

# Tracking the Limits of Knowledge Propagation: How LLMs Fail at Multi-Step Reasoning with Conflicting Knowledge

Yiyang Feng<sup>✦✦</sup>

yiyang.feng@stonybrook.edu

Zeming Chen<sup>♣</sup>

zeming.chen@epfl.ch

Haotian Wu<sup>♣</sup>

haotian.wu@epfl.ch

Jiawei Zhou<sup>✦</sup>

jiawei.zhou.1@stonybrook.edu

Antoine Bosselut<sup>♣</sup>

antoine.bosselut@epfl.ch

♣ EPFL ✦ Stony Brook University

## Abstract

A common solution for mitigating outdated or incorrect information in Large Language Models (LLMs) is to provide updated facts in-context or through knowledge editing. However, these methods introduce knowledge conflicts when the knowledge update fails to overwrite the model’s parametric knowledge, which propagate to faulty reasoning. Current benchmarks for this problem, however, largely focus only on single knowledge updates and fact recall without evaluating how these updates affect downstream reasoning. In this work, we introduce TRACK (*Testing Reasoning Amid Conflicting Knowledge*), a new benchmark for studying how LLMs propagate new knowledge through multi-step reasoning when it conflicts with the model’s initial parametric knowledge. Spanning three reasoning-intensive scenarios (WIKI, CODE, and MATH), TRACK introduces multiple, realistic conflicts to mirror real-world complexity. Our results on TRACK reveal that providing updated facts to models for reasoning can worsen performance compared to providing no updated facts to a model, and that this performance degradation exacerbates as more updated facts are provided. We show this failure stems from both inability to faithfully integrate updated facts, but also flawed reasoning even when knowledge is integrated. TRACK provides a rigorous new benchmark to measure and guide future progress on propagating conflicting knowledge in multi-step reasoning.<sup>1</sup>

## 1 Introduction

Large Language Models (LLMs) demonstrate remarkable reasoning capabilities (Sprague et al., 2025), leading to their widespread use in scenarios such as multi-hop question answering (Zhu et al., 2024), coding (Jiang et al., 2024), and mathematics (Ahn et al., 2024; Poiroux et al., 2025). However, much of these abilities come from their

<sup>1</sup>Code and data are available at <https://github.com/Wind-2375-like/crack>.

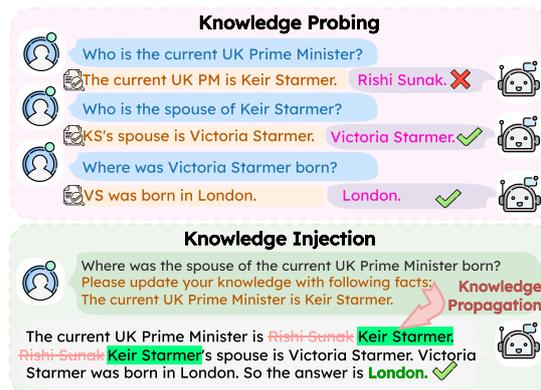


Figure 1: An illustration of TRACK’s two-stage evaluation framework using a multi-hop QA example. (i) Knowledge Probing: We first decompose a complex question into atomic facts to identify a model’s specific knowledge gaps. (ii) Knowledge Injection: We then provide correct facts as conflicting knowledge to test LLM’s ability to propagate it within its reasoning.

parametric knowledge learned during pretraining, which can become outdated and inaccurate. For example, an LLM might hold outdated information about a current head of state (Zhong et al., 2023), or may generate non-executable codes using deprecated function signatures from older version libraries (Liu et al., 2025), or misremember a specific mathematical theorem (Singh et al., 2024).

To mitigate this issue, a common practice is to provide updated factual knowledge, either through in-context learning (Longpre et al., 2021) or by directly editing model parameters (De Cao et al., 2021; Mitchell et al., 2022a). However, because the provided new knowledge introduces knowledge conflicts that contradict the model’s internal, parametric beliefs (Xu et al., 2024b), models fail to overwrite their parametric knowledge under such conflicts, which challenges the reliability of multi-step reasoning. Investigating *how effectively LLMs propagate new knowledge through multi-step reasoning amidst such conflicts* offers fundamental insights into how LLMs integrate knowledge and

perform reasoning.

Unfortunately, existing benchmarks are insufficient for evaluating knowledge propagation in multi-step reasoning. First, they largely focus on simple fact recall (Mitchell et al., 2022b), rather than assessing whether a model can faithfully leverage new knowledge in multi-step reasoning. Second, they typically evaluate single knowledge updates in isolated scenarios (Xie et al., 2024), which oversimplifies the real-world challenge of managing multiple, interacting conflicts. Finally, many of these benchmarks rely on counterfactual data (e.g., Eiffel Tower being in Rome), which fails to reflect the complexity of real-world knowledge evolution (Liu et al., 2025), and can also contain unintentional factual errors (Zhong et al., 2025).

To address these limitations, we introduce TRACK (*Testing Reasoning Amid Conflicting Knowledge*), a new benchmark to evaluate how LLMs propagate conflicting knowledge through multi-step reasoning, illustrated in Fig. 1. The benchmark is defined by a novel evaluation framework consisting of two stages. The first *knowledge probing* stage identifies a model’s specific knowledge gaps for a given problem. Then, in the *knowledge injection* stage, we provide the updated facts as conflicting knowledge and evaluate if the model can successfully solve a downstream reasoning problem requiring this information. We compare performance between two settings: a closed-book setting where the model relies on its internal knowledge, and an open-book setting where the new facts are provided. Our framework moves beyond simple fact recall and explicitly tests knowledge propagation because each injected fact is a necessary piece of intermediate knowledge, requiring the model to perform further reasoning steps to arrive at the final answer.

We instantiate the TRACK benchmark with three datasets constructed from realistic conflicts in challenging scenarios: (i) multi-hop QA on recent Wiki-data (WIKI); (ii) code generation with external APIs (CODE); and (iii) multi-step mathematical reasoning (MATH). To evaluate reasoning quality on TRACK, we introduce novel metrics: *Full Knowledge Entailment* (FKE), which measures if the reasoning is faithful to all facts required for the solution, and *Holistic Pass* (HP), a strict metric requiring both a correct answer and a faithful reasoning chain.

Our experiments across a wide range of LLMs, including open-source and closed-source, and

thinking and non-thinking, models reveal that current models struggle significantly on TRACK. Specifically, we find that: (i) providing models with correct, conflicting facts (the open-book setting) yields surprisingly limited gains and can even backfire, with performance sometimes falling below the closed-book baseline; (ii) performance degrades as more conflicting facts are provided; and (iii) this failure stems from both an inability to faithfully integrate the new facts, but also flawed reasoning even when knowledge integration is successful.

Our contributions are threefold. First, we formalize a novel evaluation framework for rigorously measuring reasoning under knowledge conflicts, centered on a two-stage methodology of knowledge probing and injection. Second, we introduce TRACK, a new benchmark that instantiates this framework with three diverse datasets (WIKI, CODE, MATH) featuring realistic, multi-fact conflicts, and novel evaluation metrics (FKE, HP) for a more fine-grained analysis. Finally, we conduct a comprehensive experimental analysis on a wide range of LLMs, revealing a critical failure where providing correct conflicting facts shows limited gains and even degrades performance, rooted in both unfaithful knowledge integration and flawed reasoning.

## 2 Related Work

**Knowledge Conflict.** We define knowledge conflicts as discrepancies between internal knowledge and external updates (context or edits), distinct from inter-context or intra-memory conflicts (Xu et al., 2024b). Since Longpre et al. (2021) introduced this domain, benchmarks have analyzed preferences between internal and external knowledge (Xie et al., 2024; Wu et al., 2024), misinformation susceptibility (Xu et al., 2024a), and conflict resolution (Wang et al., 2024), yet generally neglect knowledge propagation, the use of new facts in multi-step reasoning. While recent studies explore propagation (Feng et al., 2025) and real-world conflicts (Kortukov et al., 2024), they remain limited in counterfactual data and reasoning variability. TRACK addresses these gaps by employing multiple realistic conflicts from various scenarios to evaluate complex knowledge propagation.

**Knowledge Editing.** Knowledge editing aims to efficiently update a model’s parametric knowledge (Yao et al., 2023). A wide range of methods have been proposed, from early hypernetwork-

Benchmark	Knowledge Propagation	Multiple Conflicts	Reasoning Scenarios			Real-World Knowledge	Size
			Multi-Hop QA	Coding	Math		
<i>Knowledge Conflict Benchmarks</i>							
(Longpre et al., 2021)	□	□	■	□	□	□	315,203
(Xie et al., 2024)	□	□	■	□	□	□	16,557
(Kortukov et al., 2024)	■	■	■	□	□	■	58,281
(Wang et al., 2024)	□	□	■	□	□	□	9,083
(Xu et al., 2024a)	□	□	■	□	□	□	1,500
(Wu et al., 2024)	□	□	■	□	□	■	1,278
(Ying et al., 2024)	□	■	□	□	□	■	11,684
(Feng et al., 2025)	■	■	□	□	■	□	400
<i>Knowledge Editing Benchmarks</i>							
(Mitchell et al., 2022b)	□	■	■	□	□	■	15,000
(Zhong et al., 2023)	■	■	■	□	□	■	11,086
(Onoe et al., 2023)	■	□	■	□	□	■	1,000
(Cohen et al., 2024)	■	□	■	□	□	□	5,000
(Hua et al., 2024)	■	□	■	□	□	□	5,010
(Liu et al., 2025)	■	□	□	■	□	□	670
(Huang et al., 2025)	■	□	■	□	□	■	40,195
(Rosati et al., 2024)	■	□	■	□	□	□	7,389
(Thede et al., 2025)	■	□	■	□	□	■	502,382
<i>Our Benchmark: Testing Reasoning Amid Conflicting Knowledge (TRACK)</i>							
TRACK (ours)	■	■	■	■	■	■	1,500

Table 1: **Comparison between TRACK and prior knowledge conflict and editing benchmarks.** Our benchmark, TRACK, (i) requires true knowledge propagation<sup>2</sup> to achieve multi-step reasoning; (ii) presents multiple conflicts and spans multiple scenarios (multi-hop QA, coding, and math); and (iii) uses conflicts derived from real-world knowledge rather than counterfactuals. Symbols denote full (■), partial (▣), or no (□) support for a feature.

based approaches (De Cao et al., 2021; Hase et al., 2021) and influential locate-then-edit techniques like ROME/MEMIT (Meng et al., 2023a,b), to demonstrating that standard fine-tuning can be a strong baseline (Gangadhar and Stratos, 2024). However, the conceptual foundations of this field are heavily debated. Studies have questioned the core assumptions of localization (Hase et al., 2023), pointed out fundamental problems with the belief revision paradigm itself (Hase et al., 2024), and highlighted unintended side effects, such as the amplification of social biases (Halevy et al., 2024). Consequently, the evaluation of editing has also matured, with new benchmarks assessing multi-hop reasoning (Zhong et al., 2023; Onoe et al., 2023; Hua et al., 2024), coding (Liu et al., 2025), performance on long-form generation (Rosati et al., 2024), the correction of verified real-world hallucinations (Huang et al., 2025), and large-scale lifelong updates (Thede et al., 2025). However, existing work is often limited to single reasoning scenarios or conflict types. In contrast, TRACK provides the first unified benchmark to test propagation over multiple, realistic conflicts across three diverse domains. Tab. 1 presents a comprehensive comparison of these existing datasets.

### 3 Testing Reasoning Amid Conflicting Knowledge (TRACK)

In this section, we introduce our benchmark, *Testing Reasoning Amid Conflicting Knowledge*

(TRACK), to evaluate how LLMs perform multi-step reasoning when faced with knowledge conflicts. Specifically, we outline our problem formulation (Sec. 3.1), three instantiated scenarios (Sec. 3.2), and evaluation framework (Sec. 3.3). The overall pipeline of our benchmark is illustrated in Fig. 2.

#### 3.1 Problem Setup

We investigate how an LLM  $f$  reasons when newly introduced facts contradict its parametric knowledge. The benchmark comprises two stages: (i) knowledge probing to identify the model’s knowledge gaps, and (ii) knowledge injection to integrate new relevant facts into reasoning.

**Knowledge Probing.** Given a question  $q$ , we first decompose it into a set of atomic facts,  $K_q = \{k_1, \dots, k_n\}$ , required to answer the question. To identify the model’s knowledge gap ( $K_g$ ), we convert each fact  $k_i \in K_q$  into a probing question-answer pair  $(q_i, a_i)$  and query the model. The knowledge gap is then defined as the subset of facts unknown by a model (detailed in Sec. 3.3).  $K_g$  pinpoints exact locations where providing correct external information will induce a knowledge conflict, making it necessary for the subsequent injection stage.

<sup>2</sup>We define knowledge propagation as using a new fact as a premise in a reasoning chain, not simply recalling a fact that is the final answer. For instance, the conflicting facts in Ying et al. (2024) are final answers (see their Fig. 2), thus testing recall rather than propagation.

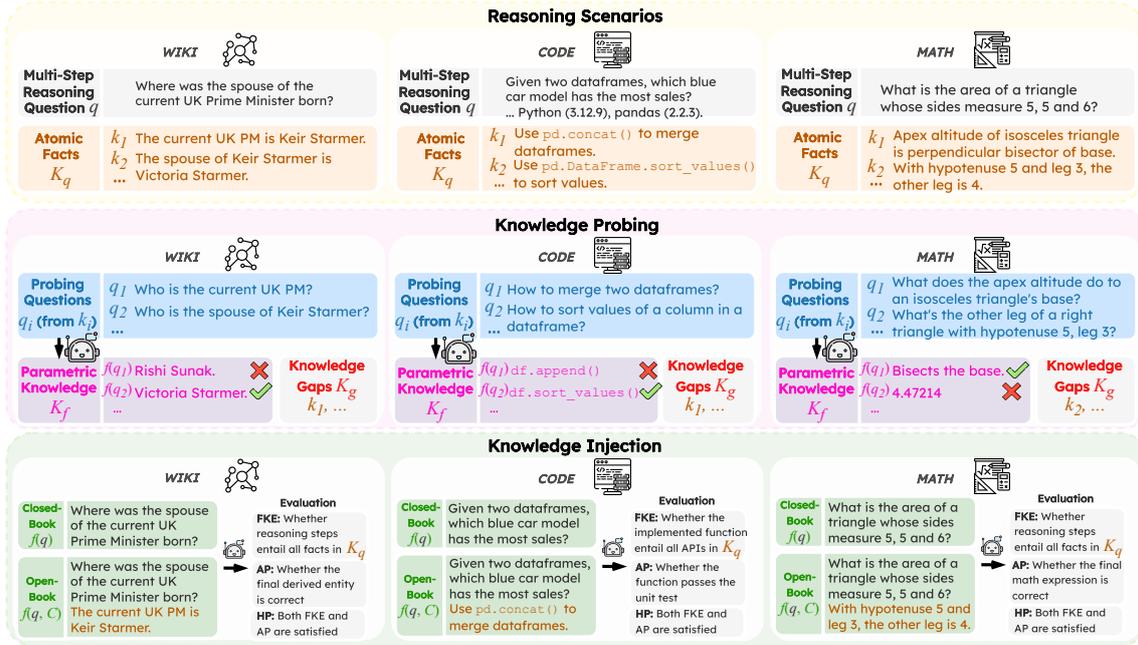


Figure 2: **Full illustration of the TRACK benchmark across our three reasoning scenarios:** Multi-Hop QA (WIKI), Code Generation (CODE), and Mathematical Reasoning (MATH). The benchmark follows a two-stage process. (i) Knowledge Probing: we identify the model’s knowledge gaps by testing it on the required atomic facts, the individual pieces of knowledge needed to solve the complex question. (ii) Knowledge Injection: we evaluate reasoning by comparing a closed-book setting (using only the model’s internal knowledge) with an open-book setting (where the identified knowledge gaps are provided). Performance is assessed using our metrics: Answer Pass (AP), Full Knowledge Entailment (FKE), and Holistic Pass (HP).

**Knowledge Injection.** This stage evaluates how the model  $f$  integrates facts from the knowledge gap  $K_g$ , generating a final answer  $\hat{a}$  and reasoning  $\hat{R}$ . We compare a *closed-book* baseline where the model only receives the question  $q$ , denoted  $(\hat{R}_p, \hat{a}_p) = f(q)$ ; and an *open-book* setting where the question is accompanied by a set of new information  $C$  derived from  $K_g$ , denoted  $(\hat{R}_c, \hat{a}_c) = f(q, C)$ . To simulate shared update environment where models must discern relevant facts from concurrent inputs, the *knowledge aggregation scope* (KAS) parameter controls how this set of information  $C$  is constructed. We aggregate knowledge gaps from a batch  $\mathcal{B}$  of KAS questions to construct the set:  $C_{\mathcal{B}} = \bigcup_{q \in \mathcal{B}} K_g$ . For example, A KAS of 1 is instance-specific (each question receives only its own missing facts), while a KAS of 10 requires the model to reason using the combined knowledge from 10 different questions. We retain samples without knowledge gaps to ensure a fixed test set and avoid inductive bias (Si et al., 2023). As shown in App. D, excluding them has negligible impact on scores and leaves conclusions unchanged.

