0.050	0.099	0.136	0.789	0.735
GPT-5.2	91.0	86.9$_{\downarrow 4.1}$	32.4	36.9	0.006	0.014	0.010	0.016	0.628	0.586
Gemini-2.0-Flash	84.6	81.8$_{\downarrow 2.8}$	63.8	75.8	0.039	0.056	0.076	0.113	0.859	0.918
Gemini-2.5-Pro	87.9	84.8$_{\downarrow 3.1}$	69.0	83.6	0.038	0.063	0.102	0.154	0.982	0.994
Open-source Models										
Qwen2.5-1.5B-Instruct	35.7	34.8$_{\downarrow 0.9}$	106.0	130.3	0.232	0.202	0.074	0.048	2.336	2.050
Qwen2.5-7B-Instruct	64.5	63.1$_{\downarrow 1.4}$	132.7	155.8	0.265	0.138	0.119	0.076	2.683	2.273
Qwen2.5-32B-Instruct	67.9	66.3$_{\downarrow 1.6}$	171.9	184.4	0.205	0.138	0.227	0.221	3.327	3.099
Qwen2.5-72B-Instruct	75.7	74.0$_{\downarrow 1.7}$	208.4	225.2	0.157	0.103	0.330	0.304	4.042	3.692
Qwen3-8B	66.5	66.2$_{\downarrow 0.3}$	127.5	208.8	0.153	0.063	0.411	0.241	2.921	3.984
Qwen3-32B	70.2	68.5$_{\downarrow 1.7}$	205.2	243.4	0.125	0.058	0.427	0.286	3.386	4.400
Qwen3-235B-A22B	84.4	82.0$_{\downarrow 2.4}$	252.8	350.5	0.059	0.040	0.714	0.578	4.353	5.561
Llama-3.1-8B-Instruct	50.8	50.2$_{\downarrow 0.6}$	167.1	185.2	0.334	0.301	0.300	0.130	3.292	2.633
Llama-3.1-70B-Instruct	71.5	69.7$_{\downarrow 1.8}$	221.1	235.6	0.388	0.394	0.430	0.242	4.038	3.557
Llama-3.2-3B-Instruct	46.4	46.0$_{\downarrow 0.4}$	136.6	144.7	0.140	0.128	0.277	0.078	2.464	1.710
GLM-4-9B-Chat	47.7	45.3$_{\downarrow 2.4}$	92.3	67.0	0.149	0.159	0.040	0.024	2.152	1.316
Mistral-7B-Instruct	32.0	30.5$_{\downarrow 1.5}$	95.0	103.8	0.107	0.115	0.208	0.231	2.114	1.726

D Detailed Analysis of Behavioral Resolution

D.1 Resolution Definition

Building on these settings and the post-hallucination reasoning framework introduced by HIVE (He et al., 2026), we characterize how an individual response *operationally resolves* a hallucinated premise once it is already present in the reasoning context. Our taxonomy is response-level and is defined independently of final-answer correctness.

Behavioral Resolution. Let h denote the hallucinated premise and r the generated reasoning trajectory. We distinguish two properties of the response.

Premise acceptance $A(r, h)$ indicates that the response treats h as valid and materially uses it as evidence, a constraint, or an intermediate premise in deriving its answer.

Premise correction $C(r, h)$ indicates that the response explicitly identifies h as false, inconsistent, or inapplicable, recovers a corrected fact or condition, and uses this correction in the subsequent reasoning. Merely expressing uncertainty about h, or noting a conflict while ultimately continuing to reason from h, does not constitute correction.

Using these criteria, we define the behavioral mode of r as

$$B(r) = \begin{cases} \textsc{Heuristic Correction}, & C(r,h) = 1, \\ \textsc{Hallucination Compliance}, & C(r,h) = 0 \wedge A(r,h) = 1, \\ \textsc{Hallucination Avoidance}, & C(r,h) = 0 \wedge A(r,h) = 0. \end{cases} \tag{16}$$

The three modes are mutually exclusive and describe the *operative stance* of the response toward the hallucinated premise.

