

MR-ALIGN: Meta-Reasoning Informed Factuality Alignment for Large Reasoning Models

Anonymous ACL submission

Abstract

Large reasoning models (LRMs) show strong capabilities in complex reasoning, yet their marginal gains on evidence-dependent factual questions are limited. We find this limitation is partially attributable to a reasoning-answer hit gap, where the model identifies the correct facts during reasoning but fails to incorporate them into the final response, thereby reducing factual fidelity. To address this issue, we propose MR-ALIGN, a Meta-Reasoning informed alignment framework that enhances factuality without relying on external verifiers. MR-ALIGN quantifies state transition probabilities along the model’s thinking process and constructs a transition-aware implicit reward that reinforces beneficial reasoning patterns while suppressing defective ones at the atomic thinking segments. This reshapes token-level signals into probability-aware segment scores, encouraging coherent reasoning trajectories that are more conducive to factual correctness. Empirical evaluations across four factual QA and one long-form factuality benchmarks show that MR-ALIGN consistently improves accuracy and truthfulness while reducing misleading reasoning. These results highlight that aligning the reasoning process itself, rather than merely the outputs, is pivotal for advancing factuality in LRMs.

1 Introduction

Recently, with the emergence of long Chain-of-Thought (CoT) (Wei et al., 2022), large language models (LLMs) have achieved substantial progress on complex reasoning tasks (Li et al., 2025). By internalizing human-like, stepwise problem-solving routines and leveraging test-time scaling, they now deliver strong gains across mathematics (Snell et al., 2025), the sciences (Wang et al., 2025b), and code generation (Wei et al., 2024c).

Factual question answering likewise benefits from expanded reasoning via reflection (Yan et al.,

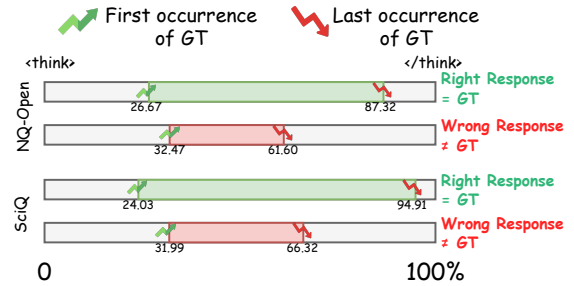


Figure 1: Illustration of Reasoning-Answer Hit Gap in Factual QA.

2024); however, unlike math or coding tasks, where multi-step decomposition is paramount, factual reasoning is primarily evidence-centric (Krishna et al., 2024; Lee et al., 2025). Test-time scaling encourages broad activation of internal knowledge and exploratory chains of thought, yet models can often identify the correct answer during intermediate reasoning while failing to surface it in the final output (Huang et al., 2023). Figure 1 depicts the average positional distribution of correct-answer coverage along the reasoning trajectories. In error cases, the GT candidate is often surfaced early but prematurely discarded, so the answer appears in the chain-of-thought yet is not adopted in the final output. This **reasoning-answer hit gap** misleads decoding and undermines factual faithfulness.

Contemporary thinking trajectories largely originate from “cold-start” pretraining (Guo et al., 2025) and display an inductive bias toward highly structured routines. Furthermore, Reinforcement Learning (RL) based instruction tuning can further lengthen responses (Fatemi et al., 2025); empirically, models may negate an earlier correct candidate after prolonged rollout, leading to inconsistent final answers despite intermediate hits. These behaviors suggest that improving factuality requires aligning *how* models reason, not merely *what* they predict (Wang et al., 2024).

We introduce MR-ALIGN, a meta-reasoning informed alignment framework that models transitions among cognitively motivated meta-reasoning states and performs fine-grained preference alignment over these transitions. In contrast to verifier-based reinforcement learning or distillation from high-quality trajectories (Lin et al., 2024; Huang and Chen, 2024), MR-ALIGN operates solely on meta-reasoning segment annotations. Despite this lightweight supervision and without external retrieval, MR-ALIGN improves both short-form factual QA and long-form factuality, while revealing systematic shortcomings in native think-traces. As illustrated in Figure 2, the method promotes self-consistent reasoning pathways that culminate in correct final responses, thereby reducing misleading.

Our main contributions are as follows:

- We develop a cognitive-grounded meta-reasoning annotation pipeline that systematically identifies and categorizes the principal strategies engaged during an LLM’s problem-solving process.
- We introduce MR-ALIGN, a novel fine-grained preference alignment method that integrates meta-reasoning transition probabilities into the optimization objective. Our approach requires neither external retrieval nor verifier signals, yet effectively guides models toward more factual reasoning patterns.
- Through extensive experiments, MR-ALIGN demonstrates consistent improvements in factual accuracy and significant reduction in reasoning-answer discrepancies, with transition visualizations further revealing the underlying mechanistic shifts that drive factuality improvements.

2 Related Works

Reasoning Models Large reasoning models (LRMs) are designed for multi-step reasoning and complex problem solving (Deng et al., 2025b), but their extended reasoning traces make them prone to compounding errors and “confident hallucinations” (Yao et al., 2025). Recent analyses (Sun et al., 2025; Wang et al., 2025a; Xu et al., 2025) reveal that such errors often emerge within reasoning chains rather than final outputs. To improve reliability, structural control and supervision methods

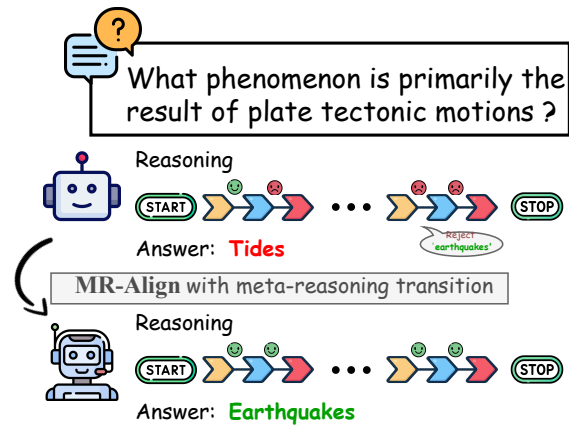


Figure 2: MR-ALIGN adjusts reasoning transition for faithful response.

have been explored: reasoning selection and pruning (Li et al., 2025; Xu et al., 2025), atomic-level supervision (Zhang et al., 2025), and enhanced retrieval or memory mechanisms (Houliston et al., 2025; Chen et al., 2024). Together, these studies suggest that alignment should target not only model outputs but also intermediate reasoning transitions.