### 3.2 TRACK Scenarios

We instantiate the TRACK benchmark across three diverse reasoning scenarios (WIKI, CODE,

MATH).

**WIKI: Wikipedia Multi-Hop QA.** A complex question  $q$  like “Where was the spouse of the current UK Prime Minister born?” requires reasoning over a chain of relational facts from Wikipedia (Vrandečić and Krötzsch, 2014). Each relation (e.g., “The current UK Prime Minister is Keir Starmer.”) serves as an atomic fact  $k_i$ . A conflict arises from the model’s outdated world knowledge (e.g., believing Rishi Sunak is the PM).

**CODE: Code Generation with External APIs.** Given a coding problem  $q$  like “Given two dataframes, find which blue car model has the most sales? ... Python (3.12.9), pandas (2.2.3).”, the model must generate a correct function. Each atomic fact  $k_i$  is the proper usage of an external API. Outdated API knowledge can introduce conflicts, such as using a deprecated function (e.g., `df.append`) instead of its modern equivalent (`pd.concat`).

**MATH: Multi-Step Mathematical Reasoning.** For a standard math problem  $q$  like “What is the area of a triangle whose sides measure 5, 5, and 6?”, each atomic fact  $k_i$  is a procedural step (“What’s the other leg of a right triangle with hypotenuse 5 and leg 3?”). A conflict arises when the model’s in-

ternal parametric knowledge contains an incorrect or hallucinated procedural step.

### 3.3 Evaluation Metrics

We evaluate each of the two stages of the TRACK benchmark. Details are in App. A.

**Evaluate Knowledge Probing.** For each atomic fact, we sample 10 responses to its probing question. Using an LLM judge (GPT-5-mini),<sup>3</sup> we count the number of responses equivalent to the ground-truth answer. A fact is considered *known* if its correct answer is the most frequent response; otherwise, it is *unknown*. We define *knowledge confidence* (KConf) as the proportion of correct responses.

**Evaluate Knowledge Injection.** We assess the model’s reasoning output using three metrics: (i) *Answer Pass* (AP) is a binary metric for the final answer’s correctness. Answer correctness check is scenario-specific: an equivalence check for WIKI and MATH using GPT-5-mini, and Pass@1 (Chen et al., 2021) for CODE, (ii) *Full Knowledge Entailment* (FKE) is a binary metric checking if the reasoning is faithful to all required atomic facts. FKE holds if the reasoning chain entails every fact in the required set. (iii) *Holistic Pass* (HP) is a stricter metric requiring both correctness (Answer Pass) and faithfulness (Full Knowledge Entailment).

## 4 TRACK Benchmark Datasets

The TRACK benchmark comprises 1,500 examples, with 500 for each of our three reasoning scenarios (WIKI, CODE, MATH). We constructed these scenarios by sourcing and adapting data from Wikidata (Vrandečić and Krötzsch, 2014), BigCodeBench (Zhuo et al., 2025), and PRM800K (Lightman et al., 2024), respectively. Each example contains five key components (Fig. 5): a reasoning question ( $q$ ), its required atomic facts ( $K_q$ ), corresponding probing question-answer pairs ( $q_i, a_i$ ), and the final answer ( $a$ ). The detailed data generation pipeline, which combines heuristic rules and the use of GPT-5-mini, is detailed in App. B. Note that our pipeline supports regenerating the benchmark with future knowledge to ensure its continued relevance for evaluating new models. For example, researchers can run the pipeline on new Wikipedia dumps to generate recent WIKI questions, use it to parse up-to-date

<sup>3</sup>We use gpt-5-mini-0807 for all following GPT-5-mini models.

Statistic	WIKI	CODE	MATH
# Multi-Step Questions ( $q$ )	500	500	500
Avg. # Atomic Facts ( $k_i$ )	4.00	3.82	5.16
Avg. Tokens in Question ( $q$ )	38.02	198.68	97.02
Avg. Tokens in Probing Q ( $q_i$ )	19.80	46.88	28.12
Avg. Tokens in Atomic Fact ( $k_i$ )	12.32	26.17	39.62

Table 2: Key statistics for the constructed datasets.

code pieces for new CODE challenges, and source MATH problems from recent competitions.

We install rigorous quality controls to ensure the factuality of probing question-answer pairs and the necessity of atomic facts. To enforce factuality, we filter out all non-factual probing pairs and corresponding atomic facts using GPT-5-mini. Additionally, WIKI questions are time-stamped (e.g., “By Sep. 2025”), and CODE problems specify library versions (e.g., “Python (3.12.9), pandas (2.2.3)”). For necessity, we compose WIKI questions directly from the underlying fact chains, while CODE and MATH problems explicitly instruct the model to use the functions or procedures derived from the provided atomic facts. We further validated our data by checking whether each probing question-answer pair ( $q_i, a_i$ ) is factual and each fact in atomic facts  $K_q$  is necessary in the multi-step reasoning question  $q$ . Using both gpt-5-mini and gemini-2.5-pro<sup>4</sup> as judges, two models annotate 94.3% as factual and 88.0% as necessary on average, and their annotations have high average F1 scores with one human annotator (98.0% for factuality and 93.5% for necessity). Details of the quality check are in App. E.2.

While automated pipelines guarantee factuality, we ensure high reasoning variability both across and within domains. Cross-domain structures span graph traversal (WIKI), procedural composition (CODE), and symbolic derivation (MATH), while within-domain diversity arises from adapting sources like BigCodeBench and PRM800K to create varied reasoning chains. As highlighted by the statistics in Tab. 2, this diversity in reasoning complexity and question length ensures TRACK provides a comprehensive evaluation of model capabilities across various scenarios.

## 5 Knowledge Probing

In the knowledge probing stage, we identify each model’s internal knowledge gaps. These identified gaps, which we release as a new diagnostic

<sup>4</sup>Released on June 17, 2025

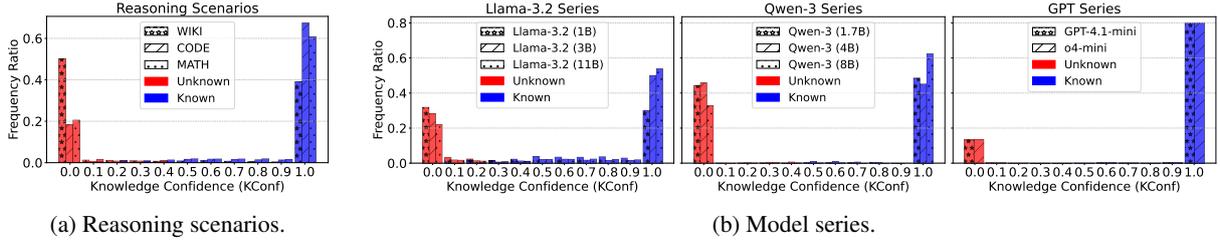


Figure 3: **Knowledge Confidence (KConf) distributions** of Known and Unknown facts across (a) reasoning scenarios, and (b) model series. A consistent color scheme is used across both panels (Blue: Known facts, Red: Unknown facts).

dataset, are necessary for preparing model-specific knowledge conflicts evaluated in the knowledge injection stage. This section outlines our probing setup (Sec. 5.1) and analyzes the results (Sec. 5.2).

## 5.1 Experimental Setup

**Backbone LLMs.** We select a wide range of LLMs, including both open-source and closed-source, thinking and non-thinking models, across various scales. Specifically, we test Llama-3.2 (1B, 3B, 11B)<sup>5</sup> (Grattafiori et al., 2024), Qwen-3 (1.7B, 4B, 8B) (Yang et al., 2025), GPT-4.1-mini<sup>6</sup> (OpenAI, 2025a), and o4-mini<sup>7</sup> (OpenAI, 2025b).

**Knowledge Probing.** To identify each model’s knowledge gaps, we probe it with questions based on atomic facts. For each question, we generate  $M = 10$  responses. Detailed prompts and model-specific hyperparameters are in Appendix C.1.

## 5.2 Knowledge Probing Results

**Models exhibit polarized KConf.** As shown in Fig. 3, models exhibit a highly polarized knowledge confidence (KConf). The KConf for known facts is sharply concentrated at 1, while the confidence for unknown facts is skewed towards 0. This polarization indicates that models are rarely uncertain: they either know a fact with high confidence or are clearly aware they do not know it at all.

**KConf distributions vary across scenarios and models.** The distribution of known versus unknown facts also varies significantly by both scenario (Fig. 3a) and model (Fig. 3b). The WIKI scenario, built from up-to-date Wikidata (Sep. 2025), contains the most unknown facts likely due to many of its facts post-dating the models’ knowledge cut-offs.<sup>8</sup> In contrast, CODE and MATH contain more

known facts, as they are reconstructed from established benchmarks likely included in the models’ pre-training data. Knowledge also differs by model family and scale, with closed-source models (GPT series) and larger scales (e.g., Llama-3.2 11B v.s. 1B) generally possessing more known facts.

## 6 Knowledge Injection

In the knowledge injection stage, we evaluate how various models and methods perform when presented with knowledge conflicts. These knowledge conflicts are constructed from identified knowledge gaps in the knowledge probing stage. This section details the experimental setup (Sec. 6.1), presents the main results (Sec. 6.2), and provides an in-depth failure analysis (Sec. 6.3).

### 6.1 Experimental Setup

**Knowledge Injection.** Building on the same set of backbone LLMs (Sec. 5.1), we evaluate how they perform multi-step reasoning when injected with conflicting knowledge. We compare a closed-book Base Model baseline against a diverse set of open-book injection methods representing paradigms like in-context learning, test-time inference, fine-tuning, and retrieval-augmentation. These methods include Append,<sup>9</sup> which provides new facts in-context; Append-Thinking (Append-T), a variant that enables longer thinking for Qwen-3 and o4-mini; Fine-tuning on Conflicting Knowledge (FT-CK); and MeLLo (Zhong et al., 2023), which uses an external memory to store injected facts and a question-decomposition approach for retrieval and reasoning. For all open-book methods, we vary the number of KAS (Sec. 3.1) to 1, 10, 100, and 500. Full implementation specifics are in Appendix C.2.

<sup>9</sup>Early experiments showed that appending outperforms prepending facts in-context, potentially because the question provides context for the model to attend to relevant knowledge.

<sup>5</sup>The Llama-3.2 11B is a multimodal model.

<sup>6</sup>We use gpt-4.1-mini-2025-04-14.

<sup>7</sup>We use o4-mini-2025-04-16.

<sup>8</sup>Llama-3.2: Dec. 2023; Qwen-3: Sep. 2024; GPT-4.1-mini and o4-mini: Jun. 2024

Backbone Model	Method	WIKI			CODE			MATH		
		HP	AP	FKE	HP	AP	FKE	HP	AP	FKE
Llama-3.2 (1B)	Base Model	0.9 ± 0.7	6.1 ± 2.1	0.9 ± 0.7	5.7 ± 2.1	9.2 ± 2.6	36.4 ± 4.2	20.0 ± 3.4	26.0 ± 4.0	26.3 ± 3.7
	Append	1.1 ± 0.9	7.6 ± 2.4	1.0 ± 0.8	4.1 ± 1.7	7.0 ± 2.4	38.5 ± 4.3	22.5 ± 3.7	27.8 ± 3.8	28.6 ± 4.0
	FT-CK	0.9 ± 0.7	6.8 ± 2.2	0.9 ± 0.7	5.1 ± 1.9	8.5 ± 2.3	36.5 ± 4.3	21.8 ± 3.6	28.2 ± 3.8	28.1 ± 3.9
	MeLLO	2.4 ± 1.2	11.5 ± 2.7	2.4 ± 1.4	3.9 ± 1.7	5.4 ± 2.0	33.3 ± 4.3	18.0 ± 3.2	24.5 ± 3.7	24.6 ± 3.8
Llama-3.2 (3B)	Base Model	5.3 ± 1.9	23.1 ± 3.5	5.6 ± 2.0	15.3 ± 3.1	22.0 ± 3.4	54.7 ± 4.5	40.5 ± 4.1	52.9 ± 4.1	44.6 ± 4.2
	Append	7.3 ± 2.1	23.3 ± 3.7	7.8 ± 2.4	16.8 ± 3.2	23.3 ± 3.7	52.9 ± 4.5	40.4 ± 4.2	52.7 ± 4.3	43.9 ± 4.3
	FT-CK	7.4 ± 2.2	23.0 ± 3.6	8.2 ± 2.4	12.8 ± 2.8	20.1 ± 3.5	54.8 ± 4.4	41.6 ± 4.2	53.7 ± 4.3	46.1 ± 4.3
	MeLLO	7.3 ± 2.3	19.3 ± 3.3	8.2 ± 2.2	12.2 ± 2.8	17.4 ± 3.2	45.1 ± 4.7	37.9 ± 4.3	47.5 ± 4.7	40.5 ± 4.3
Llama-3.2 (11B)	Base Model	12.5 ± 2.9	32.6 ± 4.0	12.9 ± 2.9	18.7 ± 3.3	29.3 ± 3.9	50.5 ± 4.5	34.8 ± 4.0	40.4 ± 4.2	48.8 ± 4.4
	Append	13.1 ± 2.9	34.0 ± 4.2	13.2 ± 3.0	18.1 ± 3.3	27.1 ± 3.7	51.9 ± 4.5	33.4 ± 4.0	40.2 ± 4.4	46.5 ± 4.5
	FT-CK	12.7 ± 2.9	35.2 ± 4.4	12.7 ± 2.9	19.1 ± 3.5	28.1 ± 3.9	51.9 ± 4.5	32.6 ± 4.2	38.7 ± 4.3	49.6 ± 4.4
	MeLLO	11.0 ± 2.6	22.0 ± 3.6	12.4 ± 2.8	9.4 ± 2.6	17.1 ± 3.3	25.9 ± 3.9	20.3 ± 3.3	23.4 ± 4.0	36.3 ± 4.1
Qwen-3 (1.7B)	Base Model	4.5 ± 1.7	17.1 ± 3.1	4.5 ± 1.7	17.4 ± 3.2	24.3 ± 3.9	54.4 ± 4.2	48.3 ± 4.3	58.6 ± 4.2	64.0 ± 4.4
	Append	<b>83.6</b> ± 3.2	90.9 ± 2.5	<b>87.2</b> ± 3.0	18.3 ± 3.3	26.3 ± 3.7	59.6 ± 4.2	49.2 ± 4.0	57.6 ± 4.6	66.8 ± 4.0
	Append-T	6.4 ± 2.2	51.2 ± 4.2	6.2 ± 2.2	15.5 ± 3.1	23.1 ± 3.7	35.1 ± 4.1	7.1 ± 2.3	75.6 ± 3.8	19.3 ± 3.5
	FT-CK	4.8 ± 1.8	16.8 ± 3.2	4.7 ± 1.7	16.7 ± 3.3	22.4 ± 3.6	51.2 ± 4.2	45.6 ± 4.4	56.3 ± 4.3	60.9 ± 4.3
	MeLLO	5.5 ± 1.9	16.5 ± 3.3	5.8 ± 2.0	12.5 ± 2.9	18.6 ± 3.4	49.4 ± 4.4	27.8 ± 3.8	40.2 ± 4.4	31.8 ± 4.0
Qwen-3 (4B)	Base Model	7.7 ± 2.3	26.2 ± 3.8	8.1 ± 2.3	24.7 ± 3.7	35.8 ± 4.2	59.9 ± 4.3	63.0 ± 4.2	74.1 ± 3.9	73.5 ± 3.9
	Append	80.5 ± 3.5	<b>94.1</b> ± 2.1	82.6 ± 3.2	28.2 ± 3.8	34.7 ± 4.1	69.7 ± 3.9	62.3 ± 4.1	74.2 ± 3.8	72.9 ± 4.1
	Append-T	46.9 ± 4.3	63.4 ± 4.2	47.6 ± 4.4	22.8 ± 3.6	30.2 ± 3.8	42.0 ± 4.2	36.5 ± 4.1	87.5 ± 2.9	43.8 ± 4.2
	FT-CK	9.0 ± 2.4	30.3 ± 3.9	9.0 ± 2.4	25.6 ± 3.6	36.9 ± 4.3	60.6 ± 4.4	62.2 ± 4.4	74.7 ± 3.9	72.7 ± 3.9
	MeLLO	7.9 ± 2.3	29.4 ± 4.0	7.8 ± 2.2	21.4 ± 3.8	32.2 ± 4.0	55.2 ± 4.4	57.0 ± 4.4	74.4 ± 4.0	59.2 ± 4.4
Qwen-3 (8B)	Base Model	10.3 ± 2.5	26.8 ± 3.8	10.5 ± 2.5	27.2 ± 3.8	39.7 ± 4.3	63.1 ± 4.3	68.8 ± 4.0	85.0 ± 3.0	70.3 ± 3.9
	Append	77.7 ± 3.7	91.6 ± 2.4	78.9 ± 3.5	27.8 ± 3.8	40.6 ± 4.4	63.4 ± 4.2	70.7 ± 3.9	87.0 ± 3.0	71.8 ± 4.0
	Append-T	73.6 ± 3.8	93.7 ± 2.1	74.3 ± 3.7	26.5 ± 3.9	32.1 ± 4.1	43.8 ± 4.4	70.7 ± 3.9	82.2 ± 3.2	81.7 ± 3.5
	FT-CK	10.4 ± 2.6	28.4 ± 4.0	10.3 ± 2.5	28.5 ± 3.9	38.6 ± 4.4	65.5 ± 4.1	68.7 ± 3.9	84.6 ± 3.0	70.4 ± 4.2
	MeLLO	10.6 ± 2.6	27.8 ± 4.0	11.5 ± 2.9	24.8 ± 3.8	34.2 ± 4.2	59.7 ± 4.3	49.5 ± 4.3	65.3 ± 4.3	51.6 ± 4.2
GPT-4.1-mini & o4-mini	Base Model	22.9 ± 3.7	49.8 ± 4.6	22.6 ± 3.6	35.9 ± 4.1	48.0 ± 4.4	71.9 ± 4.1	79.3 ± 3.5	94.0 ± 2.0	80.9 ± 3.5
	Append	70.8 ± 4.0	93.7 ± 2.1	71.5 ± 3.9	36.8 ± 4.2	<b>49.2</b> ± 4.2	72.7 ± 3.9	86.4 ± 3.0	<b>98.7</b> ± 0.9	87.2 ± 2.8
	Append-T	79.6 ± 3.4	93.8 ± 2.2	79.9 ± 3.5	<b>40.6</b> ± 4.2	48.1 ± 4.1	<b>77.0</b> ± 3.6	<b>87.2</b> ± 2.8	<b>98.7</b> ± 0.9	<b>87.3</b> ± 2.9