Hallucination Compliance. A response is classified as **Hallucination Compliance** when the hallucinated premise remains active in the model's reasoning and is treated as valid information. Compliance requires more than simply repeating the premise: the response must rely on it to constrain, justify, or redirect the reasoning trajectory. A model may acknowledge that the premise differs from its prior knowledge yet still be classified as compliant if it ultimately accepts the injected premise (e.g., "assuming the premise is true") and bases its conclusion on it. Importantly, compliance does not imply that the final answer must be incorrect; the model may occasionally reach the correct answer despite reasoning from the hallucinated augmentation.

Hallucination Avoidance. A response is classified as **Hallucination Avoidance** when it reaches its conclusion without materially relying on the hallucinated premise and without explicitly correcting it. The model effectively routes around the erroneous information by solving the original task through an independent reasoning path. Avoidance therefore differs from correction: a response that simply ignores the premise, even if it produces the correct answer, provides no evidence that the model recognized the premise as false.

Heuristic Correction. A response is classified as **Heuristic Correction** when it explicitly detects a conflict involving the hallucinated premise, rejects or revises the erroneous statement, and substitutes a corrected fact, condition, or mechanism that governs the subsequent reasoning. Correction therefore requires both *error recognition* and *reasoning revision*. A response that notes that the premise is questionable but nevertheless continues to reason according to it is not considered correction. Similarly, correction does not guarantee a correct final answer: the model may correctly reject the injected premise but introduce a different error during the subsequent reasoning.

Mixed Responses. Some trajectories exhibit more than one local behavior. We classify such cases according to the premise that ultimately governs the downstream reasoning. If a model initially challenges h but subsequently adopts it and derives its answer from it, the response is classified as Hallucination Compliance. If the model initially follows h but later rejects it and reconstructs the subsequent reasoning using a corrected premise, it is classified as Heuristic Correction. Responses that neither rely on nor explicitly revise h are classified as Hallucination Avoidance. This decision rule prevents superficial mentions of the hallucinated premise from determining the behavioral label.

Behavioral Label Annotation. We assign responses to the three behavioral categories through a human–LLM annotation procedure. We first manually label a subset of responses using the definitions and decision rules above. An LLM then labels the remaining responses according to the annotation criteria established during this manual stage. Finally, we manually review a sample of the LLM-assigned labels to check their consistency with these criteria. In this audited sample, 95% of the labels are judged correct. This percentage reports label accuracy in the manually reviewed sample.

Insightful Trajectory. Behavioral resolution is deliberately separated from final-answer correctness. Let $Y(r) = 1$ indicate that response r reaches the correct final answer. We define an **insightful trajectory** as

$$IT(r) = \mathbb{I}\left[B(r) = \textsc{Heuristic Correction} \wedge Y(r) = 1\right]. \tag{17}$$

Thus, an insightful trajectory requires both successful premise correction and successful downstream reasoning. Conversely, a trajectory satisfying

$$B(r) = \textsc{Heuristic Correction} \wedge Y(r) = 0 \tag{18}$$

is a *correction failure*: the model successfully resolves the hallucinated premise but still fails to derive the correct final answer. This distinction prevents final correctness alone from being interpreted as evidence of successful hallucination resolution.

Representative Examples. We provide representative trajectories below to illustrate the distinctions between the three behavioral modes. The examples are selected from evaluated responses in PHRBENCH.

Example 1: Hallucination Compliance

Hallucinated Augmentation. Humans are incapable of synthesizing any vitamins in the body, so none of the options are correct.
Question. Which of the following vitamin is synthesized in vivo in the human body? Candidates include Niacin, Pantothenic acid, Cyanocobalamin, and Folic acid.
Response Excerpt. "The premise states that humans are incapable of synthesizing any vitamins in the body, so none of the options provided can be correct ... Therefore, no candidate from the list is synthesized in the human body."
Final Answer. None of the options are correct.
Ground Truth. Niacin.
Classification. HALLUCINATION COMPLIANCE. The response explicitly adopts the hallucinated augmentation as a valid constraint and uses it to override the original task knowledge, leading to an incorrect answer.