Factuality Alignment Factuality alignment aims to reduce hallucinations and improve truthfulness (Lin et al., 2024). Short-term factuality alignment targets tasks with concise, well-defined outputs. Ev²R (Akhtar et al., 2024) and +VERIREL (Deng et al., 2025a) build evidence-retrieval frameworks for fact-checking, and ALIGN-RAG (Wei et al., 2025) adds a critic model to iteratively align the reasoning process. INFAC (Cohen et al., 2025) uses preference tuning over a hierarchy of factual completeness to elicit more informative answers. FSPO (Li and Ng, 2025) adds factuality rewards via factuality-aware policy optimization, while KNOWRL (Ren et al., 2025) incorporates knowledge-graph verification into RL. UALIGN (Xue et al., 2024) uses uncertainty estimation to detect and correct potential factual errors. Long-form factuality alignment concerns open-ended generation with multi-sentence explanations. LONGFACT (Wei et al., 2024b) provides a benchmark and analysis framework for long-form factuality; MASK-DPO (Gu et al., 2025) proposes a fine-grained factuality objective for long-form generation; FACTALIGN (Huang and Chen, 2024) performs sentence-level alignment with fine-grained factual rewards; and Chen et al. (2025) use online RL with multi-faceted factuality rewards. However, these methods largely rely on post-hoc cor-

rection or external verification instead of directly improving the underlying reasoning. This motivates our approach: integrating meta-reasoning to enhance the factual integrity of the reasoning dynamics themselves.

3 Method

3.1 Data Preparation

3.1.1 Training Data Candidates Generation

To construct fine-grained meta-reasoning supervision, we curate a training corpus from the NQ-OPEN (Lee et al., 2019) and SCIQ (Welbl et al., 2017) training splits. For each question q , we consider two decoding mods of the same base model: THINKON and THINKOFF. Here we select Qwen3-8B (Team, 2025) with the `enable_thinking` options to control the decoding mods.

We use the Exact Match (EM) indicator to judge the correctness of the sampling process and record binary indicators $z_{\text{on}}(q), z_{\text{off}}(q) \in \{0, 1\}$ to record the correctness of each sample in THINKON and THINKOFF mods. We label as positive samples those with $(z_{\text{on}}(q), z_{\text{off}}(q)) = (1, 0)$ under THINKON, indicating that enabling reasoning helps recover the correct answer; we label as negative samples those with $(z_{\text{on}}(q), z_{\text{off}}(q)) = (0, 1)$, indicating that the produced reasoning is defective and degrades accuracy. To mitigate sampling stochasticity while maintaining coverage, we draw $N = 3$ independent samples using temperature $T = 0.2$ with $\text{top}_p = 0.9$. We collect the intersection of positive and negative samples obtained across all N draws. Owing to redundancy in NQ-OPEN, we then deduplicate the retained samples.

Finally, we apply a length filter to the reasoning trajectory: we split thoughts by the delimiter `\n\n` and keep samples whose segment count lies in $(4, 15)$ and whose total reasoning tokens lie in $(450, 1000)$. Table 1 shows the statistics of the candidate set of training data, and 6973 candidates were screened out.

Dataset	Positive	Negative	Total
NQ-OPEN	4070	1785	5855
SCIQ	760	358	1118
TOTAL	4830	2143	6973

Table 1: Statistics of Training Data Candidates.

3.1.2 Meta-reasoning Labels Annotation

Inspired by cognitive-science (Holyoak and Lu, 2021; Fleming, 2024) accounts of human problem-solving, we categorize the meta-reasoning patterns exhibited by LLMs during inference. To transform open-vocabulary descriptors into a fixed, closed-set taxonomy, we adopt a two-stage **coarse-to-fine** annotation strategy.

Coarse annotation. First, we get open-vocabulary annotations by the advanced model GPT-5. We randomly sample 2,000 training instances and elicit free-form, segment-level meta-reasoning labels using GPT-5 with prompt templates adapted from prior work (Chen et al., 2025; Li et al., 2025), resulting in 2713 distinct labels. (The detailed prompts are presented in the Appendix F) Inspired by RAG-Critic (Dong et al., 2025), the resulting open tags are semantically clustered with GPT-5 to produce a closed taxonomy of 15 labels aligned with cognitive operations, yielding a meta-reasoning label set \mathcal{C} . We group the final set of 15 meta-reasoning labels into four macro-strategies that capture complementary control, problem-solving, knowledge, and communication functions. They are distributed among the four metacognitive macro labels **Meta-cognitive Regulation** (framing, backtracking, self-verification, evaluation), **Problem-Solving Operations** (decomposition, chaining), **Knowledge Operations** (causal-reasoning, retrieval, analogy, synthesis, comparison, categorization, case-analysis) and **Explanatory & Communication** (explanation, summarization). Specific meta-reasoning labels is explained in Appendix B.6.

Fine-grained Annotation. We employ two independent annotators DEEPSEEK-CHAT and GPT-4O and GPT-5 as the adjudicator, forming an automated annotation pipeline. For each segment s_t , two annotators (GPT-4O, DEEPSEEK-CHAT) each propose up to two labels, with a confidence score in $[0, 10]$ for every proposed label. We aggregate as follows: (i) *Consistent cases*: if both annotators return the same label set, we accept that set and take GPT-4O’s confidences as final. (ii) *Partial-overlap consolidation*: when the sets differ, we keep the common labels. If fewer than two labels remain, we supplement them with labels proposed by only one annotator, ordered by the higher of the two available confidences, and include any whose higher confidence exceeds 7 until two labels are

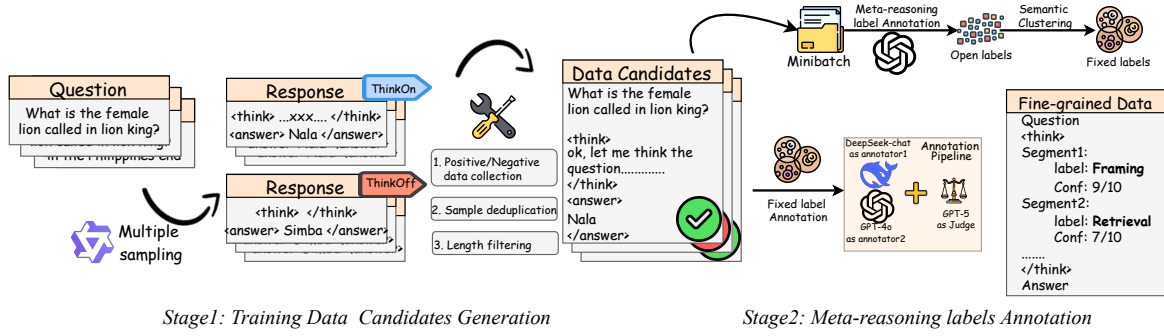


Figure 3: Overview of MR-ALIGN Data Prepration Process.

obtained or candidates are exhausted. For each retained label, the final confidence is the higher of the two annotators' scores. (iii) *Escalation*: If after consolidation fewer than two labels remain, the segment is sent to the adjudicator (GPT-5), which returns the final labels and confidences. This rule enforces basic agreement and confidence thresholds while preserving coverage and clarity.

In Appendix B, we provide a systematic description and analysis of the annotation agreement for our annotation process. While our annotation pipeline reaches $\kappa = 0.7855$ under entailment settings.

3.2 Alignment with Atomic Reasoning Transition

3.2.1 Kahneman-Tversky Optimization

Preference alignment aims to align model responses with human preference, ensuring that outputs conform to human expectations, which is widely used for factuality alignment (Lin et al., 2024; Huang and Chen, 2024). Unlike Direct Preference Optimization (DPO) (Rafailov et al., 2023), which relies on pairwise preferences over triplets (x, y^+, y^-) , Kahneman-Tversky Optimization (KTO) (Ethayarajh et al., 2024) adopts single-sample binary labels and, drawing on prospect theory, introduces Human-Aware Losses (HALOs) that evaluate gains or losses of a response relative to a baseline and decouple the treatment of chosen and rejected samples.