Table 3: **Main results of knowledge injection methods** on the WIKI, CODE, and MATH scenarios. We report Holistic Pass (HP), Answer Pass (AP), and Full Knowledge Entailment (FKE) results. We compare the Base Model (closed-book) against several open-book knowledge injection methods (Append, FT-CK, MeLLO) across Llama-3.2, Qwen-3, and GPT series. We set KAS to 1 in the open-book setting so each question receives only its missing facts. We report 95% confidence intervals (CIs) in the  $\pm$  sign and **bold** the best scores per column.

## 6.2 Knowledge Injection Effectiveness

**Most methods show limited reasoning improvements or degraded performance.** We evaluate knowledge injection performance at KAS=1. As shown in Tab. 3, the poor performance of the Base Model in WIKI and CODE shows that reasoning struggles without the necessary atomic facts. In the open-book setting, the simple Append method yields the most substantial gains for Qwen-3 and GPT series on the WIKI scenario. For example, Qwen-3 1.7B’s HP score on WIKI jumps from 4.5% to 83.6%. However, for other scenarios (CODE and MATH), models (Llama-3.2 series, GPT series), and methods (FT-CK, MeLLO), the improvement is limited or even worse than the Base Model without injection. Notably, longer thinking has a limited improvement on o4-mini and even decreases Qwen-3 models’ performance.

**Some models and methods show performance degradation under increased knowledge scopes.**

We increase the KAS (Sec. 3.3), which aggregates unknown facts from multiple questions. As shown in Tab. 4, performance with the Append and Append-T methods degrade significantly for some models as KAS increases. For instance, on WIKI, the HP score for Qwen-3 1.7B Append method drops from 83.6% (KAS=1) to 22.2% (KAS=500), and for Qwen-3 8B Append-T method drops from 73.6% (KAS=1) to 11.2% (KAS=500). This suggests that as the provided information becomes saturated with more irrelevant facts, the model’s attention mechanism struggles to identify and use the pertinent knowledge. In contrast, some models or methods exhibit greater resilience: GPT-4.1-mini remains highly stable with slight performance drops for both methods, and MeLLO on Llama-3.2 and Qwen-3 can even improve with a larger KAS (e.g., on WIKI scenario). This indicates that longer reasoning or explicit retrieval is better at handling more distractor information.

Backbone Model	Method	WIKI				CODE				MATH			
		KAS=1	KAS=10	KAS=100	KAS=500	KAS=1	KAS=10	KAS=100	KAS=500	KAS=1	KAS=10	KAS=100	KAS=500
Llama-3.2 (1B)	Append	1.1 ± 0.9	1.1 ± 0.9	0.7 ± 0.7	0.9 ± 0.7	4.1 ± 1.7	4.4 ± 1.8	5.4 ± 2.0	4.3 ± 1.9	22.5 ± 3.7	18.6 ± 3.4	18.1 ± 3.3	20.6 ± 3.4
	FT-CK	0.9 ± 0.7	1.1 ± 0.9	1.3 ± 0.9	1.3 ± 0.9	5.1 ± 1.9	4.1 ± 1.7	4.9 ± 1.9	5.7 ± 1.9	21.8 ± 3.6	19.3 ± 3.5	18.4 ± 3.4	19.8 ± 3.4
	MeLLO	2.4 ± 1.2	2.1 ± 1.3	1.9 ± 1.1	2.3 ± 1.3	3.9 ± 1.7	5.8 ± 2.0	5.3 ± 2.1	3.4 ± 1.6	18.0 ± 3.2	17.5 ± 3.1	20.9 ± 3.3	18.6 ± 3.4
Llama-3.2 (3B)	Append	7.3 ± 2.1	6.9 ± 2.3	5.6 ± 2.0	6.7 ± 2.3	16.8 ± 3.2	16.2 ± 3.4	14.1 ± 3.1	14.3 ± 3.1	40.4 ± 4.2	40.1 ± 4.5	41.1 ± 4.3	40.4 ± 4.2
	FT-CK	7.4 ± 2.2	7.7 ± 2.3	6.1 ± 2.1	7.3 ± 2.3	12.8 ± 2.8	16.4 ± 3.4	16.1 ± 3.1	14.7 ± 3.1	41.6 ± 4.2	40.9 ± 4.3	41.0 ± 4.2	41.6 ± 4.4
	MeLLO	7.3 ± 2.3	6.9 ± 2.3	6.7 ± 2.1	7.4 ± 2.2	12.2 ± 2.8	11.4 ± 2.8	11.4 ± 2.8	12.9 ± 2.9	37.9 ± 4.3	41.2 ± 4.4	37.5 ± 4.3	35.3 ± 4.1
Llama-3.2 (11B)	Append	13.1 ± 2.9	12.2 ± 2.8	12.3 ± 2.7	14.2 ± 3.0	18.1 ± 3.3	18.2 ± 3.6	19.5 ± 3.7	19.2 ± 3.4	33.4 ± 4.0	32.2 ± 4.2	33.0 ± 4.2	30.6 ± 4.2
	FT-CK	12.7 ± 2.9	13.6 ± 3.0	13.2 ± 3.0	12.5 ± 2.9	19.1 ± 3.5	19.4 ± 3.6	18.7 ± 3.5	20.3 ± 3.5	32.6 ± 4.2	33.9 ± 4.3	33.8 ± 4.2	34.6 ± 4.2
	MeLLO	11.0 ± 2.6	10.5 ± 2.5	10.3 ± 2.5	11.4 ± 2.6	9.4 ± 2.6	9.2 ± 2.6	8.7 ± 2.5	8.6 ± 2.4	20.3 ± 3.3	20.3 ± 3.5	18.7 ± 3.3	18.6 ± 3.2
Qwen-3 (1.7B)	Append	<b>83.6</b> ± 3.2	59.3 ± 4.1	46.7 ± 4.3	22.2 ± 3.6	18.3 ± 3.3	16.7 ± 3.3	15.0 ± 3.2	17.8 ± 3.4	49.2 ± 4.0	50.8 ± 4.2	47.6 ± 4.2	48.6 ± 4.2
	Append-T	6.4 ± 2.2	3.6 ± 1.6	1.3 ± 0.9	1.4 ± 1.0	15.5 ± 3.1	15.3 ± 3.1	14.8 ± 3.2	15.6 ± 3.2	7.1 ± 2.3	7.3 ± 2.3	7.1 ± 2.1	7.1 ± 2.1
	FT-CK	4.8 ± 1.8	0.7 ± 0.7	9.3 ± 2.5	1.0 ± 0.8	16.7 ± 3.3	16.8 ± 3.4	17.3 ± 3.3	1.7 ± 1.1	45.6 ± 4.4	47.9 ± 4.5	43.3 ± 4.1	48.1 ± 4.1
	MeLLO	5.5 ± 1.9	6.3 ± 2.1	9.4 ± 2.4	9.3 ± 2.5	12.5 ± 2.9	14.0 ± 3.0	13.3 ± 3.1	12.2 ± 2.8	27.8 ± 3.8	27.8 ± 4.0	26.7 ± 3.7	25.3 ± 3.5
Qwen-3 (4B)	Append	80.5 ± 3.5	75.5 ± 3.7	67.0 ± 4.0	58.8 ± 4.4	28.2 ± 3.8	25.2 ± 3.8	25.3 ± 3.7	26.0 ± 3.8	62.3 ± 4.1	78.0 ± 3.6	64.7 ± 4.1	62.8 ± 4.2
	Append-T	46.9 ± 4.3	11.4 ± 2.8	10.6 ± 2.6	3.5 ± 1.7	22.8 ± 3.6	23.0 ± 3.8	23.1 ± 3.7	22.8 ± 3.6	36.5 ± 4.1	32.4 ± 4.2	32.6 ± 3.8	33.4 ± 4.2
	FT-CK	9.0 ± 2.4	0.7 ± 0.7	0.0 ± 0.0	0.0 ± 0.0	25.6 ± 3.6	26.0 ± 3.6	8.4 ± 2.4	0.0 ± 0.0	62.2 ± 4.4	65.8 ± 4.2	62.7 ± 4.3	66.2 ± 4.0
	MeLLO	7.9 ± 2.3	9.6 ± 2.6	17.4 ± 3.2	26.6 ± 3.8	21.4 ± 3.8	21.1 ± 3.5	22.5 ± 3.7	18.7 ± 3.5	57.0 ± 4.4	57.0 ± 4.4	55.8 ± 4.4	55.0 ± 4.4
Qwen-3 (8B)	Append	77.7 ± 3.7	62.3 ± 4.1	57.0 ± 4.4	51.5 ± 4.1	27.8 ± 3.8	29.4 ± 4.0	28.6 ± 4.0	26.0 ± 3.6	70.7 ± 3.9	67.7 ± 4.1	69.7 ± 3.9	70.2 ± 3.8
	Append-T	73.6 ± 3.8	45.6 ± 4.2	24.9 ± 3.7	11.2 ± 2.6	26.5 ± 3.9	25.9 ± 3.7	26.1 ± 3.9	27.6 ± 3.8	70.7 ± 3.9	69.0 ± 4.0	71.4 ± 4.0	69.4 ± 4.0
	FT-CK	10.4 ± 2.6	10.8 ± 2.6	19.0 ± 3.4	4.1 ± 1.7	28.5 ± 3.9	29.1 ± 3.7	28.0 ± 3.8	28.3 ± 3.9	68.7 ± 3.9	69.9 ± 3.9	70.2 ± 3.8	72.6 ± 3.8
	MeLLO	10.6 ± 2.6	11.5 ± 2.7	24.7 ± 3.7	24.7 ± 3.7	24.8 ± 3.8	23.3 ± 3.7	22.8 ± 3.6	23.8 ± 3.6	49.5 ± 4.3	50.5 ± 4.5	51.1 ± 4.3	48.8 ± 4.4
GPT-4.1-mini & o4-mini	Append	70.8 ± 4.0	69.8 ± 4.0	71.4 ± 4.0	70.8 ± 4.0	36.8 ± 4.2	33.7 ± 4.1	37.6 ± 4.2	34.8 ± 4.2	86.4 ± 3.0	82.3 ± 3.3	82.6 ± 3.2	<b>81.7</b> ± 3.3
	Append-T	79.6 ± 3.4	<b>80.3</b> ± 3.5	<b>77.5</b> ± 3.5	<b>78.2</b> ± 3.6	<b>40.6</b> ± 4.2	<b>38.3</b> ± 4.3	<b>38.6</b> ± 4.4	<b>39.4</b> ± 4.2	<b>87.2</b> ± 2.8	<b>85.5</b> ± 3.1	<b>83.3</b> ± 3.1	80.7 ± 3.5

Table 4: **Impact of Knowledge Aggregation Scope (KAS) on Holistic Pass (HP).** We inject a batch of conflicted knowledge associated with KAS = 1, 10, 100, 500, comparing open-book knowledge injection methods (Append, Append-T, FT-CK, MeLLO). We report 95% CIs in the  $\pm$  sign and **bold** the best scores per column.

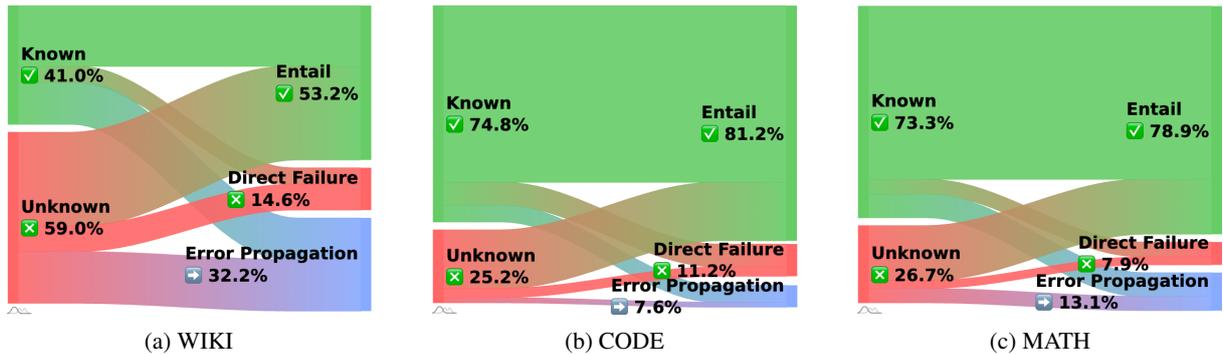


Figure 4: **Failure analysis at the atomic fact level for the Append method**, averaged across all models. The flow shows the percentage of all facts that were initially known or unknown (left) and their outcome in the reasoning (right): successfully integrated (Entail), being the first point of failure (Direct Failure), or failing due to a previous error (Error Propagation).

### 6.3 Analysis of Knowledge Injection Failure

To diagnose why performance remains low even with injected knowledge, we analyze the failure modes. We identify two primary causes: (i) the model’s inability to faithfully use the provided atomic facts, and (ii) flawed reasoning, even when all facts are correctly used.

**Models struggle to integrate atomic facts faithfully.** A primary failure mode is the model’s lack of faithfulness to the provided facts, evidenced by a large gap between achieving a correct answer (AP) and using faithful reasoning (FKE) in Llama-3.2 models on WIKI scenario (Tab. 3). This disparity suggests that models often arrive at the correct answer by coincidence; for instance, a model might use an outdated entity (e.g., a former president of the US) yet reach the correct final answer because

the incorrect entity happens to share the queried attribute (e.g., a spouse’s birthplace) with the correct one. This lack of grounding is a critical failure, proving the model cannot be reliably steered by new information.

We further dissect these faithfulness failures into *direct failures* (the first required atomic fact that is not entailed in reasoning) and *error propagation* (subsequent facts not entailed). As shown in Fig. 4, direct failures are a significant issue, accounting for 7.9–14.6% of outcomes across all scenarios. These

<sup>9</sup>We treat the Append method for the o4-mini model as the Append-T method for the GPT-4.1-mini model, because evidence shows that GPT-4.1-mini and o4-mini share the same knowledge (according to <https://platform.openai.com/docs/models/gpt-4.1-mini>, <https://platform.openai.com/docs/models/o4-mini>, and results in Fig. 3b). We hypothesize that o4-mini and GPT-4.1-mini use similar architectures, but o4-mini is optimized for longer thinking.

failures occur on both known and unknown facts, revealing an inconsistency in how models apply their knowledge during multi-step reasoning. More critically, these initial errors trigger a massive cascade of error propagation failures (7.6%–32.2%), demonstrating a compounding effect where a single mistake derails the entire reasoning process.