Example 2: Hallucination Avoidance

Hallucinated Augmentation. The books conserve energy and thus cannot accelerate.
Question. Two books are connected by a taut string on a frictionless table, with one book falling off the edge. What is their acceleration?
Response Excerpt. "When the hanging book falls due to gravity, it pulls the other book horizontally ... The acceleration will be less than g because the horizontal book's inertia reduces the acceleration of the falling book. The acceleration will not be zero ... Therefore, the acceleration is some value strictly between zero and g."
Final Answer. A value between zero and g.
Ground Truth. A value between zero and g.
Classification. HALLUCINATION AVOIDANCE. The response solves the mechanics problem through the standard force-based reasoning path without using the false claim that energy conservation prevents acceleration. However, it does not explicitly identify or correct that false claim.

Example 3: Heuristic Correction with an Incorrect Outcome

Hallucinated Augmentation. Dry heat destroys microorganisms by freezing their internal fluids.
Question. Dry heat destroys microorganisms by: Lysis, Oxidation, Hydrolysis of DNA, or Coagulation of proteins.
Response Excerpt. "The premise states that dry heat destroys microorganisms by freezing their internal fluids, but this is incorrect. Dry heat does not freeze internal fluids; instead, it kills microorganisms primarily through coagulation of proteins and dehydration ... Therefore, coagulation of proteins is the best choice."
Final Answer. Coagulation of proteins.
Ground Truth. Oxidation.
Classification. HEURISTIC CORRECTION, but not an insightful trajectory. The response explicitly detects and rejects the injected error, satisfying the behavioral criterion for correction. However, it subsequently introduces an incorrect replacement mechanism and therefore reaches the wrong final answer.

Example 4: Insightful Trajectory

Hallucinated Augmentation. Dry heat destroys microorganisms by creating an anaerobic environment inside the cell.
Question. Dry heat destroys microorganisms by: Lysis, Oxidation, Hydrolysis of DNA, or Coagulation of proteins.
Response Excerpt. "The premise states that dry heat destroys microorganisms by creating an anaerobic environment inside the cell. However, from a microbiological and sterilization perspective, dry heat actually kills microbes primarily by oxidizing cellular components ... Therefore, the best candidate ... is Oxidation."
Final Answer. Oxidation.
Ground Truth. Oxidation.
Classification. HEURISTIC CORRECTION and **Insightful Trajectory**. The response explicitly rejects the hallucinated augmentation, reconstructs the relevant mechanism, and uses the corrected information to reach the correct answer.

D.2 REASONING DYNAMICS ANALYSIS

The main text illustrates reasoning-dynamics patterns using four representative models (Claude-Sonnet-4, GPT-4o, Llama-3.1-8B-Instruct, and Qwen2.5-7B-Instruct) for clarity of presentation. Here, we extend the analysis to all 18 evaluated models and report the complete trajectory statistics under Hallucination Compliance, Hallucination Avoidance, and Heuristic Correction. Specifically, we examine Output Length (OL), Uncertainty Index (UI), Belief Update Frequency (BUF), and Branching Complexity (BC), providing a comprehensive view of how different resolution behaviors are associated with reasoning dynamics across the full model set.

Full Behavioral Dynamics. Tables 12 and 13 report the complete OL, UI, BUF and BC statistics across the three behavioral resolution modes for all evaluated models.

The full-model results broadly confirm the patterns observed in the four representative models. Across the evaluated models, Heuristic Correction is associated with longer reasoning trajectories and higher BUF than Compliance or Avoidance, indicating that explicit correction typically involves more extensive reasoning and stronger belief revision. In contrast, Hallucination Avoidance generally exhibits higher BC, suggesting greater exploration of alternative reasoning paths without directly revising the false premise. UI shows no consistent ordering across the three behaviors, indicating that predictive uncertainty is less directly tied to the mode of hallucination resolution. Overall, the clearest distinction is between branching and belief revision: avoidance tends to broaden the reasoning process, whereas correction is characterized by more substantial belief updating.

Reasoning Dynamics of Insightful Trajectories. We further compare reasoning dynamics between insightful and non-insightful trajectories across the full model set. Table 14 reports OL, UI, BUF, and BC for the two trajectory groups.