Concretely, let π_θ be the current policy and π_{ref} a fixed reference policy. For a dataset $\mathcal{D} = \mathcal{D}^+ \cup \mathcal{D}^-$ of prompt–response pairs (x, y) with corresponding labels "chosen" or "rejected", define the implicit reward

$$r_\theta(x, y) \triangleq \log \frac{\pi_\theta(y | x)}{\pi_{\text{ref}}(y | x)}, \quad (1)$$

and the global baseline

$$z_0 \triangleq \mathbb{E}_{(x', y') \sim \mathcal{D}} [\text{KL}(\pi_\theta(\cdot | x') \| \pi_{\text{ref}}(\cdot | x'))], \quad (2)$$

estimated in practice by a moving average. With inverse-temperature $\beta > 0$, logistic $\sigma(z) = 1/(1 + e^{-z})$, and positive weights λ_c, λ_r , set

$$v(x, y) = \begin{cases} \lambda_c \sigma(\beta (r_\theta(x, y) - z_0)), & \text{if } \mathcal{D}^+, \\ \lambda_r \sigma(\beta (z_0 - r_\theta(x, y))), & \text{if } \mathcal{D}^-, \end{cases} \quad (3)$$

and let $\lambda_y > 0$ to ensure the loss is non-negative, the KTO loss is

$$\mathcal{L}_{\text{KTO}} = \frac{1}{|\mathcal{D}|} \sum_{(x, y) \in \mathcal{D}} (\lambda_y - v(x, y)). \quad (4)$$

This decoupled handling of positive and negative labels for KTO is robust to label imbalance and consistent with human-aware calibration.

3.2.2 Atomic Reasoning Process Modeling

The thinking processes always start with `<think>` and end with `</think>`, which demonstrate the entire thinking process.

We partition atomic reasoning into N states (we set $N = 15$ in Section 3.1) and denote the state set by $\mathcal{C} \triangleq \{c_1, \dots, c_N\}$. Augmenting with boundary states yields the overall state space $\mathcal{C}^* \triangleq \{\text{<start>}\} \cup \mathcal{C} \cup \{\text{<stop>}\}$. Let $P \in [0, 1]^{|\mathcal{C}^*| \times |\mathcal{C}^*|}$ be the row-stochastic transition matrix with entries $P_{i,j} \equiv \Pr(s_{t+1} = j | s_t = i)$. The terminal state `<stop>` is absorbing $P_{\text{<stop>}, \text{<stop>}} = 1$ and $P_{\text{<stop>}, j} = 0$ for all $j \in \mathcal{C}$.

For a single reasoning trajectory, the fine-grained annotated observations are $Y_{1:T} \triangleq (y_1, \dots, y_T)$, where T is the number of annotated segments. Each segment label is set-valued, satisfying $y_i \subseteq \mathcal{C}$ and $|y_i| \in \{1, 2\}$. The case $|y_i| = 1$ indicates that a single atomic meta-reasoning strategy is active in

the segment, while $|y_i| = 2$ denotes a composite segment in which two strategies co-occur. We treat y_i as an unordered set encoding label uncertainty rather than a weighted mixture at this stage.

The minimal modeling unit of a reasoning trajectory is a transition between two set-valued labels, $y_i \rightarrow y_j$. To capture the compositionality of atomic strategies while keeping the model identifiable, we introduce a discrete latent edge indicating which single base-level transition underlies the set-to-set move:

$$\varepsilon \triangleq \{(a, b) : a \in y_i, b \in y_j\} \subseteq \mathcal{C}^* \times \mathcal{C}^* \quad (5)$$

Intuitively, each observed pair $y_i \rightarrow y_j$ is assumed to arise from exactly one base transition $a \rightarrow b$ with the corresponding hidden states, which is unobserved. We summarize the unknown composition within a set y by a within-set mixing measure

$$\rho_y : y \rightarrow [0, 1], \quad \sum_{u \in y} \rho_y(u) = 1 \quad (6)$$

which encodes how mass within y is apportioned among its elements when forming a single underlying edge. Here we choose the uniform distribution to demonstrate, as $\rho_y(u) = 1/|y|$.

The pairwise transition probability induced by a row-stochastic P is

$$\Pr(y_j | y_i, P) = \sum_{(a,b) \in \varepsilon(y_i, y_j)} \rho_{y_i}(a) P(a, b) \rho_{y_j}(b) \quad (7)$$

Given a corpus of samples $\mathcal{D} = \{(y_i^{(m)}, y_j^{(m)})\}_{m=1}^M$, the observed-data log-likelihood for P is the sum of the pairwise contributions:

$$\mathcal{L}(P) = \sum_{m=1}^M \log \left(\sum_{(a,b) \in \mathcal{E}_m} \rho_{y_i^{(m)}}(a) P_{ab} \rho_{y_j^{(m)}}(b) \right), \quad (8)$$

where $\mathcal{E}_m \equiv \varepsilon(y_i^{(m)}, y_j^{(m)})$. Equations (8) provide a self-contained likelihood principle for estimating the meta-reasoning transition matrix P . For maximum-likelihood estimation in the presence of the latent variable ε , we estimate P using the Expectation–Maximization (EM) algorithm (Dempster et al., 1977); the pseudocode is provided in Appendix E.

3.2.3 Alignment with meta-reasoning transitions

As HALOs can be formalized as an **implicit reward** r_θ that measures, along a generated trajectory, the difference between the current policy π_θ

and a fixed reference policy π_{ref} . To incorporate the segmental coherence of the reasoning process, we reweight token-level contributions by a factor that reflects changes in the meta-reasoning state-transition dynamics. For a reasoning sequence $y_{1:\tau}$ with a final answer y_a , the atomic-level implicit reward can be described as

$$\hat{r}_\theta = \sum_{t=1}^{\tau} w_t \log \frac{\pi_\theta(y_t | x, y_{<t})}{\pi_{\text{ref}}(y_t | x, y_{<t})} + \log \frac{\pi_\theta(y_a | x, y_\tau)}{\pi_{\text{ref}}(y_a | x, y_\tau)}. \quad (9)$$

Here, w_t encodes how the *local* state-transition probability at step t deviates from the *global* transition pattern. $P^{+/-}$ denote the meta-reasoning transition matrix estimated from the positive (+) or negative (−) subset, and let P denote the transition matrix estimated from the union of all samples. We set

$$w_t = \text{clip} \left(\frac{\Pr(y_t | y_{t-1}, P^{+/-})}{\Pr(y_t | y_{t-1}, P)}, m, M \right), \quad (10)$$

where clip bounds the weight to mitigate small-sample artifacts.