**Successful fact integration does not guarantee accurate reasoning.** Even when models successfully incorporate all atomic facts, they can still fail due to flawed reasoning. This is evident from the consistent gap between the FKE and HP scores in Tab. 3. This is most evident in the CODE scenario. For instance, o4-mini with the Append-T method achieves a 76.7% FKE, but this collapses to just a 40.4% HP. This means that in nearly half of the cases where the model correctly utilized all API information, it still fails to produce a functionally correct program. Performance drops are also observed in other scenarios. For MATH, Qwen-3 4B’s FKE of 73.0% results in an HP of only 62.1%, and on WIKI, its 82.5% FKE leads to an 80.6% HP. This demonstrates that incorporating facts into correct reasoning and deriving correct final answers remains challenging.

## 7 Conclusion

We introduce TRACK, a benchmark that instantiates a two-stage framework (probing and injection) with realistic, multi-fact conflicts across diverse scenarios and novel metrics for evaluating reasoning faithfulness. Our experiments reveal a critical failure: providing correct facts often yields limited gains and can even degrade performance, due to both poor knowledge integration and flawed reasoning. TRACK provides a rigorous new tool for the community to measure and guide progress on this crucial challenge.

## Limitations

First, our study is restricted to knowledge propagation within English, text-based reasoning scenarios. Consequently, we do not evaluate performance in more complex environments, such as multi-modal reasoning, multi-lingual contexts, or agent settings involving tool use, which may exhibit different propagation dynamics. Additionally, we conduct experiments on a specific set of open-source and proprietary models (Llama-3.2, Qwen-3, GPT-4.1-mini, and o4-mini). While we include diverse architectures and capabilities (e.g., think-

ing vs. non-thinking variants), this selection is not exhaustive, and our findings may not strictly generalize to significantly larger models or alternative architectures not tested here. Finally, our benchmark represents a static snapshot of knowledge. Unlike mathematical domains, facts in multi-hop QA and code libraries are subject to temporal drift and may become outdated. Although our dynamic data construction pipeline allows for regeneration using updated sources (e.g., new Wikidata dumps), the current evaluation does not cover such dynamics. This built-in extensibility paves the way for future studies on continual and life-long knowledge propagation.

## Ethical Considerations

Our work probes a specific dimension of model reliability: the ability to faithfully reason with new, correct information against internal knowledge, which is foundational for trustworthy AI, particularly in real-world scenarios where knowledge is constantly evolving. However, we acknowledge the ethical risks of blindly following user inputs. Models excelling on our benchmark could be vulnerable to malicious prompt injection attacks (Liu et al., 2024; Xu et al., 2024a) or amplify unintentional misinformation (Kumar et al., 2023; Xu et al., 2025; Feng et al., 2025), for instance in domains like mathematical reasoning. Our paper focuses on the first step: ensuring models can reliably process trustworthy information. We argue this is a prerequisite for the more advanced challenge of discerning when to trust an input. Future work should build upon our benchmark by incorporating non-factual or malicious data, thereby creating a more comprehensive testbed for developing models that can both faithfully utilize correct information and critically reject harmful input.

## Acknowledgments

We are grateful to EPFL NLP lab and specifically Gail Weiss, Silin Gao, and Beatriz Borges for their invaluable feedback and insightful suggestions. We gratefully acknowledge the financial and IT support provided by EPFL, which allows us to use the server and computing resources. We also gratefully acknowledge the support of the Swiss National Science Foundation (No. 215390), Innosuisse (PFFS-21-29), the EPFL Center for Imaging, Sony Group Corporation, and a Meta LLM Evaluation Research Grant.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). arXiv preprint arXiv:2303.08774.
- Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. 2024. [Large language models for mathematical reasoning: Progresses and challenges](#). In [Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop](#), pages 225–237, St. Julian’s, Malta. Association for Computational Linguistics.
- amCharts. 2025. [amcharts 5: Javascript charts](#). Accessed: 2025-02-15.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. [Evaluating large language models trained on code](#). Preprint, arXiv:2107.03374.
- Roi Cohen, Eden Biran, Ori Yoran, Amir Globerson, and Mor Geva. 2024. [Evaluating the ripple effects of knowledge editing in language models](#). [Transactions of the Association for Computational Linguistics](#), 12:283–298.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). arXiv preprint arXiv:2507.06261.
- Tri Dao. 2024. [Flashattention-2: Faster attention with better parallelism and work partitioning](#). In [The Twelfth International Conference on Learning Representations](#).
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. [Editing factual knowledge in language models](#). In [Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing](#), pages 6491–6506, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yiyang Feng, Yichen Wang, Shaobo Cui, Boi Faltings, Mina Lee, and Jiawei Zhou. 2025. [Unraveling misinformation propagation in llm reasoning](#). Preprint, arXiv:2505.18555.
- Govind Krishnan Gangadhar and Karl Stratos. 2024. [Model editing by standard fine-tuning](#). In [Findings of the Association for Computational Linguistics: ACL 2024](#), pages 5907–5913, Bangkok, Thailand. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. [The llama 3 herd of models](#). Preprint, arXiv:2407.21783.
- Karina Halevy, Anna Sotnikova, Badr AlKhamissi, Syrielle Montariol, and Antoine Bosselut. 2024. [“flex tape can’t fix that”: Bias and misinformation in edited language models](#). In [Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing](#), pages 8690–8707, Miami, Florida, USA. Association for Computational Linguistics.
- Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. 2020. [Array programming with numpy](#). [Nature](#), 585(7825):357–362.
- Peter Hase, Mohit Bansal, Been Kim, and Asma Ghandeharioun. 2023. [Does localization inform editing? surprising differences in causality-based localization vs. knowledge editing in language models](#). In [Thirty-seventh Conference on Neural Information Processing Systems](#).
- Peter Hase, Mona Diab, Asli Celikyilmaz, Xian Li, Zornitsa Kozareva, Veselin Stoyanov, Mohit Bansal, and Srinivasan Iyer. 2021. [Do language models have beliefs? methods for detecting, updating, and visualizing model beliefs](#). Preprint, arXiv:2111.13654.
- Peter Hase, Thomas Hofweber, Xiang Zhou, Elias Stengel-Eskin, and Mohit Bansal. 2024. [Fundamental problems with model editing: How should rational belief revision work in LLMs?](#) [Transactions on Machine Learning Research](#).
- Wenyue Hua, Jiang Guo, Mingwen Dong, Henghui Zhu, Patrick Ng, and Zhiguo Wang. 2024. [Propagation and pitfalls: Reasoning-based assessment of knowledge editing through counterfactual tasks](#). In [Findings of the Association for Computational Linguistics: ACL 2024](#), pages 12503–12525, Bangkok, Thailand. Association for Computational Linguistics.
- Baixiang Huang, Canyu Chen, Xiong Xiao Xu, Ali Payani, and Kai Shu. 2025. [Can knowledge editing really correct hallucinations?](#) In [The](#)

- Thirteenth International Conference on Learning Representations.
- John Hunter and Darren Dale. 2007. The matplotlib user’s guide. Matplotlib 0.90.0 user’s guide.
- Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. A survey on large language models for code generation. Preprint, arXiv:2406.00515.
- Evgenii Kortukov, Alexander Rubinstein, Elisa Nguyen, and Seong Joon Oh. 2024. Studying large language model behaviors under context-memory conflicts with real documents. In First Conference on Language Modeling.
- Harsh Kumar, David M Rothschild, Daniel G Goldstein, and Jake M Hofman. 2023. Math education with large language models: Peril or promise? Available at SSRN 4641653.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. Let’s verify step by step. In The Twelfth International Conference on Learning Representations.
- Yi Liu, Gelei Deng, Yuekang Li, Kailong Wang, Zihao Wang, Xiaofeng Wang, Tianwei Zhang, Yepang Liu, Haoyu Wang, Yan Zheng, and Yang Liu. 2024. Prompt injection attack against llm-integrated applications. Preprint, arXiv:2306.05499.
- Zeyu Leo Liu, Shrey Pandit, Xi Ye, Eunsol Choi, and Greg Durrett. 2025. Codeupdatearena: Benchmarking knowledge editing on api updates. Preprint, arXiv:2407.06249.
- Shayne Longpre, Kartik Perisetla, Anthony Chen, Nikhil Ramesh, Chris DuBois, and Sameer Singh. 2021. Entity-based knowledge conflicts in question answering. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, pages 7052–7063, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Wes McKinney et al. 2011. pandas: a foundational python library for data analysis and statistics. Python for high performance and scientific computing, 14(9):1–9.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2023a. Locating and editing factual associations in gpt. Preprint, arXiv:2202.05262.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2023b. Mass-editing memory in a transformer. Preprint, arXiv:2210.07229.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D. Manning. 2022a. Fast model editing at scale. Preprint, arXiv:2110.11309.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher D Manning, and Chelsea Finn. 2022b. Memory-based model editing at scale. In Proceedings of the 39th International Conference on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pages 15817–15831. PMLR.
- Yasumasa Onoe, Michael Zhang, Shankar Padmanabhan, Greg Durrett, and Eunsol Choi. 2023. Can LMs learn new entities from descriptions? challenges in propagating injected knowledge. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 5469–5485, Toronto, Canada. Association for Computational Linguistics.
- OpenAI. 2025a. Introducing gpt-4.1 in the api.
- OpenAI. 2025b. Introducing openai o3 and o4-mini.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. Advances in neural information processing systems, 32.
- Auguste Poiroux, Gail Weiss, Viktor Kunčák, and Antoine Bosselut. 2025. Reliable evaluation and benchmarks for statement autoformalization. In Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, pages 17958–17980, Suzhou, China. Association for Computational Linguistics.
- Domenic Rosati, Robie Gonzales, Jinkun Chen, Xuemin Yu, Yahya Kayani, Frank Rudzicz, and Hassan Sajjad. 2024. Long-form evaluation of model editing. In Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), pages 3749–3780, Mexico City, Mexico. Association for Computational Linguistics.
- Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. 2023. Measuring inductive biases of in-context learning with underspecified demonstrations. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 11289–11310, Toronto, Canada. Association for Computational Linguistics.
- Joykirat Singh, Akshay Nambi, and Vibhav Vineet. 2024. Exposing the achilles’ heel: Evaluating llms ability to handle mistakes in mathematical reasoning. Preprint, arXiv:2406.10834.
- Zayne Rea Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2025. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. In The Thirteenth International Conference on Learning Representations.

- Lukas Thede, Karsten Roth, Matthias Bethge, Zeynep Akata, and Thomas Hartvigsen. 2025. [Wikibigedit: Understanding the limits of lifelong knowledge editing in LLMs](#). In [Forty-second International Conference on Machine Learning](#).
- TogetherAI. 2025. [Together ai: The ai acceleration cloud](#).
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). [Commun. ACM](#), 57(10):78–85.
- Yike Wang, Shangbin Feng, Heng Wang, Weijia Shi, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2024. [Resolving knowledge conflicts in large language models](#). In [First Conference on Language Modeling](#).
- Michael L Waskom. 2021. [Seaborn: statistical data visualization](#). [Journal of Open Source Software](#), 6(60):3021.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In [Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations](#), pages 38–45, Online. Association for Computational Linguistics.
- Kevin Wu, Eric Wu, and James Zou. 2024. [Clash: Quantifying the tug-of-war between an LLM’s internal prior and external evidence](#). In [The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track](#).
- Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2024. [Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts](#). In [The Twelfth International Conference on Learning Representations](#).
- Rongwu Xu, Brian Lin, Shujian Yang, Tianqi Zhang, Weiyan Shi, Tianwei Zhang, Zhixuan Fang, Wei Xu, and Han Qiu. 2024a. [The earth is flat because...: Investigating LLMs’ belief towards misinformation via persuasive conversation](#). In [Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 16259–16303, Bangkok, Thailand. Association for Computational Linguistics.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024b. [Knowledge conflicts for LLMs: A survey](#). In [Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing](#), pages 8541–8565, Miami, Florida, USA. Association for Computational Linguistics.
- Tianlong Xu, Yi-Fan Zhang, Zhendong Chu, Shen Wang, and Qingsong Wen. 2025. [Ai-driven virtual teacher for enhanced educational efficiency: leveraging large pretrained models for autonomous error analysis and correction](#). In [Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence, AAAI’25/IAAI’25/EAAI’25](#). AAAI Press.
- An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#). [Preprint, arXiv:2505.09388](#).
- Yunzhi Yao, Peng Wang, Bozhong Tian, Siyuan Cheng, Zhoubo Li, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023. [Editing large language models: Problems, methods, and opportunities](#). In [Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing](#), pages 10222–10240, Singapore. Association for Computational Linguistics.
- Jiahao Ying, Yixin Cao, Kai Xiong, Long Cui, Yidong He, and Yongbin Liu. 2024. [Intuitive or dependent? investigating LLMs’ behavior style to conflicting prompts](#). In [Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 4221–4246, Bangkok, Thailand. Association for Computational Linguistics.
- Shaochen Zhong, Yifan Lu, Lize Shao, Bhargav Bhushanam, Xiaocong Du, Yixin Wan, Yucheng Shi, Daochen Zha, Yiwei Wang, Ninghao Liu, Kaixiong Zhou, Shuai Xu, Kai-Wei Chang, Louis Feng, Vipin Chaudhary, and Xia Hu. 2025. [MQuAKE-remastered: Multi-hop knowledge editing can only be advanced with reliable evaluations](#). In [The Thirteenth International Conference on Learning Representations](#).
- Zexuan Zhong, Zhengxuan Wu, Christopher Manning, Christopher Potts, and Danqi Chen. 2023. [MQuAKE: Assessing knowledge editing in language models via multi-hop questions](#). In [Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing](#), pages 15686–15702, Singapore. Association for Computational Linguistics.

Andrew Zhu, Alyssa Hwang, Liam Dugan, and Chris Callison-Burch. 2024. [Fanoutqa: A multi-hop, multi-document question answering benchmark for large language models](#). [Preprint](#), arXiv:2402.14116.

Terry Yue Zhuo, Vu Minh Chien, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen GONG, James Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kadour, Ming Xu, Zhihan Zhang, Prateek Yadav, Naman Jain, Alex Gu, Zhoujun Cheng, Jiawei Liu, Qian Liu, Zijian Wang, David Lo, Binyuan Hui, Niklas Muennighoff, Daniel Fried, Xiaoning Du, Harm de Vries, and Leandro Von Werra. 2025. [Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions](#). In [The Thirteenth International Conference on Learning Representations](#).

## A Evaluation Metrics

### A.1 Knowledge Probing Evaluation

We classify an atomic fact  $k_i$  as *known* by verifying LLM’s responses to the corresponding probing question  $q_i$  against the probing answer  $a_i$  (i.e.,  $\text{Verify}(f(q_i), a_i) = \text{True}$ ). The verification process is as follows. First, we sample  $M = 10$  responses for the probing question  $q_i$  and group them by identical text to obtain unique answers  $\{\hat{a}_{i,1}, \dots, \hat{a}_{i,M}\}$  with corresponding frequencies  $\{c_{i,1}, \dots, c_{i,M}\}$ . The equivalence ( $\equiv$ ) between each unique response and the ground-truth probing answer  $a_i$  is judged by gpt-5-mini with a tailored prompt for each scenario. Based on these judgments, we compute the total count for the merged group of correct answers,  $c_{\text{correct}} = \sum_j \text{s.t. } \hat{a}_{i,j} \equiv a_i c_{i,j}$ . A fact  $k_i$  is **Known** if  $c_{\text{correct}} = \max(c_{\text{correct}}, c_{i,1}, \dots, c_{i,M})$ . Facts that do not meet this criterion are classified as *unknown*. We also define *Knowledge Confidence (KConf)* as the total correct count normalized by the number of samples:  $\text{KConf} = \frac{c_{\text{correct}}}{N} \times 100\%$ .

Our prompt includes detailed instructions and demonstrations tailored to each complex reasoning scenario. The prompts are as follows:

#### Scenario WIKI.

##### System:

You are given a question, a response, and a ground truth answer. The ground truth answer is an entity, and each response might contain an entity. Your task is to use commonsense knowledge to evaluate whether the response most probably refers the same entity as the ground truth.

If they are equivalent, answer 'Yes' and provide an explanation. Otherwise, answer 'No' and provide an explanation.