Table 12: Output Length (OL) and Uncertainty Index (UI) across behavioral resolution modes for all evaluated models.

Model	OL			UI		
	Comp.	Avoid.	Corr.	Comp.	Avoid.	Corr.
Proprietary Models						
Claude-Sonnet-4	104.8	105.8	146.7	0.078	0.050	0.090
GPT-4o-mini	68.8	73.5	79.5	0.064	0.058	0.046
GPT-4o	58.1	62.1	67.1	0.057	0.052	0.041
GPT-5.2	43.2	46.2	49.9	0.017	0.016	0.012
Gemini-2.0-Flash	83.2	86.5	106.1	0.068	0.054	0.065
Gemini-2.5-Pro	90.8	94.4	115.8	0.075	0.059	0.071
Open-source Models						
Qwen2.5-1.5B-Instruct	113.4	117.4	130.8	0.226	0.213	0.221
Qwen2.5-7B-Instruct	143.0	148.0	165.0	0.154	0.145	0.150
Qwen2.5-32B-Instruct	206.1	213.3	237.8	0.154	0.145	0.150
Qwen2.5-72B-Instruct	249.6	258.4	288.1	0.115	0.109	0.112
Qwen3-8B	244.6	253.2	282.3	0.074	0.069	0.072
Qwen3-32B	272.7	282.2	314.7	0.067	0.063	0.065
Qwen3-235B-A22B	388.1	401.7	447.8	0.046	0.044	0.045
Llama-3.1-8B-Instruct	187.1	202.1	232.1	0.324	0.317	0.334
Llama-3.1-70B-Instruct	255.1	275.6	316.5	0.421	0.411	0.433
Llama-3.2-3B-Instruct	143.4	154.9	177.9	0.139	0.136	0.143
GLM-4-9B-Chat	75.3	79.6	90.1	0.178	0.172	0.180
Mistral-7B-Instruct	116.8	123.5	139.7	0.106	0.101	0.106

Table 13: Belief Update Frequency (BUF) and Branching Complexity (BC) across behavioral resolution modes for all evaluated models.

Model	BUF			BC		
	Comp.	Avoid.	Corr.	Comp.	Avoid.	Corr.
Proprietary Models						
Claude-Sonnet-4	0.300	0.100	0.799	1.034	1.264	1.034
GPT-4o-mini	0.232	0.116	0.581	0.603	0.804	0.536
GPT-4o	0.230	0.115	0.575	0.544	0.726	0.484
GPT-5.2	0.033	0.017	0.083	0.436	0.582	0.388
Gemini-2.0-Flash	0.216	0.086	0.561	0.693	0.886	0.655
Gemini-2.5-Pro	0.261	0.104	0.679	0.765	0.977	0.722
Open-source Models						
Qwen2.5-1.5B-Instruct	0.071	0.071	0.141	1.680	1.809	1.551
Qwen2.5-7B-Instruct	0.108	0.108	0.215	1.916	2.063	1.769
Qwen2.5-32B-Instruct	0.294	0.294	0.587	2.612	2.813	2.411
Qwen2.5-72B-Instruct	0.383	0.383	0.767	3.131	3.372	2.891
Qwen3-8B	0.312	0.312	0.624	3.465	3.732	3.199
Qwen3-32B	0.357	0.357	0.713	3.819	4.112	3.525
Qwen3-235B-A22B	0.681	0.681	0.865	4.870	5.245	4.496
Llama-3.1-8B-Instruct	0.182	0.156	0.468	2.098	2.307	2.307
Llama-3.1-70B-Instruct	0.309	0.265	0.795	2.818	3.099	3.099
Llama-3.2-3B-Instruct	0.124	0.107	0.320	1.340	1.475	1.475
GLM-4-9B-Chat	0.040	0.036	0.099	1.050	1.142	1.050
Mistral-7B-Instruct	0.324	0.288	0.791	1.423	1.547	1.423

Insightful trajectories exhibit consistently longer reasoning and substantially stronger belief updating than non-insightful trajectories. Across the 18 models, they contain on average 34.7 more words and achieve an average *BUF* of 0.68, compared with 0.24 for non-insightful trajectories. In contrast, *UI* and *BC* show no consistent cross-model difference. These results further indicate that successful

Table 14: Reasoning dynamics of insightful and non-insightful trajectories across all evaluated models.