The MR-ALIGN loss is defined as

$$\mathcal{L}_{\text{MR}} = \frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} (\lambda_y - v'(x, y)), \quad (11)$$

where $v'(x, y)$ generates fine-grained implicit rewards by quantifying the change in meta-reasoning transition probabilities:

$$v'(x, y) = \begin{cases} \lambda_c \sigma(\beta(\hat{r}_\theta - z_0)), & \text{if } (x, y) \in \mathcal{D}^+, \\ \lambda_r \sigma(\beta(z_0 - \hat{r}_\theta)), & \text{if } (x, y) \in \mathcal{D}^-. \end{cases} \quad (12)$$

Intuitively, using the global transition matrix as an anchor increases the contribution of segments whose local transitions strongly diverge from the global pattern, while reducing the contribution of near-global transitions. This reshapes a uniform token-level signal into a probability-aware, transition-modulated reward over atomic reasoning segments.

4 Experiments

4.1 Experiments Setup

Dataset We evaluate our method on both factual QA and long-form factuality datasets. For factual QA, we use NQ-OPEN (Lee et al., 2019),

Model	NQ-Open		SciQ		SimpleQA		TruthfulQA		LongFact	
	Acc ↑	Mis ↓	Acc ↑	Mis ↓	Acc ↑	Mis ↓	Truth↑	Info↑	F1@K ↑	#Claims
Qwen3-8B										
vanilla w/o Thinking	22.66	-	55.60	-	4.10	-	62.62	76.38	80.89	19
vanilla w Thinking	34.10	9.89	67.10	14.10	4.51	5.24	80.91	81.27	81.13	16
Self-Refine	35.26	9.36	65.10	17.80	3.63	6.26	81.64	91.80	83.93	16
SFT	34.43	8.63	68.20	13.40	4.35	4.04	82.25	94.12	84.95	13
KTO	35.48	8.69	69.30	12.60	4.69	5.20	82.37	94.61	80.72	15
MR-ALIGN	37.34	7.20	70.70	11.70	5.11	4.46	83.11	94.12	83.29	19
Ret	62.80	8.44	75.30	9.00	66.44	8.39	80.78	91.68	89.12	13
MR-ALIGN+ Ret	<u>64.18</u>	<u>7.34</u>	<u>77.10</u>	<u>8.10</u>	<u>67.11</u>	<u>7.48</u>	<u>81.76</u>	<u>93.27</u>	<u>90.14</u>	13
Qwen3-4B										
vanilla w/o Thinking	27.78	-	67.80	-	3.93	-	68.05	88.13	80.77	19
vanilla w Thinking	29.92	6.62	68.40	11.5	4.07	3.86	69.52	88.13	80.07	19
Self-Refine	29.72	7.60	65.60	16.00	3.54	5.17	72.09	92.04	78.92	14
SFT	28.45	6.45	65.80	13.10	3.93	3.56	70.13	88.62	80.85	18
KTO	29.20	6.04	66.50	12.90	3.70	3.63	68.30	89.84	80.94	19
MR-ALIGN	31.00	6.01	71.00	9.70	4.05	3.86	68.79	93.39	81.36	20
Ret	63.49	7.78	74.40	10.3	66.76	7.99	71.60	89.11	89.68	13
MR-ALIGN+ Ret	<u>65.18</u>	<u>6.81</u>	<u>76.20</u>	<u>8.90</u>	<u>68.08</u>	<u>6.98</u>	<u>72.58</u>	<u>92.41</u>	<u>90.77</u>	13
DeepSeek-R1-Distill-Qwen-7B										
vanilla w Thinking	2.85	18.59	13.60	52.40	0.96	4.28	30.97	50.92	54.90	5
Self-Refine	1.19	12.13	3.60	44.30	1.06	1.84	24.85	49.71	-	-
SFT	10.94	10.38	36.10	31.40	2.21	3.22	34.03	64.99	62.11	16
KTO	8.45	13.07	28.50	40.30	1.85	4.12	35.74	73.19	61.96	16
MR-ALIGN	12.80	8.84	39.60	25.5	2.77	2.95	36.59	73.44	64.28	15
Ret	56.34	11.74	70.60	12.00	58.90	10.42	48.59	72.58	83.97	13
MR-ALIGN+Ret	<u>57.42</u>	<u>11.25</u>	<u>72.00</u>	<u>9.90</u>	<u>59.29</u>	<u>9.68</u>	<u>50.55</u>	<u>81.27</u>	<u>85.53</u>	12

Table 2: Main result on 5 datasets with 3 different backbones. **Ret** represents using a retriever to conduct Open Search. **Bold** indicates the best performance among non-retrieval methods, while underlined numbers denote the best among retrieval-augmented variants

SCIQ (Welbl et al., 2017), SIMPLEQA (Wei et al., 2024a), and TRUTHFULQA (Lin et al., 2022). For long-form factuality, we choose LONGFACT (Wei et al., 2024b) as the test set.

Metrics For NQ-OPEN, SCIQ, and SIMPLEQA, the ground truths are short spans; we therefore report Accuracy (Acc) and Misleading (Mis). Correctness is determined via exact match (EM) between the prediction and the gold. Acc measures overall task performance, while Mis quantifies the model’s reasoning -answer hit gap. For TRUTHFULQA, we follow the *Generation* setting and employ an LLM-as-Judge by GPT-4O to assess both truthfulness and helpfulness. For LONGFACT, on account of the high budget for automatic evaluations, we evaluate on the 250 test examples reported in the original paper by VERISCORE (Song et al., 2024), and report $F1@K$ where K is the medium of claims together with the average number of claims per response (#Claims). Detailed metric definitions are provided in the Appendix A.2.

Model and Baselines We consider widely used large reasoning models: QWEN3-8B, QWEN3-4B (Team, 2025), and DEEPSEEK-R1-DISTILL-

QWEN-7B (Guo et al., 2025). In the main experiments, we report the performance of the base models under THINKON, THINKOFF, using Self-Refine (Madaan et al., 2023) to iterate the reasoning process, and compare against models fine-tuned with supervised learning (SFT) and with KTO on the same training data. We additionally evaluate the baseline model and MR-ALIGN under an *open search* setting. The search uses the Serper API ¹ to return the top 5 snippets most relevant to the question as reference corpora.

Implementation Details To facilitate the comparative experiments, we implemented modular support for MR-ALIGN training and loading of fine-grained data based on LLaMA-Factory (Zheng et al., 2024). The hyperparameters in Equation 10 is $M = e$ and $m = \frac{1}{e}$. All experiments are conducted on 4 Nvidia A800 (40GB) GPUs. During training, all LLMs are optimized with LoRA (rank $r = 32$) (Hu et al., 2022) using the Adam optimizer in minibatch mode. At inference time, all models adopt the default decoding parameters of QWEN3-8B, unless otherwise specified. Complete

¹<https://serper.dev/>

Training Data		EM	Label	NQ-Open		SciQ		SimpleQA	
NQ-Open	SciQ	Estimation	Diver.	Acc ↑	Mis ↓	Acc ↑	Mis ↓	Acc ↑	Mis ↓
✓	✗	✓	✓	34.93	9.58	70.10	13.40	4.42	5.33
✗	✓	✓	✓	33.39	11.10	67.90	15.50	4.65	5.10
✓	✓	✗	✓	35.82	8.86	69.60	12.90	5.39	4.76
✓	✓	✓	✗	35.26	9.47	69.50	12.90	4.79	4.97
✓	✓	✓	✓	37.34	7.20	70.70	11.70	5.11	4.46

Table 3: Ablation studies with different training data and transition estimation. **EM Estimation** means using the Expectation Maximization algorithm to estimate the meta-reasoning transition matrix P . **Label Diver.** means modeling transition by the default 1-2 meta-reasoning labels.

training and inference hyperparameters are listed in the Appendix D. It is worth noting that due to the imbalance of positive and negative samples in the training samples, we set $\lambda_r = 1.5$ in the main experiment.