Note that if the response does not contain an entity, it should be treated as 'N/A' and not equivalent to the answer.

Examples:

Question:  
Who is the current US president?

Response:  
Therefore, the answer is Donald Trump.

Answer:  
Donald J. Trump

Equivalence:  
Yes, Donald Trump is the same person as Donald J. Trump.

Question:  
Who is the current US president?

Response:  
Therefore, the answer is Donald Trump.

Answer:  
Joe Biden

Equivalence:  
No, Donald Trump is a different person from Joe Biden. They belong to different political parties.

Question:  
Who is Albuquerque’s head of government?

Response:  
Based on my cutoff knowledge, Albuquerque’s head of government is Tim Keller.

Answer:  
Timothy M. Keller

Equivalence:  
Yes, 'Tim' is a common short form of 'Timothy'.

Question:  
Who is Albuquerque’s head of government?

Response:  
I cannot provide an exact answer based on my cutoff knowledge.

Answer:  
Timothy M. Keller

Equivalence:  
No, the response fails to provide an entity (N/A), while the answer provides the entity 'Timothy M. Keller'.

##### User:

Question:  
{a probing question}

Response:  
{a model's response to the probing question}  
Answer:  
{the ground-truth probing answer}  
Equivalence:

## Scenario CODE.

### System:

You are given a question, a canonical function from a library for the question, and the model's response. Each response might also contain a function call from a library. Your task is to use basic Python coding knowledge to evaluate whether the model's response is most probably correct.

If the answer is correct, answer 'Yes' and provide an explanation. Otherwise, answer 'No' and provide an explanation.

Note that if the response does not contain a function call, it should be treated as 'N/A' and not correct.

– Examples 1 –

Question:

Given the library pandas, how can we create a DataFrame by explicitly passing the input data (such as an ndarray, Iterable, dict, or DataFrame) using the 'data' parameter?

Function:

```
pandas.DataFrame(data)
```

Response:

```
“python  
pandas.DataFrame(arr)  
“
```

Correct:

Yes, the response contains the same function call as the ground truth function call. The function call 'pandas.DataFrame(arr)' is equivalent to the ground truth function call, which creates a DataFrame from the provided data.

– Example 2 –

Question:

Given the library pandas, how can we create a DataFrame by explicitly passing the input data (such as an ndarray, Iterable, dict, or DataFrame) using the 'data' parameter?

Function:

```
pandas.DataFrame(data)
```

Response:

```
“python  
pandas.DataFrame("id": [0, 1, 2, 3,  
4], "val": [100, 200, -2, 34, 45.2],  
dtype=None)  
“
```

Correct:

Yes, the response contains the same function call as the ground truth function call. The function call 'pandas.DataFrame("id": [0, 1, 2, 3, 4], "val": [100, 200, -2, 34, 45.2], dtype=None)' is equivalent to the ground truth function call, which creates a DataFrame from the provided data. The 'dtype' parameter is optional and defaults to None, so it does not change the equivalence.

– Example 3 –

Question:

Given the library pandas, how can we create a DataFrame by explicitly passing the input data (such as an ndarray, Iterable, dict, or DataFrame) using the 'data' parameter?

Function: pandas.DataFrame(data)

Response:

```
“python  
pandas.DataFrame(data,  
dtype="float")  
“
```

Correct:

No, the response contains a different function call than the ground truth function call. The function call 'pandas.DataFrame(data,

`dtype="float")` specifies a dtype of "float", which is not equivalent to the ground truth function call that does not specify a dtype. The ground truth function call creates a DataFrame from the provided data without any specific dtype.

**User:**

Question:

{a probing question}

Function:

{the ground-truth API calling function}

Response:

{the model's response of API calling}

Correct:

**Scenario MATH.**

**System:**

You are given a question, a response, and a ground truth answer. Your task is to use math knowledge and ground truth to evaluate whether the response most probably answers the question.

If the response answers the question, answer 'Yes' and provide an explanation. Otherwise, answer 'No' and provide an explanation.

Note that the ground truth is just a reference of correct answer. An answer is correct does not mean that it should be strictly same as the ground truth. You should perform mathematical knowledge to infer that.

Examples:

Question:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , what specific operation will eliminate the variable 'e'?

Response:

Subtracting the first equation from

the second one, then we will have  $2p = 0.58$  and eliminate 'e'.

Answer:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , subtracting the first equation from the second will eliminate the variable 'e'.

Correctness:

Yes, it mentions what it is in the ground truth that subtracting one equation from another will eliminate the variable 'e'.

Question:

If half the value of  $3x - 9$  is equal to  $x + 37$ , how can this relationship be written as an equation?

Response:

$3x - 9 = 2(x + 37)$

Answer:

If half the value of  $3x - 9$  is equal to  $x + 37$ , it can be written as  $(1/2)(3x - 9) = x + 37$ .

Correctness:

Yes, it mentions  $3x - 9 = 2(x + 37)$  while the ground truth mentions  $(1/2)(3x - 9) = x + 37$ . Though they are different, they are mathematically equivalent and the response is also correct.

Question:

How do you combine the like imaginary terms  $-3i$  and  $12i$  in a sum?

Response:

To combine the like imaginary terms, add or subtract their coefficients, which are the numbers in front of the 'i'.

Answer:

Like terms with the same imaginary unit  $i$  are combined by adding their coefficients, so  $-3 + 12 = 9$ , giving  $9i$ .

Correctness:

Yes, though the response does not explicitly give  $9i$ , the answer "add or subtract their coefficients, which are the numbers in front of the 'i'" is also valid because it

tells how to imaginary terms  $-3i$  and  $12i$ .

Question:

According to the associative property of addition, how can we regroup the sum  $6 - 3i + 12i$ ?

Response:

The associative property of addition states that the grouping of terms does not affect their sum, so  $6 - 3i + 12i$  can be written as  $(6 - 3i) + 12i$ .

Answer:

We can regroup the sum as  $6 + (12i - 3i)$ .

Correctness:

Yes, though the response does not match the specific regrouping given in the answer, both are valid regroupings due to the associative property.

Question:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , what specific operation will eliminate the variable 'e'?

Response:

Adding the first equation to the second one will eliminate 'e'.

Answer:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , subtracting the first equation from the second will eliminate the variable 'e'.

Correctness:

No, it mentions adding one equation to another, different from the ground truth answer which subtracts the first equation from the second one.

Question:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , what specific operation will eliminate the variable 'e'?

Response:

After eliminating 'e', we have  $p = 0.29$ .

Answer:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , subtracting the first equation from the second will eliminate the variable 'e'.

Correctness:

No, it fails to mention that subtracting the first equation from the second one will eliminate 'e'.

**User:**

Question:

{a probing question}

Response:

{a model's response to the probing question}

Answer:

{the ground-truth probing answer}

Correctness:

## A.2 Knowledge Injection Evaluation

For each complex reasoning question  $q$ , we evaluate a model  $f$ 's response in both closed-book  $f(q)$  and open-book settings  $f(q, C)$ , which consists of reasoning steps  $\hat{R}$  and a final answer  $\hat{a}$ , using three metrics: Answer Pass (AP) for final answer correctness, Full Knowledge Entailment (FKE) for reasoning faithfulness, and Holistic Pass (HP), which requires both AP and FKE to be true ( $HP = AP \wedge FKE$ ).

The specific implementation of AP and FKE is tailored to each scenario, though the FKE check always involves using gpt-5-mini to perform Natural Language Inference (NLI).

- For the WIKI and MATH scenarios: The AP is determined by using gpt-4.1-mini to assess the equivalence between the model's extracted final answer  $\hat{a}$  and the ground-truth answer  $a^*$ . The FKE check verifies that the reasoning steps  $\hat{R}$  entail each required atomic fact.
- For the CODE scenario: The evaluation is based on functional correctness. The model's generated function is treated as the reasoning steps  $\hat{R}$ . The AP is true if this function passes the predefined unit tests from BigCodeBench (Zhuo et al., 2025). The FKE check verifies that the implemented code in  $\hat{R}$  is faithful to the functionality described in the provided API documentation (the atomic facts).

Below are detailed prompts for AP and FKE metrics together with WIKI, CODE, and MATH scenarios.

### A.2.1 Answer Pass (AP)

AP is a binary metric of correctness for the final answer  $\hat{a}$ :  $AP = \text{Is\_Correct}(\hat{a}, a)$ . Its verification method is domain-specific. For WIKI and MATH, the gpt-5-mini model assesses the equivalence between the model's extracted final answer and the ground truth  $a$ , i.e.,  $\text{Is\_Correct}(\hat{a}, a) = \mathbb{I}(\hat{a} \equiv a)$ . For CODE, AP is determined by functional correctness, using standard Pass@1 (Chen et al., 2021): AP is true if the executed function passes a predefined suite of unit tests once, i.e.,  $\text{Is\_Correct}(\hat{a}, a) = \text{Pass@1}(\hat{a})$ . Given some implementation issues of building unit test environments in the original BigCodeBench, we use GPT-5-mini as a judge to detect whether the LLM-generated code can pass the unit test compared with the ground truth code.

**Scenario WIKI.** We use the same prompt as Scenario WIKI in § A.1.

**Scenario CODE.** We pass the function to predefined unit tests from BigCodeBench (Zhuo et al., 2025). The unit test result is treated as the final answer  $\hat{a}$ , and is correct if it passes the unit test.

**Scenario MATH.**

#### System:

You are given a question, a final step of response, and a ground truth answer. The final step may contain a step number and "the answer is ...". PLEASE IGNORE THE STEP NUMBER. Your task is to use math knowledge to evaluate whether the response is mathematically equivalent to the ground truth answer.

If they are equivalent, answer 'Yes' and provide an explanation. Otherwise, answer 'No' and provide an explanation.

Examples:

Question:

Evaluate  $(1 + 2i)6 - 3i$ .

Response:

7. The answer is  $9i + 6$ .

Ground Truth:

$6 + 9i$

Correctness:

Yes, the answer from response  $9i + 6$  is equivalent to the ground truth  $6 + 9i$ .

Question:

Evaluate  $(1 + 2i)6 - 3i$ .

Response:

7. The answer is  $9 + 6i$ .

Ground Truth:

$6 + 9i$

Correctness:

No, the answer from response is different from ground truth. The real part and the imaginary part are reversed.

Question:

The lengths of two opposite sides of a square are decreased by 40% while the lengths of the other two sides are increased by 50% to form a rectangle. By what percent does the square's area decrease?

Response:

Therefore, the square's area decreases by 10%.

Ground Truth:

10

Correctness:

Yes, the response "10%" is equivalent to the ground truth "10" because the question asks for a percentage.

#### User:

Question:

{a complex reasoning question}

Response:

{a model's final step to the complex reasoning question}

Answer:

{the ground-truth final answer}

Correctness:

### A.2.2 Full Knowledge Entailment (FKE)

FKE is a binary metric that evaluates whether the reasoning process is faithful to the atomic facts. The entailment check is performed by gpt-5-mini, guided by domain-specific NLI prompts. FKE is true if the reasoning steps  $\hat{R}$  together with final

answer  $\hat{a}$  entail every required atomic fact  $k_i$  from the full set  $K_q$ :  $\text{FKE} = \bigwedge_{k_i \in K_q} \mathbb{I}(\text{NLI}(k_i, \hat{R}) = \text{Entailment})$ .

### Scenario WIKI.

#### System:

You are an expert in natural language inference and commonsense reasoning. You will be given a "Context" (the model's reasoning response) and a "Statement" (a piece of knowledge). Your task is to determine if the Context finally entails, contradicts, or is neutral with respect to the Statement.

Answer "Entailment", "Contradiction", or "Neutral" and provide a brief explanation of your reasoning.

Note that if the Context mentions some knowledge is "unknown", it should be treated as "N/A" and contradictory to the Statement.

Examples:

Context:

1. As of my knowledge cutoff in December 2023, the current Prime Minister of the United Kingdom is Rishi Sunak, and Keir Starmer's spouse is Victoria Starmer.
2. The user provided that Keir Starmer is the current head of government. As of my knowledge, he is the leader of the Opposition and the leader of the Labour Party in the UK. So it is possible that he is the Prime Minister now.
3. Therefore, the current British Prime Minister is Keir Starmer.
4. Victoria Starmer is married to British Prime Minister.
5. The answer is Victoria Starmer.

Statement:

The name of the current head of the British government is Keir Starmer.

NLI:

Entailment. The context initially stated that Rishi Sunak is the current Prime Minister, but later

it corrected that Keir Starmer is the current head of government.

Context:

... (the same context as above)

Statement:

The name of the current head of the British government is Rishi Sunak.

NLI:

Contradiction. The context initially stated that Rishi Sunak is the current Prime Minister, but later it corrected that Keir Starmer is the current head of government.

Context:

... (the same context as above)

Statement:

The name of the current head of the US government is Donald Trump.

NLI:

Neutral. The context does not provide any information about the current head of the US government.

Context:

1. I don't know the current head of the British government.
2. I don't have the information about the answer.

Statement:

The name of the current head of the British government is Keir Starmer.

NLI:

Contradiction. The context states that it does not know the current head of the British government, which contradicts the statement that Keir Starmer is the current head of the British government.

Context:

... (the same context as above)

Statement:

The name of the current head of the US government is Donald Trump.

NLI:

Neutral. The context does not provide any information about the current head of the US government.

**User:**

Context:

{a model's generated reasoning steps}

Statement:

{one of the atomic facts required to answer the complex reasoning question}

NLI:

**Scenario CODE.****System:**

You are an expert in Python programming and natural language inference. You will be given a 'Code' snippet and a 'Function'. Your task is to determine if the Code's usage of the function **\*\*Entails\*\***, **\*\*Contradicts\*\***, or is **\*\*Neutral\*\*** with respect to the correct usage of the function.

**\*\*Reasoning Framework:\*\***

1. **\*\*Check for Usage:\*\*** First, determine if the 'Code' calls the specified function. If it does not, the answer is **\*\*Neutral\*\***.

2. **\*\*Validate the Call:\*\*** If the function is called, analyze the arguments used in the 'Code' against the 'Function Signature'.

\* **\*\*Entailment:\*\*** The usage is valid. This means:

\* All required arguments are provided.

\* Any keyword arguments used are valid (i.e., they exist in the function signature).

\* **\*\*Crucially,** omitting optional arguments (e.g., those with default values like 'r=None') is a valid use case.\*\*

\* **\*\*Contradiction:\*\*** The usage is invalid. This means:

\* A required argument is missing.

\* An incorrect or non-existent keyword argument

is used (e.g., 'datatype=' when it should be 'dtype=').

– Example 1 –

Code:

```

“python
import pandas as pd
def task_func(dealer_sales_data):
    # Step 1: Create DataFrame &
    # Step 2: Handle Empty Input (if
    dealer_sales_data is empty)
    df =
pd.DataFrame(dealer_sales_data,
dtype=None)
    if not dealer_sales_data:
        return []
    # Ensure 'id' and 'num_sold'
columns exist, otherwise it's
malformed input
    if 'id' not in df.columns or
'num_sold' not in df.columns:
        return []
    # Step 3: Find Max Sales
    max_sold =
df['num_sold'].max()
    # Step 4: Identify Top Sellers
    top_selling_cars =
df[df['num_sold'] == max_sold]
    # Step 5: Extract and Sort IDs
    top_selling_ids =
top_selling_cars['id'].tolist()
    sorted_ids =
sorted(top_selling_ids)
    # Step 6: Return Result
    return sorted_ids
“

```

Function: pandas.DataFrame(data)

NLI:

Entailment. The code contains the function call 'pd.DataFrame(dealer\_sales\_data)'. The usage 'pd.DataFrame(dealer\_sales\_data)' is entailed in the provided function. Only one positional parameter and the optional parameter 'dtype' is set to 'None', which is by default.

– Example 2 –

```
Code:
“python
import pandas as pd
def task_func(dealer_sales_data):
    # Step 1: Create DataFrame &
    Step 2: Handle Empty Input (if
    dealer_sales_data is empty)
    df =
pd.DataFrame(dealer_sales_data,
dtype="float")
    if not dealer_sales_data:
        return []
    # Ensure 'id' and 'num_sold'
    columns exist, otherwise it's
    malformed input
    if 'id' not in df.columns or
    'num_sold' not in df.columns:
        return []
    # Step 3: Find Max Sales
    max_sold =
df['num_sold'].max()
    # Step 4: Identify Top Sellers
    top_selling_cars =
df[df['num_sold'] == max_sold]
    # Step 5: Extract and Sort IDs
    top_selling_ids =
top_selling_cars['id'].tolist()
    sorted_ids =
sorted(top_selling_ids)
    # Step 6: Return Result
    return sorted_ids
“
```

Function: pandas.DataFrame(data)

NLI:  
 Contradiction. The code contains a related function call 'pd.DataFrame(dealer\_sales\_data, datatype="float")', but it uses a different keyword argument 'datatype' which does not exist. The usage of the function contradicts with the correct usage of of the provided function.