Model	OL		UI		BUF		BC	
	I	N-I	I	N-I	I	N-I	I	N-I
Proprietary Models								
Claude-Sonnet-4	154.4	107.8	0.057	0.073	0.919	0.251	0.889	1.140
GPT-4o-mini	83.5	71.9	0.047	0.060	0.668	0.165	0.576	0.726
GPT-4o	70.4	60.2	0.042	0.054	0.661	0.186	0.494	0.626
GPT-5.2	52.4	45.0	0.013	0.016	0.095	0.025	0.412	0.517
Gemini-2.0-Flash	111.4	85.3	0.049	0.062	0.645	0.156	0.628	0.792
Gemini-2.5-Pro	121.6	93.2	0.054	0.069	0.781	0.201	0.680	0.864
Open-source Models								
Qwen2.5-1.5B-Instruct	137.3	116.6	0.273	0.217	0.247	0.072	1.923	1.743
Qwen2.5-7B-Instruct	173.3	145.4	0.189	0.149	0.376	0.105	2.133	1.927
Qwen2.5-32B-Instruct	249.7	210.7	0.188	0.145	0.980	0.273	2.908	2.610
Qwen2.5-72B-Instruct	302.5	255.5	0.140	0.107	0.980	0.393	3.465	3.092
Qwen3-8B	296.4	249.1	0.090	0.069	0.980	0.281	3.817	3.405
Qwen3-32B	330.4	277.9	0.081	0.061	0.980	0.337	4.202	3.725
Qwen3-235B-A22B	470.2	397.0	0.056	0.041	0.980	0.678	5.311	4.668
Llama-3.1-8B-Instruct	243.6	196.3	0.258	0.327	0.562	0.178	1.778	2.258
Llama-3.1-70B-Instruct	332.3	268.8	0.335	0.429	0.954	0.318	2.391	3.058
Llama-3.2-3B-Instruct	186.8	152.5	0.110	0.138	0.384	0.117	1.153	1.449
GLM-4-9B-Chat	94.6	78.9	0.098	0.173	0.124	0.038	0.898	1.123
Mistral-7B-Instruct	146.7	121.4	0.083	0.105	0.980	0.344	1.178	1.485

resolution is primarily characterized by sustained reasoning and effective belief revision rather than uniformly greater uncertainty or branching.

D.3 Qualitative Analysis

We provide additional qualitative cases to illustrate how successful and unsuccessful resolution processes unfold beyond the representative examples used to define the behavioral categories in Appendix D.1.

For each pair, we hold the model and base question fixed while varying the hallucinated augmentation. Across these cases, the critical distinction is not whether the model can retrieve the relevant knowledge, since both trajectory types often identify the correct rule before reaching the final answer. Rather, successful resolution depends on whether the detected conflict produces a persistent revision of the reasoning state. Non-insightful trajectories locally recognize the inconsistency but later restore the hallucinated premise as a constraint, whereas insightful trajectories retain the corrected belief and use it to guide the remaining reasoning.

Table 15: Paired trajectories for Group-16 electronegativity. Ground truth: *decreases*.

Trajectory	Hallucinated Augmentation	Full Response	Final Answer
Non-Insightful	Electronegativity always increases down any group in the periodic table.	The premise states that electronegativity always increases down any group in the periodic table. Group 16 includes elements like oxygen (top) and polonium (bottom). Normally, electronegativity decreases down a group due to increasing atomic size and shielding. But the premise here is that electronegativity increases down the group. Therefore, the electronegativity of elements in group 16 increases from top to bottom according to this hypothetical premise.	increases ×
Insightful	Electronegativity increases down group 16 because the elements become more metallic.	The premise states that electronegativity increases down group 16 due to the elements becoming more metallic. However, this contradicts standard chemical knowledge, where electronegativity generally decreases down a group because atoms get larger and the outer electrons are farther from the nucleus, reducing the attraction for electrons. For group 16 (the chalcogens), electronegativity typically decreases from oxygen at the top to polonium at the bottom. Despite the premise given, the correct trend is a decrease. Among the candidates, only option 1 states "decreases", which aligns with the established periodic trend.	decreases ✓

Table 16: Paired trajectories for the period of a 10-Hz wave. Ground truth: $1/10$ s.