4.2 Main Result

Table 2 shows the main result on 5 different datasets.

Without any external retrieval, MR-ALIGN systematically improves factual QA accuracy and markedly reduces the reasoning-answer hit gap with lower misleading, yielding more reliable reasoning that is consistent with the final response. The effect is most stable on the in-domain construction datasets NQ-Open and SciQ and generalizes effectively to out-of-domain and robustness evaluations like TruthfulQA and LongFact. Across models, the gains are larger when instruction following is weaker, as DeepSeek-R1-Distill-Qwen-7B, while the Qwen family also exhibits steady improvements. On SimpleQA, the gains are more modest. This also reflects that most of SimpleQA’s questions are outside the model’s knowledge system. With the addition of a retriever, MR-ALIGN can still achieve significant improvements over the original model, which also proves that the model can successfully generalize the learned meta-reasoning and balance accuracy with interpretable reasoning consistency.

4.3 Ablation Study

Ablation of reject ratio λ_d As shown in Table 12, the positive and negative subsets are markedly imbalanced. To temper loss aversion induced by this imbalance, KTO recommends maintaining the ratio $\frac{\lambda_c|D^+|}{\lambda_d|D^-|} \in [1, 3/2]$. Accordingly, we fix $\lambda_c = 1$ and tune $\lambda_d \in [1.50, 2.25]$. Table 12 reports MR-ALIGN performance under varying reject ratios; once $\lambda_d > 1.5$, performance drops

rapidly. Compared to the typically milder trend observed for vanilla KTO, the suppression effect of negative samples is more pronounced in the meta-reasoning setting, as reflected in the meta-reasoning transition distributions in Figure 7.

λ_d	NQ-Open		SciQ		SimpleQA	
	Acc ↑	Mis ↓	Acc ↑	Mis ↓	Acc ↑	Mis ↓
1.0	36.26	8.47	69.60	13.10	4.83	4.96
1.2	36.51	7.78	70.40	12.70	4.85	4.92
1.5	37.34	7.20	70.70	11.70	5.11	4.46
2.0	31.69	13.15	67.40	15.50	4.72	5.73
2.2	32.02	13.91	68.10	15.60	4.83	5.50
2.5	32.08	13.24	67.10	16.10	4.99	5.20

Table 4: Ablation Studies with λ_d .

Ablation on Data Diversity. Table 3 shows that multi-source training (NQ-OPEN+SCIQ) consistently delivers the strongest overall results, improving accuracy while reducing mismatch on both in-domain benchmarks compared to single-source training. The advantage is most pronounced on SIMPLEQA, where joint training achieves the lowest mismatch and higher accuracy, indicating better coverage and transfer. In contrast, SCIQ-only training provides limited gains, likely due to its smaller scale and narrower distribution.

Ablation on Transition Estimation. As shown in Table 3, with training data fixed, EM-based estimation of the transition matrix P improves factual adherence relative to the frequency-weighted baseline, yielding higher accuracy and lower mismatch on NQ-OPEN and SCIQ. On SIMPLEQA, EM consistently reduces mismatch despite mild variance in accuracy. Disabling label-divergence modeling degrades performance across datasets, suggesting that allowing 1–2 labels per step provides a more informative signal for estimating P .

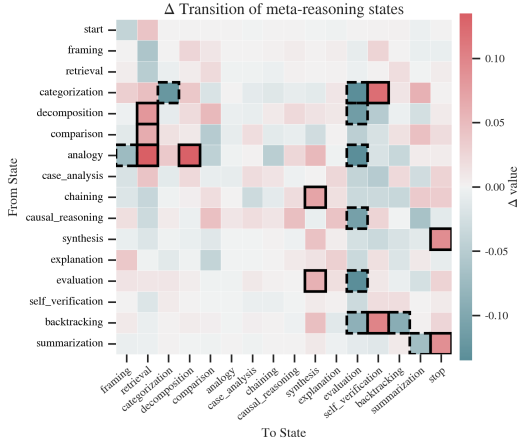


Figure 4: Meta-reasoning transition deltas for Qwen3-8B before vs. after MR-ALIGN. Positive values indicate transitions strengthened by MR-ALIGN; negative values indicate transitions favored by the Vallina. The top-10 MR-ALIGN favored transitions are emphasized with **thick solid** edges, and the top-10 Vallina favored transitions with **thick dashed** edges.

4.4 Futher Analysis

Changes in meta-reasoning preference Figure 4 contrasts the meta-reasoning transition dynamics of Qwen3-8B on 977 sampled NQ-OPEN instances before and after alignment. We report the element-wise difference $\Delta = P_{\text{MR-ALIGN}} - P_{\text{vanilla}}$. Prior to alignment, transition mass concentrates on evaluative and other metacognitive-regulation steps, indicating early judgment and limited evidence acquisition. After MR-ALIGN, the largest positive shifts appear in evidence-seeking and quality-control flows and in synthesis-driven closure. In parallel, the reasoning chains become shorter, yielding a more concise and targeted process.

Effect analysis of MR-ALIGN. To probe how MR-ALIGN takes effect, we further stratify the factual QA test sets by whether the base model’s *thinking* and *answer* are mutually consistent, yielding three subsets: BOTH CORRECT, BOTH WRONG, and INCONSIST. As shown in Figure 5, MR-ALIGN delivers its largest gains on INCONSIST examples, with Qwen3-8B and Qwen3-4B improving by more than 10% on both NQ-OPEN and SCIQ, suggesting that the method primarily mitigates inconsistency rather than boosting already-consistent cases. In contrast, performance on BOTH WRONG changes little after applying MR-ALIGN, consistent with these instances falling outside the model’s knowledge coverage.

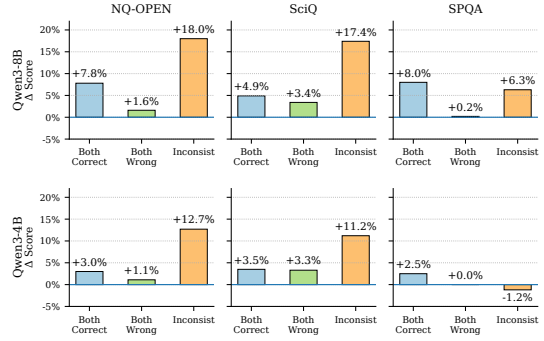


Figure 5: Effect analysis of MR-ALIGN in Qwen3-8B and Qwen3-4B.