– Example 3 –

Code:  
 ... (the same code as above)

Function:  
 sklearn.linear\_model.LinearRegress

ion()  
 NLI:  
 Neutral. The code does not contain any function call related to the provided function.

– Example 4 –

Code:  
 N/A  
 Function: ... (any function)

NLI:  
 Neutral. The code does not contain any function call related to the provided function.

**User:**

Code:  
 {a model's generated code}  
 Function: {one of the required external function to solve the coding problem}  
 NLI:

**Scenario MATH.**

**System:**

You are an expert in natural language inference and math reasoning. You will be given a "Context" (the model's reasoning response) and a "Statement" (a piece of knowledge). Your task is to determine if the Context finally entails, contradicts, or is neutral with respect to the Statement.

Answer "Entailment", "Contradiction", or "Neutral" and provide a brief explanation of your reasoning.

Note that if the Statement can be indirectly implied from the Context with math reasoning, it should also be treated as entailment.

Examples:

Context:  
 Three pencils and a jumbo eraser

cost \$1.24. Five pencils and a jumbo eraser cost \$1.82. No prices include tax. In cents, what is the cost of a pencil?

1. Let's call the price of a pencil  $x$  and the price of a jumbo eraser  $y$ . Then we can write two equations.

2. We have  $3x + y = 1.24$  and  $5x + y = 1.82$ .

3. To solve this system, let's subtract the first equation from the second equation. This will eliminate  $y$ .

4. This simplifies to  $x = 0.29$ .

5. That means a pencil costs 29 cents.

6. The answer is 29 cents.

Statement:

After subtracting  $3p + e = 1.24$  from  $5p + e = 1.82$ , we will have  $2p = 0.58$ , which solves to  $p = 0.29$ .

NLI:

Entailment. The statement is a direct summary of the mathematical reasoning presented in steps 2, 3, and 4 in the context. It describes the exact same process and reaches the identical conclusion. The use of different variables ( $p$  and  $e$  instead of  $x$  and  $y$ ) is a superficial change that doesn't affect the logic, and missing  $2p = 0.58$  does not matter because  $x = 0.29$  already implies it.

Context:

Three pencils and a jumbo eraser cost \$1.24. Five pencils and a jumbo eraser cost \$1.82. No prices include tax. In cents, what is the cost of a pencil?

1. Let's call the price of a pencil  $x$  and the price of a jumbo eraser  $y$ . Then we can write two equations.

2. We have  $3x + y = 1.24$  and  $5x + y = 1.82$ .

3. To solve this system, let's add the first equation to the second equation. This will eliminate  $y$ .

4. This simplifies to  $2x = 0.58$ . So  $x = 0.29$ .

5. That means a pencil costs 29 cents.

6. The answer is 29 cents.

Statement:

After subtracting  $3p + e = 1.24$  from  $5p + e = 1.82$ , we will have  $2p = 0.58$ , which solves to  $p = 0.29$ .

NLI: Contradiction. The context states in step 3 that the two equations should be added. The statement, however, describes the process using subtraction. Since adding and subtracting are opposite operations, the statement directly contradicts the method described in the context.

Context:

Three pencils and a jumbo eraser cost \$1.24. Five pencils and a jumbo eraser cost \$1.82. No prices include tax. In cents, what is the cost of a pencil?

1. Let's call the price of a pencil  $x$  and the price of a jumbo eraser  $y$ . Then we can write two equations.

2. We have  $3x + y = 1.24$  and  $5x + y = 1.82$ .

3. To solve this system, let's add the first equation to the second equation. This will eliminate  $y$ .

4. This simplifies to  $2x = 0.58$ . So  $x = 0.29$ .

5. That means a pencil costs 29 cents.

6. The answer is 29 cents.

Statement:

The commutative rule, also known as the commutative property, states that the order of numbers in addition and multiplication doesn't change the result.

NLI:

Neutral. The context does not provide any information about the commutative rules. The statement is completely irrelevant to the context.

**User:**

Context:

{a model’s generated reasoning steps}

Statement:

{one of the atomic facts required to answer the complex reasoning question}

NLI:

**A.2.3 Holistic Pass (HP)**

HP is a stricter binary metric requiring both a correct final answer and faithful reasoning:  $HP = AP \wedge FKE$ .

**B Dataset Generation**

The general pipeline for each scenario is visualized in Fig. 5. Here we provide additional details on the data generation pipeline for each of our three scenarios, supplementing the descriptions in Sec. 4.

**B.1 Scenario WIKI**

For this scenario, we generate multi-hop QA by sampling relational chains from recent Wikidata (Vrandečić and Krötzsch, 2014).<sup>10</sup> Our pipeline first samples a chain of varied depths (2 to 6 hops) that satisfies multiple constraints, such as ensuring entities are unique and have valid labels. From this chain, we automatically construct the five key data components. The complex reasoning question ( $q$ ) is generated by composing templated natural-language questions for each relation in the chain. The final entity in the chain serves as the final answer ( $a$ ). Each triple in the chain is converted into a declarative sentence and an interrogative sentence to form an atomic fact ( $k_i$ ) and a probing question ( $q_i$ ), respectively. The object entity of the triple serves as a probing answer ( $a_i$ ). Note that questions involving Wikidata may contain information of entities, but such information is public available and it does not include offensive content.

**Chain Sampling Constraints**

To ensure the quality, complexity, and realism of the generated data, we enforce several constraints during the chain sampling process:

- **Varied Depth:** We sample 120, 100, 100, 100, 80 chains with length of 2, 3, 4, 5, 6 hops (i.e., contain 2, 3, 4, 5, 6 triples), respectively.

<sup>10</sup>We sample chains from Wikidata in Sep. 2025.

- **No Cycles:** An entity cannot appear more than once in a given chain, preventing trivial loops.
- **Property Whitelist:** All relations (properties) in the chain must belong to a predefined set of allowed properties (e.g., ‘P39’: position held, ‘P26’: spouse), focusing the data on common and meaningful relationships.
- **Single Claim Constraint:** We only consider properties for which an entity has a single, unambiguous claim, avoiding ambiguity from multiple values.
- **Label Validity:** All entities and properties in the chain must have non-empty English labels. We also filter out entities with excessively long labels (more than 5 words or 50 characters) to maintain readability.
- **Type Constraints:** For certain properties, we programmatically verify that the subject or object entity is of the correct type using SPARQL queries (e.g., for ‘P36’ ‘capital of’, the subject must be an instance of a country ‘wd:Q6256’).

**Component Generation from Templates**

Once a valid chain is sampled, its components are generated using the templates detailed in Tab. 5.

- **Complex Reasoning Question ( $q$ ):** We recursively compose a natural language question. Starting with the head entity of the chain, we use the “attributive” phrase from our templates (e.g., “the spouse of [S]”) for each hop, nesting them to form a complex question.
- **Final Answer ( $a$ ):** This is simply the English label of the final entity in the 4-hop chain.
- **Atomic Fact ( $k_i$ ):** For each triple ‘(Subject, Property, Object)’ in the chain, we generate a declarative sentence using a cloze-style template (e.g., “[S] is married to [O].”).
- **Probing Question ( $q_i$ ):** For each triple, we use a question template corresponding to the property (e.g., “Who is [S] married to?”).
- **Probing Answer ( $a_i$ ):** This is the English label of the object entity in the triple.

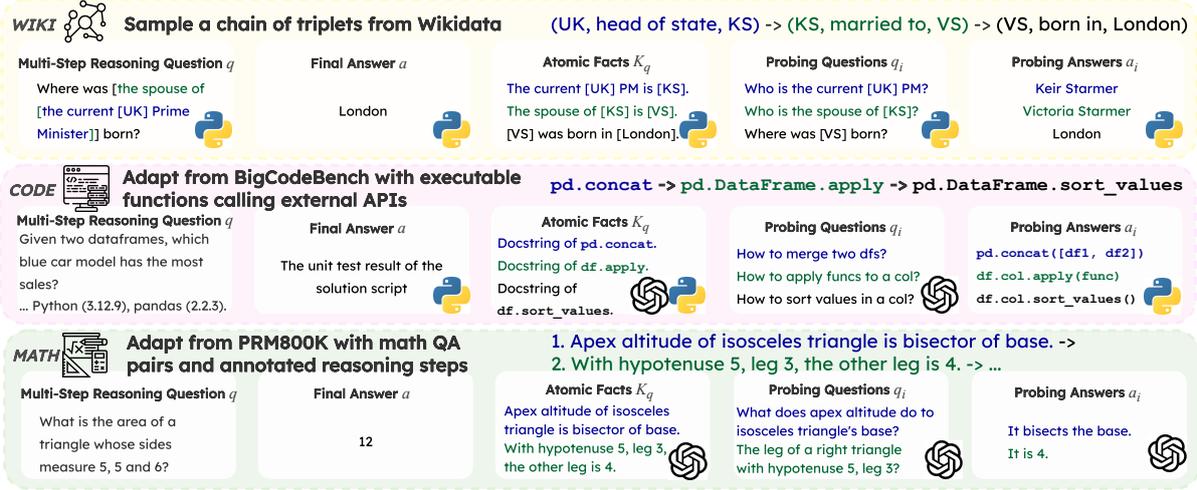


Figure 5: **The data generation pipeline for our three TRACK scenarios.** Each pipeline sources from Wikidata (Vrandečić and Krötzsch, 2014), BigCodeBench (Zhuo et al., 2025), or PRM800K (Lightman et al., 2024), and produces five key components: the multi-step reasoning question  $q$ , final answer  $a$ , atomic facts  $K_q$ , probing questions  $q_i$ , and probing answers  $a_i$ . The generation method for each component is denoted by an icon: for deterministic scripts, for LLM-based generation, and for a hybrid approach. Components with no icon are directly adapted from existing benchmarks.

Property ID	Question Template	Attributive Phrase Template	Cloze Template
P20	Which city did [S] die in?	the city where [S] died	[S] died in the city of [O].
P26	Who is [S] married to?	the person who is married to [S]	[S] is married to [O].
P27	What is the country of citizenship of [S]?	the country of citizenship of [S]	[S] is a citizen of [O].
P35	Who is the current head of state in [S]?	the person who is the current head of state in [S]	The current head of state in [S] is [O].
P36	What is the capital of [S]?	the capital of [S]	The capital of [S] is [O].
P37	What is the official language of [S]?	the official language of [S]	The official language of [S] is [O].
P39	What position is held by [S]?	the position held by [S]	The position held by [S] is [O].
P40	Who is [S]'s child?	the child of [S]	[S]'s child is [O].
P50	Who is the author of [S]?	the author of [S]	The author of [S] is [O].
P69	Which university was [S] educated at?	the university where [S] was educated	The university where [S] was educated is [O].
P106	What kind of work does [S] do?	the field of work of [S]	[S] works in the field of [O].
P108	Which organization is the employer of [S]?	the organization that employs [S]	[S] is employed by the organization [O].
P112	Who founded [S]?	the founder of [S]	[S] was founded by [O].
P136	What type of music does [S] play?	the type of music that [S] plays	The type of music that [S] plays is [O].
P140	Which religion is [S] affiliated with?	the religion that [S] is affiliated with	[S] is affiliated with the religion of [O].
P159	Which city is the headquarter of [S] located in?	the city where the headquarters of [S] is located	The headquarters of [S] is located in the city of [O].
P169	Who is the chief executive officer of [S]?	the chief executive officer of [S]	The chief executive officer of [S] is [O].
P170	Who was [S] created by?	the creator of [S]	[S] was created by [O].
P175	Who performed [S]?	the performer of [S]	[S] was performed by [O].
P176	Which company is [S] produced by?	the company that produced [S]	The company that produced [S] is [O].
P178	Who is the developer of [S]?	the developer of [S]	[S] was developed by [O].
P286	Who is the head coach of [S]?	the head coach of [S]	The head coach of [S] is [O].
P364	What is the original language of [S]?	the original language of [S]	The original language of [S] is [O].
P407	Which language was [S] written in?	the language that [S] was written in	The language that [S] was written in [O].
P413	What position does [S] play?	the position that [S] plays	[S] plays the position of [O].
P449	Who is the original broadcaster of [S]?	the original broadcaster of [S]	The original broadcaster of [S] is [O].
P488	Who is the chairperson of [S]?	the chairperson of [S]	The chairperson of [S] is [O].
P495	Which country was [S] created in?	the country where [S] was created	[S] was created in the country of [O].
P641	Which sport is [S] associated with?	the sport that [S] is associated with	[S] is associated with the sport of [O].
P740	Which city was [S] founded?	the city where [S] was founded	[S] was founded in the city of [O].
P800	What is [S] famous for?	the thing that [S] is famous for	[S] is famous for [O].
P937	Which city did [S] work in?	the city where [S] worked	[S] worked in the city of [O].
P1037	Who is the director of [S]?	the director of [S]	The director of [S] is [O].
P1412	What language does [S] speak?	the language that [S] speaks	[S] speaks the language of [O].

Table 5: Templates used for generating questions and facts in the WIKI scenario. ‘[S]’ is a placeholder for the subject entity’s label, and ‘[O]’ is for the object entity.

## B.2 Scenario CODE

For this scenario, we adapt code problems from the BigCodeBench benchmark (Zhuo et al., 2025). The process involves programmatically identifying the necessary API calls for a given coding problem and then using gpt-5-mini to generate the final data components based on this extracted knowl-

edge. Given a coding problem  $q$  as a complex reasoning question, like “Given two dataframes, find which blue car model has the most sales?”, the model must generate a complete, executable function as its reasoning process. The execution result of the function serves as the final answer  $a$ . The knowledge conflict occurs when the model has

an outdated or incorrect knowledge of a required API call, which serves as an atomic fact  $k_i$  (e.g., the new `pandas.concat` function replaces the old API `pandas.DataFrame.append`); this is probed by asking for its purpose ( $q_i$ ) and expecting the function name and required arguments as the probing answer ( $a_i$ ). We set temperature to 0.7, top P to 0.7, and max tokens to 4096. Note that the coding questions do not contain information that identifies individual people and do not include offensive content.

### Component Generation Pipeline

- **Complex Reasoning Question ( $q$ ) and Final Answer ( $a$ ):** We directly adapt these from BigCodeBench. The problem’s instruction prompt serves as the question  $q$ . The final answer  $a$  is not a static string but is determined by the result of executing the model’s generated code against the problem’s predefined unit tests.
- **Atomic Fact ( $k_i$ ) Extraction and Generalization:** This is a multi-step process.
  1. We first use a custom Python script with an Abstract Syntax Tree (AST) visitor to parse the canonical solution for each problem. This visitor identifies all calls to external libraries (e.g., `pandas`, `numpy`) and resolves their canonical names, handling import aliases.
  2. For each unique canonical function name identified, we programmatically inspect the installed library to retrieve its official function signature and docstring. This information is cached to ensure consistency.
  3. To create a generalized representation of an API call, we substitute the specific variables used in the source code with their formal parameter names from the retrieved signature. For example, two distinct calls in the source code, `pd.concat([df1, df2])` and `pd.concat([data_a, data_b])`, are both normalized to the same underlying knowledge representation corresponding to `pd.concat(objs=...)`. This ensures that the atomic fact represents the general usage of the API, not the specific variables in one instance.
  4. The summarized docstring for the API call is then provided to o4-mini with a structured prompt. The LLM’s task is to generate a final, declarative sentence that describes

the function’s purpose. This sentence becomes the atomic fact  $k_i$ .

- **Probing Question ( $q_i$ ) and Answer ( $a_i$ ):** Using the same LLM call described above, we also prompt the model to generate a natural language question about the API’s purpose ( $q_i$ ). The probing answer ( $a_i$ ) is the generalized, canonical function call itself (e.g., `pandas.concat(objs)`). This tests if a model can recall the correct API call for a given task.

### Detailed Prompts

#### System:

You are a helpful assistant that analyzes code and generates questions about library calls. Your task is to generate probe questions based on the provided code and library calls.

For the user prompt, we replace “[ANSWER]” with the generalized representation of the API, “[DOCSTRING]” with its complete docstring, and “[LIB]” with the library imported in the python script.

#### User:

You are asked to generate two items based on the function [ANSWER]:

1. A **probe question** about the function’s basic usage.
2. A declarative **answer sentence** that resolves the question.