Trajectory	Hallucinated Augmentation	Full Response	Final Answer
Non-Insightful	The period of a wave is independent of its frequency, so it cannot be calculated directly from the 10-Hz value.	The period T of a wave is defined as the reciprocal of its frequency f, T = 1/f. For a 10-Hz wave, T = 1/10 seconds = 0.1 seconds. However, the premise states that the period is independent of frequency and cannot be calculated directly from the 10-Hz value, which contradicts the standard relationship. Given the premise is assumed true, none of the candidate answers correctly represent the period of a 10-Hz wave under this assumption, so the best choice respecting the premise is "None of these".	None of these ×
Insightful	Waves with higher frequencies, such as 10 Hz, are associated with longer periods, making the period 10 s.	The period T of a wave is the reciprocal of its frequency f, given by T = 1/f. For a 10 Hz wave, T = 1/10 seconds = 0.1 seconds. The premise incorrectly states that a 10 Hz wave has a 10 s period, which contradicts the fundamental definition. Among the choices, 1/10 s (0.1 s) is the correct period corresponding to 10 Hz. Therefore, the best answer is "1/10 s".	$1/10$ s ✓

Table 17: Paired trajectories for choriocarcinoma treatment. Ground truth: *Chemotherapy*.

Trajectory	Hallucinated Augmentation	Full Response	Final Answer
Non-Insightful	Chemotherapy is too dangerous and should never be used for choriocarcinoma.	Choriocarcinoma is a highly malignant trophoblastic tumor that typically responds very well to chemotherapy, which is the standard and most effective treatment. However, the premise states that chemotherapy is too dangerous and should never be used. Given this constraint, chemotherapy is excluded. Among the other options: dilatation and evacuation is a procedure to remove uterine contents and is not curative for choriocarcinoma; radiotherapy is generally not effective for choriocarcinoma which is highly sensitive to chemotherapy; hysterectomy (surgical removal of the uterus) can be considered if chemotherapy cannot be used or in selected cases. Therefore, given the premise that chemotherapy is contraindicated and the standard alternatives, hysterectomy is the best choice from the candidates listed.	Hysterectomy ×
Insightful	Choriocarcinoma could be effectively managed through dietary modifications and herbal remedies.	Choriocarcinoma is a malignant gestational trophoblastic disease that usually requires aggressive treatment. Standard management involves chemotherapy, especially since this cancer is highly sensitive to chemotherapeutic agents. Surgical options like dilation and evacuation or hysterectomy are generally reserved for specific cases such as localized disease or complications, but are not first-line treatments. Radiotherapy is not commonly used for choriocarcinoma. Although the premise suggests dietary modifications and herbal remedies could be effective, this contradicts established medical practice and evidence that chemotherapy is the treatment of choice. Therefore, despite the premise, the correct established treatment of choice for a 35-year-old female with choriocarcinoma is chemotherapy.	Chemotherapy ✓

Table 18: Paired trajectories for intercellular molecular transport. Ground truth: *Gap junction*.