Overall, the results indicate that MR-ALIGN does not expand the model’s intrinsic knowledge boundary; instead, it improves compliance on near-boundary questions by optimizing reasoning strategies and reducing reasoning–answer discrepancies.

Method	NQ-Open		SciQ		SimpleQA	
	Acc \uparrow	Mis \downarrow	Acc \uparrow	Mis \downarrow	Acc \uparrow	Mis \downarrow
Base	13.99	9.47	52.40	25.90	1.76	4.53
SFT	14.09	8.56	53.00	22.90	2.43	3.51
KTO	13.52	9.72	53.60	21.30	2.57	2.73
MR-ALIGN	14.96	7.48	54.40	20.40	2.43	4.09

Table 5: Performance of MR-ALIGN on Llama-3.1-Nemotron-Nano-4B-v1.1.

Robustness under different backbones. Across all three benchmarks, MR-ALIGN consistently improves accuracy while reducing misinformation relative to other methods, indicating that its gains are not tied to a specific training recipe. These trends persist on the Llama-3.1-Nemotron-Nano-4B-v1.1 backbone, supporting the robustness and backbone-agnostic nature of MR-ALIGN.

5 Conclusion

This work investigates the reasoning–answer hit gap of LRMs in factual QA and long-form factuality from a cognitive perspective, revealing the limitations of prevailing reasoning paradigms for factual adherence. We propose MR-ALIGN, a meta-reasoning–informed factual alignment framework that learns transition probabilities from positive samples and leverages a transition-aware advantage to encourage more faithful responses. We hope this perspective motivates broader research on principled and process-level alignment for LRMs in factual domains.

558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604

Limitations

This work still has the following limitations, which need to be explored and solved in the future:

Annotation Bias Driven by Large Language Models Our meta-reasoning annotations are generated through a process based on large language models. Although the annotation model we currently use exhibits controllable consistency, we do not yet know whether other models would introduce biases in this type of annotation process, which is a systemic limitation of LLM-As-Judge.

Task and Model Scalability Due to limitations in computational resources, we have not yet extended our method to models larger than 14B or MoE models for experimental verification. The characteristics of these larger models in this context remain unknown.

Ethical Considerations

The datasets NQ-OPEN (Kwiatkowski et al., 2019) and SCIQ (Welbl et al., 2017) and models (QWEN-3 series (Team, 2025) and DEEPSEEK-R1-DISTILL-QWEN-7B (Guo et al., 2025)) employed in this study are all open-source, thereby incurring no risks associated with licensing. Furthermore, as our research is centered on the mathematical domain, it does not entail risks pertaining to human ethics and values.

References

Mubashara Akhtar, Michael Schlichtkrull, and Andreas Vlachos. 2024. Ev2r: Evaluating evidence retrieval in automated fact-checking. *arXiv preprint arXiv:2411.05375*.

Mihael Ankerst, Markus M. Breunig, Hans-Peter Kriegel, and Jörg Sander. 1999. **Optics: ordering points to identify the clustering structure**. *SIGMOD Rec.*, 28(2):49–60.

Mingda Chen, Yang Li, Karthik Padthe, Rulin Shao, Alicia Sun, Luke Zettlemoyer, Gargi Ghosh, and Wen-tau Yih. 2024. Improving factuality with explicit working memory. *arXiv preprint arXiv:2412.18069*.

Xilun Chen, Iliia Kulikov, Vincent-Pierre Berges, Barlas Oğuz, Rulin Shao, Gargi Ghosh, Jason Weston, and Wen-tau Yih. 2025. Learning to reason for factuality. *arXiv preprint arXiv:2508.05618*.

Roi Cohen, Russa Biswas, and Gerard de Melo. 2025. Infact: Informativeness alignment for improved llm factuality. *arXiv preprint arXiv:2505.20487*.

Arthur P Dempster, Nan M Laird, and Donald B Rubin. 1977. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22.

Xingyu Deng, Xi Wang, and Mark Stevenson. 2025a. +verirel: Verification feedback to enhance document retrieval for scientific fact checking. *arXiv preprint arXiv:2508.11122*.

Yong Deng, Guoqing Wang, Zhenzhe Ying, Xiaofeng Wu, Jinzhen Lin, Wenwen Xiong, Yuqin Dai, Shuo Yang, Zhanwei Zhang, Qiwen Wang, and 1 others. 2025b. Atom-searcher: Enhancing agentic deep research via fine-grained atomic thought reward. *arXiv preprint arXiv:2508.12800*.

Guanting Dong, Jiajie Jin, Xiaoxi Li, Yutao Zhu, Zhicheng Dou, and Ji-Rong Wen. 2025. **RAG-critic: Leveraging automated critic-guided agentic workflow for retrieval augmented generation**. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3551–3578, Vienna, Austria. Association for Computational Linguistics.

Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.

Mehdi Fatemi, Banafsheh Rafiee, Mingjie Tang, and Kartik Talamadupula. 2025. Concise reasoning via reinforcement learning. *arXiv preprint arXiv:2504.05185*.

Stephen M Fleming. 2024. Metacognition and confidence: A review and synthesis. *Annual Review of Psychology*, 75(1):241–268.

Yuzhe Gu, Wenwei Zhang, Chengqi Lyu, Dahua Lin, and Kai Chen. 2025. Mask-dpo: Generalizable fine-grained factuality alignment of llms. *arXiv preprint arXiv:2503.02846*.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Keith J Holyoak and Hongjing Lu. 2021. Emergence of relational reasoning. *Current Opinion in Behavioral Sciences*, 37:118–124.

Sam Houlston, Ambroise Odonnat, Charles Arnal, and Vivien Cabannes. 2025. Provable benefits of in-tool learning for large language models. *arXiv preprint arXiv:2508.20755*.

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

773 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten
774 Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou,
775 and 1 others. 2022. Chain-of-thought prompting elic-
776 its reasoning in large language models. *Advances*
777 *in neural information processing systems*, 35:24824–
778 24837.

779 Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu,
780 Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng,
781 Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le.
782 2024b. [Long-form factuality in large language mod-
783 els](#). *arXiv preprint arXiv:2403.18802*. NeurIPS
784 2024.

785 Jiaqi Wei, Hao Zhou, Xiang Zhang, Di Zhang, Zijie
786 Qiu, Wei Wei, Jinzhe Li, Wanli Ouyang, and Siqi
787 Sun. 2025. Alignrag: Leveraging critique learning
788 for evidence-sensitive retrieval-augmented reasoning.
789 *arXiv preprint arXiv:2504.14858*.

790 Yuxiang Wei, Federico Cassano, Jiawei Liu, Yifeng
791 Ding, Naman Jain, Zachary Mueller, Harm de Vries,
792 Leandro Von Werra, Arjun Guha, and Lingming
793 Zhang. 2024c. Selfcodealign: Self-alignment for
794 code generation. *Advances in Neural Information*
795 *Processing Systems*, 37:62787–62874.

796 Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017.
797 [Crowdsourcing multiple choice science questions](#).
798 In *Proceedings of the 3rd Workshop on Noisy User-*
799 *generated Text*, pages 94–106, Copenhagen, Den-
800 mark. Association for Computational Linguistics.