You can refer to the following docstring for context on the function’s purpose:

[DOCSTRING]

– **INSTRUCTIONS**

- 1. For the “question”:
  - It MUST start with “Given the library (libraries) [LIB], how can we ...?”.
  - It MUST be a single sentence.
  - It MUST NOT reveal the function name ‘[ANSWER]’ or its specific arguments.
  - It should describe a goal that leads to the simplest, most basic call of the function.

- If '[ANSWER]' includes specific keyword arguments (e.g., 'func(arg1=val)'), the question must be phrased to necessitate those exact arguments.

- If '[ANSWER]' is a simple call with no keyword arguments (e.g., 'func()'), the question should ask for the standard way to achieve the action.

\*\*2. For the "answer":\*\*

- It MUST be a single, complete, context-independent sentence.

- It MUST combine the premise of the question you just generated with the code snippet '[ANSWER]' to form a factual statement.

- The sentence should state that the action in the question can be accomplished using the provided code.

-

\*\*OUTPUT FORMAT\*\*

You MUST output your response as a valid JSON object and nothing else. Do not add any explanatory text before or after the JSON.

Use the following structure:

```
{
  "question": "The question you generated.",
  "answer": "The answer sentence you generated."
}
```

### B.3 Scenario MATH

For this scenario, we adapt problems from the PRM800K benchmark (Lightman et al., 2024), which contains math problems with annotated, step-by-step ground-truth solutions. The original problem is the complex reasoning question ( $q$ ), and its final numerical or symbolic result is the final answer ( $a$ ). We leverage the annotated reasoning steps to define the required atomic facts. Specifically, we prompt gpt-5-mini in a Socratic manner, where each crucial step—with a conceptual fact, a formula, or an intermediate calculation—is extracted as an atomic fact ( $k_i$ ) and a context-independent probing question ( $q_i$ ). We use the atomic fact  $k_i$  as the probing answer ( $a_i$ ) to verify a model's understanding of that specific reasoning step. Note

that the math questions do not contain information that identifies individual people and do not include offensive content. We set temperature to 0.7, top P to 0.7, and max tokens to 4096.

### Component Generation Pipeline

- **Complex Reasoning Question ( $q$ ) and Final Answer ( $a$ ):** These are taken directly from the ReasonEval dataset. The original math problem serves as the question  $q$ , and its final numerical or symbolic result is the answer  $a$ .
- **Atomic Fact ( $k_i$ ) and Probing Question ( $q_i$ ) Generation:** The core of our process is to deconstruct the ground-truth reasoning steps into a series of atomic, self-contained facts. We achieve this by prompting gpt-4.1-mini in a Socratic manner.
  1. The LLM is provided with the math problem and its complete, step-by-step solution. It is instructed to act as an expert tutor and break down the solution into a chronological sequence of question-knowledge pairs.
  2. The prompt guides the LLM to distinguish between problem-specific steps and general mathematical principles. For a step involving specific numbers from the problem, the generated question must be concrete (e.g., "Given the equations..., what operation eliminates 'e'?"). For a step involving a general theorem, the question must be abstract (e.g., "How can we convert a value from dollars to cents?").
  3. The generated "knowledge" sentence, which is a complete, declarative statement, becomes the atomic fact ( $k_i$ ). The corresponding generated "question" becomes the probing question ( $q_i$ ).
- **Probing Answer ( $a_i$ ):** For the MATH scenario, the probing answer is identical to the atomic fact ( $a_i = k_i$ ). This tests if the model can recall the specific conceptual or procedural knowledge required for a given step when prompted with a context-independent question.

### Detailed Prompts

#### System:

You are an expert AI assistant specializing in educational content creation and the Socratic method. Your primary function is to

deconstruct a given mathematical problem's solution into a series of atomic, abstract, and chronological probing questions.

**Your Goal:** Given a math question and its complete, step-by-step reasoning, you will generate a sequence of question-knowledge pairs that test a user's understanding of the solution process.

**Key Principles to Follow:**

1. **Formulate Probing Questions:** You must adjust the level of abstraction based on the reasoning step for your questions, and it should be complete, self-contained, declarative sentence.

\* **Be CONCRETE** for problem-specific steps: If a step involves translating the problem's specific text or operating on its specific numbers/equations, the question **MUST** refer to that concrete context.

\* **Be ABSTRACT** for general knowledge steps: If a step relies on a general mathematical definition, theorem, or conversion formula that exists outside the specific problem, the question should ask about that general principle.

\* **Be COMPLETE, SELF-CONTAINED, and DECLARATIVE:** The question must be independent to the provided math question, meaning it can be answered without the math question.

\* **DO NOT** ask something like "According to the problem...", which is **NOT SELF-CONTAINED**.

2. **Formulate Complete Knowledge:** The 'knowledge' field **MUST** contain a complete, self-contained, declarative sentence.

\* For **concrete questions**, the knowledge should state the

specific result of the operation.

\* For **abstract questions**, the knowledge should state the general rule or definition.

\* The sentence must be context-independent, meaning it can be understood on its own as a piece of knowledge.

\* **DO NOT** use short phrases as knowledge (e.g., "By finding the prime factorization").

\* **DO NOT** use simple "Yes/No" or short phrases as knowledge.

3. **Chronological Order:** Your questions must follow the logical sequence of the provided reasoning steps.

4. **Strict Formatting:** For each generated item, you **MUST** provide:

\* 'Question:' [The probing question. It should be complete, self-contained, declarative, and independent to the context math question.]

\* 'Knowledge:' [The complete, declarative knowledge sentence. It should be complete, self-contained, declarative, and independent to the context math question.]

5. **Only Necessary Question-Knowledge Pairs:** Provide only question-knowledge pairs that is necessary for solving the math question. [IMPORTANT] It should be as fewer as possible.

**Examples to Learn From:**

**Example 1:**

Math Question: Three pencils and a jumbo eraser cost \$1.24. Five pencils and a jumbo eraser cost \$1.82. No prices include tax. In cents, what is the cost of a pencil?

Reasoning Steps:

- Let's call the price of a pencil  $p$  and the price of a jumbo eraser  $e$ . Then we can write two equations.
- We have  $3p + e = 1.24$  and  $5p + e =$

1.82.

3. To solve this system, let's subtract the first equation from the second equation. This will eliminate  $e$ .

$$4. 5p + e - 3p - e = 1.82 - 1.24.$$

5. This simplifies to  $2p = 0.58$ . So  $p = 0.29$ .

6. That means a pencil costs 29 cents.

Generated Output:

Question: Three pencils and a jumbo eraser cost \$1.24. Five pencils and a jumbo eraser cost \$1.82. No prices include tax. If 'p' is the price of a pencil and 'e' is the price of an eraser, what two equations do we have?

Knowledge: Three pencils and a jumbo eraser cost \$1.24. Five pencils and a jumbo eraser cost \$1.82. No prices include tax. If 'p' is the price of a pencil and 'e' is the price of an eraser, the two equations we have are  $3p + e = 1.24$  and  $5p + e = 1.82$ .

Question: Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , what specific operation will eliminate the variable 'e'?

Knowledge: Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , subtracting the first equation from the second will eliminate the variable 'e'.

Question: After subtracting  $3p + e = 1.24$  from  $5p + e = 1.82$ , what is the resulting value for p?

Knowledge: After subtracting  $3p + e = 1.24$  from  $5p + e = 1.82$ , we will have  $2p = 0.58$ , which solves to  $p = 0.29$ .

Question: How can we convert a monetary value from dollars to cents?

Knowledge: To convert a value from dollars to cents, you multiply the dollar amount by 100.

\*\*Example 2:\*\*

Math Question: Compute  $58_9 - 18_9$ . Express your answer in base 9.

Reasoning Steps:

1. Subtraction works the same in base 9 as in base 10.

2. So we just find the difference of the numbers in the right column, which is  $8 - 8 = 0$ .

3. Now we need to find the difference of the numbers in the left column. This is  $5 - 1 = 4$ .

4. The answer is  $40_9$ .

Generated Output:

Question: Does the standard algorithm for column-based subtraction apply to number systems other than base 10?

Knowledge: The standard algorithm for column-based subtraction is a general method that applies to numbers in any integer base, not just base 10.

Question: Applying column-based subtraction to  $58_9 - 18_9$ , what are the results for the right and left columns respectively?

Knowledge: Applying column-based subtraction to  $58_9 - 18_9$ , the result for the right column is  $8_9 - 8_9 = 0_9$ , and the result for the left column is  $5_9 - 1_9 = 4_9$ .

Question: If the result is  $0_9$  for the right column and  $4_9$  for the left column after applying column-based subtraction in base 9, what is the final answer in base 9?

Knowledge: If the result is  $0_9$  for the right column and  $4_9$  for the left column after applying column-based subtraction in base 9, the final answer in base 9 will be  $40_9$ .

Now, await the user's input.

**User:**

Math Question: {the complex math reasoning question}

Steps:

{annotated reasoning steps}

Generated Output:

## C Experimental Setup

Our experiments are deployed on the RunAI platform. We use AMD EPYC 7543 32-Core Processor as the CPU model and multiple NVIDIA A100-SXM4-80GB as the GPU.

### C.1 Knowledge Probing

#### Hyperparameters

We set the temperature to 0.7 and top P to 0.7. For each probing question, we sample  $M = 10$  responses. The maximum output tokens is set to 4096.

#### Detailed Prompts

The prompts are tailored to each complex reasoning scenario as follows:

##### Scenario WIKI.

**System:**

Answer the question with the name of an entity. Provide only the name of the entity as your answer. Please make an educated guess and always return an entity.

[Here is one demonstration]

User:

Who is the developer of Telegram?

Assistant:

Telegram FZ-LLC

**User:**

User:

{probing question}

Assistant:

##### Scenario CODE.

**System:**

Answer the question with a Python code snippet, which requires ONLY ONE direct function or class constructor call from ONLY ONE library. Provide ONLY ONE function or constructor call itself with correct positional arguments.

- Do NOT include import statements.
- Do NOT include example data, variable assignments, or any other

code.

- For each keyword argument of the function, if the question implies specific keyword arguments, include them in the function call. If the question does not require the keyword argument explicitly or only require it with its default value, the function can be called without this keyword argument.

- Please make an educated guess and always return a function call.

[Here is one demonstration]

User:

Given the library pandas, how can we create a DataFrame by explicitly passing the input data (such as an ndarray, Iterable, dict, or DataFrame) using the 'data' parameter?

Assistant:

```
“python
pandas.DataFrame(data)
“
```

**User:**

User:

{probing question}

Assistant:

##### Scenario MATH.

**System:**

Answer the math question with a concise sentence. Provide only the direct answer to the math question and no more additional reasoning.

[Here is one demonstration]

User:

Given the equations  $3p + e = 1.24$  and  $5p + e = 1.82$ , what specific operation will eliminate the variable 'e'?

Assistant:

Subtracting the first equation from the second will eliminate the variable 'e'.

**User:**

User:  
 {probing question}  
 Assistant:

**C.2 Knowledge Injection****Hyperparameters**

Across all methods, the final reasoning response is generated with a temperature of 0.7, top-p of 0.7, and a maximum of 8192 new tokens. Method-specific hyperparameters are detailed in Tab. 6.

**Detailed Prompts**

We use one-shot Chain-of-Thought (CoT) prompting for all methods. The core difference lies in how the context  $C$  is presented to the model. The Base Model (closed-book) and Append (open-book) methods share a similar prompt structure, differing only in whether the conflicting knowledge is injected into the user prompt. The Append-T method uses the same inference prompt as Append, and FT-CK uses the same prompt as Base Model after fine-tuning.

Below are the prompt templates for the WIKI scenario. The templates for CODE and MATH follow a similar structure with domain-specific examples.

**Closed-book Prompt Template (for Base Model and FT-CK).****System:**

You are given a question. To answer the question, you should think step by step. Use line breaks between steps, but do not use line breaks within each step. The final answer to the question should start with "The answer is ...", and should be placed at the final step. Please make an educated guess and always return an entity.

[Here is one demonstration]

User:

What is the capital of the country where Plainfield Town Hall was created?

Assistant:

1. Given my knowledge, Plainfield

Town Hall was created in the United States of America.

2. Given my knowledge, the capital of United States is Washington, D.C.

3. Therefore, the capital of the country where Plainfield Town Hall was created is Washington, D.C.

4. The answer is Washington, D.C.

**User:**

User:  
 {complex reasoning question}  
 Assistant:

**Open-book Prompt Template (for Append and Append-T):.****System:**

... (the system message is the same as closed-book) ... Users may provide a set of facts or not. If they provide facts that conflict with your knowledge, you should update your knowledge and use the facts to answer the question.

[Here is one demonstration]

User:

Who is the person who is the current head of government of British married to?

Please update your knowledge with the following facts:

The name of the current head of the British government is Keir Starmer.

Assistant:

1. The user provided that Keir Starmer is the current head of government of the British government.

2. I will update my knowledge with the provided fact...

3. Given my knowledge, Keir Starmer is married to Victoria Starmer.

4. Therefore, the person who is the current head of government of British married to is Victoria Starmer.

Method	Hyperparameter	Value
FT-CK	LoRA Rank ( $r$ )	16
	LoRA Alpha ( $\alpha$ )	16
	Target Modules	q_proj, k_proj, v_proj, o_proj
	Learning Rate	2e-4
	Num. Train Epochs	4
MeLLO	Retriever Model	facebook/contriever-msmarco
	Decomposer Temperature	0.01
	Sub-Answer Generation Temperature	0.01

Table 6: Method-specific (FT-CK and MeLLO) hyperparameters used in our knowledge-injection experiments. The Append and Append-T methods do not have specific hyperparameters.

5. The answer is Victoria Starmer.

**User:**

User:  
 {complex reasoning question}  
 Please update your knowledge with  
 the following facts:  
 {context C (atomic facts to inject)}  
 Assistant:

**MeLLO.** The MeLLO method uses a multi-step process involving question decomposition and retrieval. It first generates a series of subquestions, retrieves relevant facts for each, and generates sub-answers:

**System:**

You are a machine that expertly breaks down complex reasoning questions. Your task is to produce only the VERY NEXT single-line subquestion needed to solve the main question. You MUST use the information based on the answers to the subquestions in the history provided (e.g., substitute entities, use previous answers, etc.). Do not add any extra text or explanation.

### EXAMPLES ###

[Input]  
 Main Question:  
 What is the capital city of the  
 country of citizenship of Ivanka

Trump's spouse?  
 History:  
 None  
 [Next Subquestion]  
 Who is Ivanka Trump's spouse?  
 [Input]  
 Main Question:  
 What is the capital city of the  
 country of citizenship of Ivanka  
 Trump's spouse?  
 History:  
 Subquestion:  
 Who is Ivanka Trump's spouse?  
 Generated Answer:  
 Ivanka Trump's spouse is Jared  
 Kushner.  
 [Next Subquestion]  
 What is the country of citizenship  
 of Jared Kushner?  
 [Input]  
 Main Question:  
 What is the capital city of the  
 country of citizenship of Ivanka  
 Trump's spouse?  
 History:  
 Subquestion:  
 Who is Ivanka Trump's spouse?  
 Generated Answer:  
 Ivanka Trump's spouse is Jared  
 Kushner.  
 Subquestion:  
 What is the country of citizenship  
 of Jared Kushner?  
 Generated Answer:  
 Jared Kushner is a citizen of the

United States.  
[Next Subquestion]  
What is the capital city of the United States?

**User:**

[Input]  
Main Question:  
{complex reasoning question}  
History:  
{history subquestions and answers}  
[Next Subquestion]

The final prompt is then constructed from these generated subquestion-answer pairs as follows:

**System:**

... (system message is the same as closed-book) ... Users have decomposed the question to multiple subquestions, and answer them one by one. Based on the provided subquestion-answer pairs, answer the question.

[Here is one demonstration]

User:

Who is the person who is the current head of government of British married to? Please perform reasoning with following subquestion-answer pairs:

Subquestion: Who is the current head of government of British?

Generated Answer: The name of the current head of the British government is Keir Starmer.

Subquestion: Who is the current Prime Minister of the United Kingdom married to?

Generated Answer: Keir Starmer is married to Victoria Starmer.

Given these subquestion-answer pairs, please answer user's question by reasoning step by step.

Assistant:

The subquestion-answer pairs have provided all of the required knowledge:

1. Keir Starmer is the current head of government of the British government.
2. We also know that Keir Starmer is married to Victoria Starmer.
3. Therefore, the person who is the current head of government of British married to is Victoria Starmer.
4. The answer is Victoria Starmer.

**User:**

User:

{complex reasoning question}  
Please perform reasoning with following subquestion-answer pairs:

{subquestion-answer pairs}

Given these subquestion-answer pairs, please answer user's question by reasoning step by step.