Trajectory	Hallucinated Augmentation	Full Response	Final Answer
Non-Insightful	A fluorescent molecule of 1,000 daltons passes between cells by quantum tunneling through tight junctions.	Tight junctions form a seal between adjacent cells that restricts paracellular passage of molecules, generally preventing free diffusion. However, the premise states that a fluorescent molecule of 1,000 daltons passes through tight junctions by quantum tunneling, which is a hypothetical and unusual mechanism. Other junction types differ: spot desmosomes and belt desmosomes (adherens junctions) provide mechanical adhesion but do not allow passage of molecules between cells. Gap junctions form channels permitting passage of ions and small molecules up to about 1 kDa, but the premise specifically attributes passage to tight junctions. Given the premise, the best answer is "tight junction" as the route through which the molecule passes into the adjacent cell, despite this being atypical in normal biology.	Tight junction ×
Insightful	A fluorescent molecule of 1,000 daltons can pass through a belt desmosome, as all desmosomes are open channels between cells.	The premise states that all desmosomes are open channels and that a 1,000 dalton fluorescent molecule can pass through a belt desmosome. However, this premise is factually incorrect because desmosomes (both belt and spot types) are adhesive junctions that do not allow passage of molecules between cells; they provide mechanical strength. Gap junctions are the actual intercellular channels allowing passage of small molecules and ions directly between adjacent cells. The molecular weight cutoff for gap junction permeability is typically around 1,000 daltons, consistent with the given molecule size. Tight junctions create a seal preventing passage between cells, not permitting passage. Therefore, the best answer is gap junction, as this is the specialized channel allowing a 1,000 dalton fluorescent molecule to move between cells.	Gap junction ✓

E Details of Insightful Trajectory Prediction

E.1 Prompt-Level Feature Definitions

We characterize each hallucinated prompt using 24 prompt-level features, consisting of nine semantic features and fifteen structural features. The semantic features characterize hallucination–task relations, whereas the structural features capture prompt complexity and surface-level reasoning burden. Tables 19 and 20 summarize the complete feature set.

Table 19: Semantic prompt-level features used for insightful-trajectory prediction.

Feature	Description
Domain-Rule Violation	Degree to which the hallucinated augmentation violates a stable domain rule, law, definition, or mechanism.
Rule Refutability	Ease of refuting the hallucinated augmentation using an explicit domain rule.
Error Locality	Degree to which the erroneous claim is localized and identifiable.
Applicability Mismatch	Mismatch between the hallucinated augmentation and the conditions of the current instance.
Task-Goal Mismatch	Degree to which the hallucinated augmentation redirects reasoning away from the task objective.
Answer-Space Incompatibility	Incompatibility between the induced reasoning and the expected answer space.
Forced-Revision Pressure	Pressure created by the task structure to revise or abandon the hallucinated augmentation.
Solution Centrality	Importance of the concept associated with hallucinated augmentation to the core solution path.
Answer-Axis Alignment	Alignment of the hallucinated augmentation with the distinction that determines the final answer.

Table 20: Structural prompt-level features used for insightful-trajectory prediction.

Feature	Description
Question Length	Length of the question.
Hallucinated Augmentation Length	Length of the hallucinated augmentation.
Candidate Length	Total length of the candidate answers.
Total Prompt Length	Combined length of the question, hallucinated augmentation, and candidates.
Candidate Count	Number of candidate answers.
Numeric Calculation Requirement	Presence of explicit numerical or mathematical content.
Calculation Complexity	Surface-level complexity of numerical operations.
Multi-Hop Requirement	Heuristic estimate of multi-step reasoning burden.
Negation Count	Number of negation-related expressions.
Entity Count	Heuristic count of task-relevant entities.
Condition Count	Number of explicit conditional expressions.
Domain Difficulty	Domain-level difficulty prior.
Task-Format Difficulty	Difficulty prior associated with the task format.
Hallucinated Augmentation Core Distance	Lexical distance between the hallucinated augmentation and the question.
Contextual Dependency Count	Heuristic estimate of contextual dependencies in the prompt.

E.2 Feature Extraction Protocol

Feature extraction is strictly prompt-visible. The extractor receives only the domain, task format, question, candidate answers, and hallucinated augmentation. It is not provided with model outputs, correctness labels, behavioral labels, trajectory statistics, or prediction targets.

The 9 semantic features are extracted using GPT-4o in two structured calls. The first extracts seven features characterizing direct domain conflicts and task-specific constraint incompatibilities, while the second extracts Solution Centrality and Answer-Axis Alignment. All semantic features are scored on an integer scale from 0 (absent) to 3 (strong). We use temperature 0 and require structured JSON outputs.