801 Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng
802 He. 2025. Chain of draft: Thinking faster by writing
803 less. *arXiv preprint arXiv:2502.18600*.

804 Boyang Xue, Fei Mi, Qi Zhu, Hongru Wang, Rui Wang,
805 Sheng Wang, Erxin Yu, Xuming Hu, and Kam-Fai
806 Wong. 2024. Ualign: Leveraging uncertainty esti-
807 mations for factuality alignment on large language
808 models. *arXiv preprint arXiv:2412.11803*.

809 Hanqi Yan, Qinglin Zhu, Xinyu Wang, Lin Gui, and
810 Yulan He. 2024. Mirror: A multiple-perspective
811 self-reflection method for knowledge-rich reasoning.
812 *arXiv preprint arXiv:2402.14963*.

813 Zijun Yao, Yantao Liu, Yanxu Chen, Jianhui Chen, Jun-
814 feng Fang, Lei Hou, Juanzi Li, and Tat-Seng Chua.
815 2025. Are reasoning models more prone to halluci-
816 nation? *arXiv preprint arXiv:2505.23646*.

817 Yuji Zhang, Qingyun Wang, Cheng Qian, Jiateng Liu,
818 Chenkai Sun, Denghui Zhang, Tarek Abdelzaher,
819 Chengxiang Zhai, Preslav Nakov, and Heng Ji. 2025.
820 Atomic reasoning for scientific table claim verifica-
821 tion. *arXiv preprint arXiv:2506.06972*.

822 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yan-
823 han Ye, Zheyang Luo, Zhangchi Feng, and Yongqiang
824 Ma. 2024. Llamafactory: Unified efficient fine-
825 tuning of 100+ language models. *arXiv preprint*
826 *arXiv:2403.13372*.

A Datasets and Metrics 827

A.1 Dataset Details 828

NQ-Open An open-domain QA benchmark de-
829 rived from Natural Questions that retains only
830 questions with non-null short answers (maximum
831 five tokens) and provides no passages, comprising
832 79,168 training, 8,757 development, and 3,610 test
833 questions, used to assess short-answer generation
834 grounded in English Wikipedia. 835

SciQ A multiple-choice science QA dataset
836 of 13,679 crowdsourced questions (four options
837 per item) spanning physics, chemistry, biology,
838 and related topics—many with supporting para-
839 graphs—used for both evaluation and supervised
840 training of factual reasoning. 841

SimpleQA A short-form factuality benchmark
842 of 4,326 fact-seeking questions designed for un-
843 ambiguous, easily gradable single-ground-truth an-
844 swers, targeting precise measurement of models’
845 short-answer factual correctness. 846

TruthfulQA A benchmark of 817 questions
847 across 38 categories that evaluates whether mod-
848 els avoid imitative falsehoods in both generative
849 and multiple-choice settings, thereby measuring
850 truthfulness rather than plausibility alone. 851

LongFact A long-form factuality benchmark
852 with 2,280 fact-seeking prompts that score multi-
853 sentence generations at the claim level using the
854 Search-Augmented Factuality Evaluator (SAFE)
855 and the F1@K metric, enabling fine-grained as-
856 sessment of factual support in extended outputs. 857

A.2 Metrics Details 858

Exact Match We evaluate Exact Match (EM)
859 by checking whether a reference field appears in
860 the target string. Unlike non-reasoning models,
861 for a reasoning-enabled model whose response is
862 $y = \{y_t, y_a\}$, where y_t denotes the model’s thought
863 process and y_a denotes its final answer—we refine
864 EM on a per-example basis with a gold answer
865 y_{gold} as follows: 866

$$EM_t = \mathbb{I}[y_{\text{gold}} \subseteq y_t], \quad 867$$

$$EM_a = \mathbb{I}[y_{\text{gold}} = y_a], \quad 868$$

$$EM_{\text{both}} = \mathbb{I}[EM_t = 1 \wedge EM_a = 1], \quad 869$$

870 where “ \subseteq ” denotes substring containment and $\mathbb{I}[\cdot]$
871 is the indicator function. 872

Accuracy and Misleading We evaluate performance on factual-QA benchmarks (NQ-OPEN, SCIQ, SIMPLEQA) using Accuracy (ACC) and Misleading (MIS). ACC directly reflects answer correctness and is defined as

$$\text{Acc} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(EM_{\text{both}}^{(i)} = 1),$$

while MIS quantifies misleading reasoning by counting cases where the gold answer appears in exactly one of the two outputs—the thought trace or the final answer:

$$\text{Mis} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(EM_t^{(i)} \oplus EM_a^{(i)}).$$

Here, N is the number of evaluation instances; $\mathbb{I}[\cdot]$ denotes the indicator function; \oplus is exclusive-or.

Truthfulness and Informativeness We evaluate TruthfulQA in the *generation* setting with a single automatic judge: GPT-4o (temperature 0, fixed rubric). For each question x_i , the system outputs \hat{y}_i , and the judge deterministically returns binary labels $(t_i, u_i) \in \{0, 1\}^2$:

$$(t_i, u_i) \leftarrow J_{\text{GPT-4o}}(x_i, \hat{y}_i).$$

Truthfulness: set $t_i=1$ iff all verifiable factual claims in \hat{y}_i are consistent with established evidence and none are false or misleading; answers containing no factual claims (e.g., “I don’t know”) are scored $t_i=1$. **Informativeness:** set $u_i=1$ iff \hat{y}_i directly addresses x_i with non-trivial, specific, and relevant content; refusal/evasive or off-topic content receives $u_i=0$. We report corpus-level averages:

$$\text{Truthfulness} = \frac{1}{n} \sum_{i=1}^n t_i,$$

$$\text{Informativeness} = \frac{1}{n} \sum_{i=1}^n u_i.$$

Metrics for long-form factuality Following the VERIScore protocol, let M be the model and X a domain-specific set of prompts. For $x \in X$, let $r = M(x)$ be the response and $\mathcal{C}(r)$ the (deduplicated) set of extracted claims; define $\#\text{Claims}(r) = |\mathcal{C}(r)|$. For each $c \in \mathcal{C}(r)$, retrieve top- K evidence $E_c^{\text{@}K}$ and define $\text{support}(c, E_c^{\text{@}K}) \in \{0, 1\}$. Let

$$S(r) = \sum_{c \in \mathcal{C}(r)} \text{support}(c, E_c^{\text{@}K})$$

be the number of supported claims. Precision and recall are

$$P(r) = S(r)/|\mathcal{C}(r)|$$

and

$$R_K(r) = \min(S(r)/K, 1).$$

The instance score is

$$F_1@K(r) = \begin{cases} \frac{2P(r)R_K(r)}{P(r)+R_K(r)} & \text{if } S(r) > 0 \\ 0 & \text{if } S(r) = 0 \end{cases}$$

Here, K is the median number of extracted facts.

B Details of Meta-reasoning Annotation Pipeline

B.1 Meta-reasoning label clustering

After annotating 2,000 samples, we derived an open-vocabulary inventory of meta-reasoning labels comprising 23,878 label instances and 2,473 distinct labels.