Assistant:

## D Discard Samples Without Knowledge Conflicts

We do not discard samples without knowledge conflicts. Keeping them ensures a fixed test set for fair comparison across different models and avoids introducing inductive bias (Si et al., 2023) from pre-training models. In Tab. 7, we report the number of remaining conflicting examples, and in Tab. 8, we further report the performance of each model if we were to remove non-conflict samples. Despite the removal of examples, the relative performance ranking between models and methods remains almost consistent. For instance, even when we compare the performance of GPT-4.1-mini and o4-mini on CODE, which removes the most number of examples, we show that the exclusion of non-conflict samples does not alter our core conclusions.

## E Responsible NLP Research

### E.1 Artifacts

To foster reproducibility and open science, we will make our complete codebase and all reconstructed datasets publicly available under the MIT License. The essential artifacts used in this project, including datasets, backbone models, and major software libraries, are detailed in Tab. 9. It is important to

	Llama-3.2 (1B)	Llama-3.2 (3B)	Llama-3.2 (11B)	Qwen-3 (1.7B)	Qwen-3 (4B)	Qwen-3 (8B)	GPT-4.1-mini & o4-mini
WIKI	479	458	431	498	475	458	363
CODE	317	246	225	262	452	327	119
MATH	361	363	359	390	344	284	221

Table 7: Conflicting example counts are reported for the WIKI, CODE, and MATH domains. As samples without knowledge gaps are removed, the final numbers vary according to each model’s specific knowledge deficiencies.

Backbone Model	Method	Scenario WIKI			Scenario CODE			Scenario MATH		
		% HP	% AP	% FKE	% HP	% AP	% FKE	% HP	% AP	% FKE
Llama-3.2 (1B)	Base Model	0.7 ± 0.7	5.9 ± 2.0	0.7 ± 0.7	3.9 ± 2.1	6.8 ± 2.7	34.5 ± 5.2	13.4 ± 3.5	19.7 ± 3.9	19.4 ± 3.9
	Append	0.9 ± 0.7	7.5 ± 2.3	0.9 ± 0.7	3.9 ± 2.1	5.7 ± 2.5	31.7 ± 5.2	16.1 ± 3.6	22.0 ± 4.0	21.2 ± 4.0
	FT-CK	0.5 ± 0.5	6.6 ± 2.2	0.5 ± 0.5	3.9 ± 2.1	6.8 ± 2.7	30.3 ± 5.4	14.3 ± 3.5	20.9 ± 4.3	19.8 ± 4.0
	MeLLO	1.3 ± 1.0	10.6 ± 2.7	1.4 ± 0.9	3.9 ± 2.1	5.8 ± 2.7	30.1 ± 5.2	12.2 ± 3.3	17.6 ± 3.7	18.0 ± 3.9
Llama-3.2 (3B)	Base Model	1.9 ± 1.2	20.5 ± 3.5	2.1 ± 1.2	13.6 ± 4.3	20.7 ± 5.3	48.0 ± 6.1	32.0 ± 5.0	46.7 ± 4.8	35.8 ± 5.2
	Append	3.3 ± 1.5	19.4 ± 3.5	3.6 ± 1.6	14.6 ± 4.5	20.9 ± 5.5	49.8 ± 6.3	29.8 ± 4.7	43.5 ± 5.0	33.1 ± 5.2
	FT-CK	2.5 ± 1.4	18.6 ± 3.5	3.2 ± 1.6	11.4 ± 3.7	19.7 ± 5.1	49.0 ± 5.9	30.3 ± 4.7	44.2 ± 5.1	33.9 ± 5.0
	MeLLO	2.8 ± 1.5	15.1 ± 3.5	3.4 ± 1.6	13.0 ± 4.1	17.1 ± 4.9	40.9 ± 5.9	28.5 ± 4.8	39.4 ± 4.7	30.9 ± 4.7
Llama-3.2 (11B)	Base Model	4.5 ± 2.0	25.2 ± 4.1	4.5 ± 2.0	19.6 ± 5.3	28.2 ± 6.0	45.3 ± 6.7	27.2 ± 4.6	33.4 ± 4.7	40.8 ± 5.2
	Append	4.8 ± 2.0	26.6 ± 4.1	4.9 ± 2.1	15.1 ± 4.5	23.8 ± 5.6	47.1 ± 6.7	25.9 ± 4.5	34.1 ± 4.9	38.9 ± 4.9
	FT-CK	4.3 ± 1.7	28.2 ± 4.1	4.5 ± 2.0	16.4 ± 4.9	24.9 ± 5.8	44.9 ± 6.2	25.5 ± 4.6	31.3 ± 4.6	41.5 ± 5.3
	MeLLO	3.8 ± 1.7	14.4 ± 3.2	5.0 ± 2.0	7.3 ± 3.3	13.8 ± 4.4	22.9 ± 5.6	15.6 ± 3.6	19.2 ± 3.9	29.2 ± 4.7
Qwen-3 (1.7B)	Base Model	4.1 ± 1.7	16.7 ± 3.4	4.1 ± 1.7	15.5 ± 4.4	23.5 ± 5.2	50.6 ± 5.9	43.5 ± 5.0	56.2 ± 4.9	57.4 ± 4.9
	Append	<b>83.7</b> ± 3.2	90.7 ± 2.5	<b>87.2</b> ± 2.8	17.2 ± 4.6	25.2 ± 5.3	58.6 ± 5.9	43.3 ± 4.9	53.5 ± 5.0	60.3 ± 4.6
	Append-T	6.1 ± 2.1	50.9 ± 4.3	6.1 ± 2.1	13.6 ± 4.0	21.9 ± 5.2	32.4 ± 5.7	6.2 ± 2.3	72.3 ± 4.6	19.6 ± 3.7
	FT-CK	4.3 ± 1.7	16.6 ± 3.3	4.3 ± 1.7	17.4 ± 4.4	22.7 ± 5.2	49.0 ± 5.9	40.0 ± 4.9	51.8 ± 4.9	54.1 ± 4.6
	MeLLO	5.1 ± 1.9	16.2 ± 3.3	5.3 ± 1.9	10.1 ± 3.6	13.9 ± 4.0	46.6 ± 5.7	24.9 ± 4.4	36.8 ± 4.5	29.1 ± 4.5
Qwen-3 (4B)	Base Model	3.9 ± 1.8	22.9 ± 3.8	3.8 ± 1.7	22.2 ± 3.7	34.2 ± 4.3	58.4 ± 4.6	53.5 ± 5.2	67.3 ± 5.1	65.3 ± 5.1
	Append	80.0 ± 3.6	94.2 ± 2.0	82.2 ± 3.5	26.4 ± 3.9	33.8 ± 4.4	69.1 ± 4.3	51.3 ± 5.4	66.6 ± 4.9	66.0 ± 4.9
	Append-T	44.8 ± 4.4	62.1 ± 4.6	45.5 ± 4.8	21.5 ± 3.8	28.9 ± 4.3	40.2 ± 4.5	28.5 ± 4.7	83.7 ± 3.8	37.6 ± 5.1
	FT-CK	5.7 ± 2.1	27.3 ± 4.1	5.5 ± 2.1	23.8 ± 3.9	35.8 ± 4.6	59.2 ± 4.8	53.5 ± 5.2	68.6 ± 4.9	66.1 ± 5.4
Qwen-3 (8B)	Base Model	3.6 ± 1.6	21.1 ± 3.6	3.7 ± 1.7	25.1 ± 4.6	38.7 ± 5.4	58.0 ± 5.0	59.0 ± 5.8	78.3 ± 4.8	60.4 ± 5.8
	Append	77.3 ± 3.7	91.6 ± 2.5	78.5 ± 3.6	25.1 ± 4.6	39.4 ± 5.2	57.5 ± 5.2	60.2 ± 5.6	82.0 ± 4.6	61.3 ± 5.6
	Append-T	72.7 ± 4.1	93.9 ± 2.2	73.5 ± 3.8	23.1 ± 4.7	30.6 ± 4.9	38.4 ± 5.4	60.4 ± 5.8	73.8 ± 5.1	76.1 ± 4.9
	FT-CK	3.4 ± 1.6	22.2 ± 3.6	3.4 ± 1.6	25.7 ± 4.6	37.3 ± 5.2	58.6 ± 5.0	57.2 ± 5.8	78.0 ± 4.8	58.3 ± 5.8
	MeLLO	5.0 ± 2.0	22.9 ± 3.9	5.3 ± 2.1	24.9 ± 4.7	35.0 ± 5.0	56.1 ± 5.7	35.0 ± 5.8	52.5 ± 5.6	35.9 ± 5.3
GPT-4.1-mini & o4-mini	Base Model	7.3 ± 2.6	35.4 ± 4.8	7.3 ± 2.6	34.5 ± 8.4	51.3 ± 9.2	62.6 ± 8.8	69.3 ± 6.1	89.8 ± 3.9	72.0 ± 6.1
	Append	74.8 ± 4.5	<b>95.3</b> ± 2.2	75.5 ± 4.4	39.5 ± 8.4	<b>51.7</b> ± 8.8	67.6 ± 8.0	86.4 ± 4.5	<b>99.3</b> ± 0.7	<b>87.1</b> ± 4.3
	Append-T	82.6 ± 3.9	<b>95.3</b> ± 2.2	82.6 ± 3.9	<b>45.0</b> ± 8.8	50.4 ± 8.4	<b>73.5</b> ± 8.0	<b>86.6</b> ± 4.8	99.3 ± 0.7	86.8 ± 4.5

Table 8: Percentage of Holistic Pass (% HP), Answer Pass (% AP), and Full Knowledge Entailment (% FKE) on the WIKI (multi-hop QA on Wikidata), CODE (code generation with external APIs), and MATH (multi-step mathematical reasoning) scenarios. **We discard all examples without knowledge conflicts** and compare the Base Model (closed-book) against several open-book knowledge injection methods (Append, FT-CK, MeLLO) across Llama-3.2, Qwen-3, and GPT series. We set KAS to 1 in the open-book setting so each question receives only its missing facts. We report 95% confidence intervals (CIs) in the ± sign and **bold** the best scores per column.

note that all data is in English. We have adhered to the intended-use licenses for all artifacts, which permit non-commercial research applications.

## E.2 Usage of Artificial Intelligence Assistants

We use Artificial Intelligence (AI) assistants only to help us with data annotation, basic code completions and grammar checking.

**AI Annotators.** To validate the quality of our TRACK benchmark, we employ two AI assistants (gpt-5-mini and gemini-2.5-pro<sup>11</sup>) to annotate a sample of 50 items per domain for two dis-

tinct tasks. The first task evaluates the factual correctness of probing questions, while the second assesses the necessity of each knowledge step in solving the reasoning questions. Across all domains, gpt-5-mini identified 96.0% of items as factual and 86.7% as necessary. Similarly, gemini-2.5-pro annotated 92.7% as factual and 89.3% as necessary. The average ratio is 94.4% for factuality and 88.0% for necessity, respectively. The high scores demonstrate a high degree of consistency in the validation process. To ensure the AI annotation is reliable, we further sample 50 annotations from both GPT-5-mini and Gemini-2.5-

<sup>11</sup>Released on June 17, 2025.

Artifacts/Models/Packages	Citation	Link	License
<i>Data Artifacts</i>			
Wikidata	(Vrandečić and Krötzsch, 2014)	<a href="https://www.wikidata.org/">https://www.wikidata.org/</a>	Creative Commons BY-SA 4.0 License
BigCodeBench	(Zhuo et al., 2025)	<a href="https://bigcode-bench.github.io/">https://bigcode-bench.github.io/</a>	Apache-2.0 License
PRM800K	(Lightman et al., 2024)	<a href="https://github.com/openai/prm800k/">https://github.com/openai/prm800k/</a>	MIT License
<i>Backbone Models</i>			
LLaMA-3.2 (1B, 3B, 11B)	(Grattafiori et al., 2024)	<a href="https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices">https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices</a>	Llama 3.2 Community License Agreement
Qwen-3 (1.7B, 4B, 8B)	(Yang et al., 2025)	<a href="https://qwen3.org">https://qwen3.org</a>	Apache-2.0 License
GPT-4.1-mini	(OpenAI, 2025a)	<a href="https://openai.com/index/gpt-4-1/">https://openai.com/index/gpt-4-1/</a>	Missing
o4-mini	(OpenAI, 2025b)	<a href="https://openai.com/index/introducing-o3-and-o4-mini/">https://openai.com/index/introducing-o3-and-o4-mini/</a>	Missing
Gemini-2.5-Pro	(Comanici et al., 2025)	<a href="https://deepmind.google/models/gemini/pro/">https://deepmind.google/models/gemini/pro/</a>	Missing
<i>Packages</i>			
PyTorch	(Paszke et al., 2019)	<a href="https://pytorch.org/">https://pytorch.org/</a>	BSD 3-Clause License
transformers	(Wolf et al., 2020)	<a href="https://huggingface.co/docs/transformers/index">https://huggingface.co/docs/transformers/index</a>	Apache-2.0 License
numpy	(Harris et al., 2020)	<a href="https://numpy.org/">https://numpy.org/</a>	BSD License
pandas	(McKinney et al., 2011)	<a href="https://pandas.pydata.org/">https://pandas.pydata.org/</a>	BSD 3-Clause License
matplotlib	(Hunter and Dale, 2007)	<a href="https://matplotlib.org/">https://matplotlib.org/</a>	Python Software Foundation License
seaborn	(Waskom, 2021)	<a href="https://seaborn.pydata.org/">https://seaborn.pydata.org/</a>	BSD 3-Clause License
openai-python	(Achiam et al., 2023)	<a href="https://pypi.org/project/openai/">https://pypi.org/project/openai/</a>	Apache-2.0 License
togetherai	(TogetherAI, 2025)	<a href="https://www.together.ai/">https://www.together.ai/</a>	Apache-2.0 License
Flash Attention 2	(Dao, 2024)	<a href="https://github.com/Dao-AI/flash-attention">https://github.com/Dao-AI/flash-attention</a>	BSD-3-Clause License
amCharts 5	(amCharts, 2025)	<a href="https://www.amcharts.com/docs/v5/">https://www.amcharts.com/docs/v5/</a>	Basic License

Table 9: Data artifacts, backbone models, and major packages utilized in our study. All the reconstructed datasets and the provided code of our project are released under the MIT License to support open science and reproducibility.

Pro. One PhD student who majors in Computer Science from university in the US rigorously annotated these samples using the same prompt sent to LLMs. The annotations of LLMs have high average F1 scores with the human annotator (98.0% for factuality and 93.5% for necessity). Detailed prompts are shown below:

**User (WIKI, Factuality):**

You are given a question and an answer. Is the answer to the question factually correct? Please provide a short sentence as explanation and then answer Yes if the answer is factually correct or No if it is not.

**User (WIKI, Necessity):**

You are given a multi-hop reasoning question. You are also given a list of knowledge steps. Your task is to determine if every piece of knowledge is essential for solving the problem. A step is essential if the final answer cannot be reached without it.

**User (CODE, Factuality):**

You are given a coding question which requires a function call from Python external libraries as the ground truth answer. Is the answer factually correct? Please provide a short sentence as explanation and

then answer Yes if the answer is factually correct or No if it is not.

Question: {probing question}

Answer: {probing answer}

**User (CODE, Necessity):**

You are given a coding question which requires implementing a Python code with multiple required functions from external libraries. Does the problem description explicitly require the use of every function listed below to be considered a valid solution? Answer Yes if all functions are mandated by the prompt's text, or No otherwise.

**User (MATH, Factuality):**

You are given a math question. Is the answer to the question factually correct? Please provide a short sentence as explanation and then answer Yes if the answer is factually correct or No if it is not.

**User (MATH, Necessity):**

You are given a math question and a required solution path. You are also given a list of knowledge steps. Your task is to determine if every

piece of knowledge is essential for solving the problem **by following the required path exactly**.

Are all the listed knowledge steps required to construct the solution **as demanded by the problem's constraints**? An alternative mathematical method is irrelevant if it deviates from the specified path. Provide a short sentence as explanation. Then answer **Yes** if all knowledge is essential to the required method, or **No** if some knowledge is irrelevant or deviates from the method.

**AI Code Completions.** To streamline the development process, we leveraged [GitHub Copilot](#) for assistance. The tool was primarily used to generate routine code, such as inline comments, function header documentation, and boilerplate statements like `if __name__ == "__main__":`. The high-level software architecture and the core logic of all functions were manually designed and implemented by the authors.

**Grammar Checking.** The initial draft of this paper was composed manually. For refinement, we employed a suite of AI-powered writing assistants, including [DeepL](#) for translation, and [ChatGPT](#), [Gemini](#) for grammatical correctness.