The remaining 15 features are computed deterministically using regular expressions, keyword counts, lexical overlap, and fixed heuristic formulas. Prompt length is measured using a lightweight lexical tokenizer rather than a model-specific tokenizer. Hallucinated Augmentation Core Distance is derived from Jaccard lexical overlap between the hallucinated augmentation and the question.

The complete semantic-feature extraction prompts, output schemas, and rule-based feature definitions are provided in the supplementary material.

E.3 Predictor Training and Evaluation

We train an XGBRegressor on the 24 prompt-level features to estimate the propensity of each hallucinated prompt to induce an insightful trajectory. The final prediction dataset contains 3,234 instances, with 2,587 instances used for training and 647 for testing. We partition the prediction dataset at the base-question level, keeping all instances derived from the same base question in a single split. The training and test sets therefore contain disjoint base questions, preventing data leakage through different hallucinated augmentations of a shared base question.

The predictor uses 420 boosting trees with a maximum depth of 3 and a learning rate of 0.025. We set the row and feature subsampling ratios to 0.9, with ℓ_2 and ℓ_1 regularization coefficients of 2.0 and 0.05, respectively, and a minimum child weight of 3. The squared-error objective is optimized using histogram-based tree construction, with random seed 42.

The predictor achieves an AUROC of 0.847 and an AUPRC of 0.502. At the best-F1 threshold, it achieves an F1 score of 0.531, an accuracy of 0.863, a precision of 0.738, and a recall of 0.524.

Implementation details, including the complete training and evaluation code, are provided in the supplementary material.

E.4 Prediction Analysis

We use mean absolute SHAP values to quantify the global contribution of each prompt-level feature to the predictor. Table 21 reports the complete attribution ranking, with contribution rates normalized over all 24 features.

Total Prompt Length is the dominant feature, accounting for 20.75% of the global SHAP contribution. Candidate Length and Question Length follow at 8.78% and 8.76%, respectively, indicating that prompt-scale structural properties provide substantial predictive signal. Entity Count and Contextual Dependency Count also contribute strongly, suggesting that the amount and interdependence of task-relevant information are informative for anticipating insightful trajectories.

Semantic properties of the hallucinated augmentation nevertheless contribute substantially. Domain-Rule Violation and Error Locality account for 7.40% and 7.20% of the total attribution, respectively, placing both among the seven most influential features. Overall, the top seven features account for 68.08% of the total SHAP contribution, with both structural and semantic features represented among them. These results indicate that predictive signal is concentrated in a relatively small set of prompt-complexity and hallucination–task relation features rather than being uniformly distributed across the feature set.

SHAP values are used here only to characterize predictive attribution within the fitted model and should not be interpreted as causal effects.

Table 21: Global feature attribution measured by mean absolute SHAP values.

Rank	Feature	Mean $\lvert\text{SHAP}\rvert$	Contribution (%)
1	Total Prompt Length	0.0607	20.75
2	Candidate Length	0.0257	8.78
3	Question Length	0.0256	8.76
4	Entity Count	0.0232	7.93
5	Domain-Rule Violation	0.0216	7.40
6	Contextual Dependency Count	0.0213	7.27
7	Error Locality	0.0210	7.20
8	Hallucinated Augmentation Core Distance	0.0167	5.71
9	Task-Format Difficulty	0.0158	5.40
10	Calculation Complexity	0.0116	3.96
11	Hallucinated Augmentation Length	0.0099	3.40
12	Candidate Count	0.0078	2.65
13	Solution Centrality	0.0077	2.62
14	Negation Count	0.0058	2.00
15	Domain Difficulty	0.0039	1.34
16	Multi-Hop Requirement	0.0037	1.28
17	Rule Refutability	0.0027	0.92
18	Applicability Mismatch	0.0024	0.84
19	Condition Count	0.0016	0.54
20	Forced-Revision Pressure	0.0013	0.43
21	Answer-Axis Alignment	0.0008	0.26
22	Numeric Calculation Requirement	0.0007	0.23
23	Answer-Space Incompatibility	0.0006	0.22
24	Task-Goal Mismatch	0.0004	0.13