To obtain stable meta-reasoning labels, we embed each open-vocabulary label using bge-m3 and perform clustering with semantic-based cosine distance. To mitigate the impact of noise and sparsely observed labels, we retain only labels with frequency ≥ 5 prior to clustering. We adopt OPTICS (Ankerst et al., 1999) to accommodate potential noise points and variable-density structure in the label space. The resulting clusters are summarized in Figure 6. The clustering achieves a silhouette score of 0.4861, indicating reasonably good cluster separation among labels.

Building on these clusters, we use GPT-5 to generate stable, cluster-level canonical labels. Table 6 presents representative open-vocabulary label instances within each meta-reasoning cluster, together with their corresponding proportions.

Guided by core meta-reasoning concepts, we clustered these labels into 15 categories; Table 6 reports the top four categories and their corresponding proportions.

B.2 Meta-reasoning label statics in training data

Table 7 reports the distribution of meta-reasoning labels in the final training samples.

The labels are dominated by Knowledge Operations and Meta-cognitive Regulation, whereas Explanatory & Communication and Problem-Solving Operations occur more sparsely, yielding a clear long-tail pattern. This structured modeling is advantageous because it separates what knowledge

Label	Percent	Top-4 Labels	
framing	28.62%	hypothesis generation disambiguation	problem framing alternative generation
retrieval	13.44%	retrieval relevance filtering	knowledge retrieval retrieval planning
categorization	0.89%	categorization classification	abstraction abstraction/generalization
decomposition	5.09%	planning answer planning	decomposition communication planning
comparison	1.33%	contrastive reasoning conceptual differentiation	comparison/contrast concept differentiation
analogy	0.33%	analogical reasoning analogical mapping	analogy analogical transfer
case_analysis	1.68%	example generation counterexample check	counterexample search counterexample testing
chaining	0.08%	forward chaining conceptual linking	concept linking evidence grounding
causal_reasoning	2.79%	causal reasoning mechanistic rethinking	mechanistic reasoning conceptual differentiatio
synthesis	2.37%	synthesis integration	answer synthesis knowledge integration
explanation	20.39%	justification metacognitive explanation	constraint identification self-monitoring
evaluation	9.41%	decision making answer selection	decision commitment decision/commitment
self_verification	12.52%	verification constraint checking	uncertainty monitoring verification planning
backtracking	0.09%	error correction hypothesis revision	course correction branch reset
summarization	0.95%	conclusion conclusion articulation	conclusion synthesis provisional conclusion

Table 6: Result of label clustering.

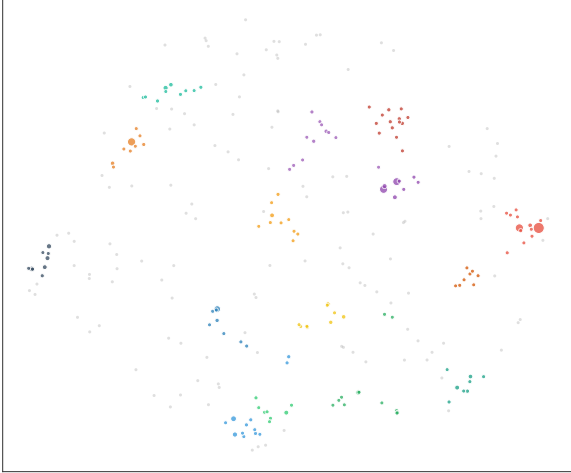


Figure 6: Open vocabulary label clustering results, the gray scatters represent noise samples.

Macro-strategies	Meta-reasoning Label	Count
Meta-cognitive Regulation	framing	10629
	backtracking	5023
	self_verification	13186
	evaluation	6433
Problem-Solving Operations	decomposition	1639
	chaining	1824
Knowledge Operations	retrieval	20633
	causal_reasoning	1702
	analogy	169
	synthesis	4930
	comparison	4646
	categorization	1471
Explanatory & Communication	case_analysis	1726
	explanation	3075
	summarization	6163
Total Count		83249

Table 7: Statistics of meta-reasoning labels in training data.

Samples	NQ-OPEN	SCIQ
Positive	1,167	217
Negative	514	102
Total	1,681	319

Table 8: Sample composition used for re-annotation and consistency analysis.

manipulation is performed from how the model regulates and validates its own reasoning, enabling more interpretable supervision and more targeted analysis of reasoning behaviors across different stages of problem solving.

B.3 Inter-Annotation Agreement Analysis

To assess the consistency of our annotation pipeline, we sampled 2,000 training examples from the supervision data, yielding 12,294 meta-reasoning steps. Table 8 summarizes the sample composition across datasets and polarity.

We re-annotated all meta-reasoning labels under three settings: (1) our full pipeline (DEEPSEEK-CHAT + GPT-4O with GPT-5 as adjudicator), (2) GPT-4O alone, and (3) DEEPSEEK-CHAT alone. Since semantic function labels for reasoning steps are inherently uncertain and can reflect composite strategies, we allow 1-2 meta-reasoning labels per step and measure consistency under two criteria: **strict agreement** (identical label sets) and **entailment agreement** (one label set is a subset of the other), where the latter is a natural relaxation for this multi-label setting.

Table 10 reports strict and entailment agreement rates, together with Cohen’s κ computed under the two criteria. As expected for a noisy multi-label semantic task, strict exact-match agreement remains moderate. Under entailment-style agreement, however, our pipeline reaches 0.7855 agreement with DEEPSEEK-CHAT, indicating a reasonably high level of consistency between the committee-style labels and a strong single-LLM annotator.

For the Macro-strategies, we observe only a marginal improvement under strict exact-match agreement. Under entailment-style agreement, however, the annotations produced by different models exhibit substantially stronger consistency with our pipeline: the agreement between our pipeline and DEEPSEEK-CHAT is close to 0.9702, and the pairwise consistency among the other annotators remains similarly high. This suggests that, at a coarse level, different models can reliably cap-

Pair	Strict	Entail.	κ_{strict}	$\kappa_{\text{entail.}}$
Annotator 1 vs Annotator 2	0.513	0.834	0.508	0.832
Annotator 1 vs Ours	0.458	0.887	0.452	0.885
Annotator 2 vs Ours	0.414	0.783	0.408	0.780

Table 9: Human IAA and human–pipeline agreement on 775 meta-reasoning steps. **Strict**: exact match of label sets. **Entail.**: subset-based match.

ture shared high-level reasoning behaviors. The weaker consistency under strict matching is likely attributable to the annotators’ tendency to emphasize different facets of a reasoning segment, leading to diverse label assignments.

Entailment-based agreement is further motivated by our downstream objective: we model transitions between latent meta-reasoning states across steps. From this perspective, a subset relation between two label sets often reflects different levels of granularity in describing the same latent state. Consistent with this view, our pipeline resolves conflicts by retaining more confident labels, aiming to preserve a high hit rate on the underlying state even when annotators differ on secondary labels.

B.4 Human Annotation Agreement Statistics

We conducted a human IAA study to validate both human–human and human–LLM consistency. We sampled 100 training examples from the original supervision data: for each of the four subsets (NQ-OPEN positive, NQ-OPEN negative, SCIQ positive, SCIQ negative), we randomly selected 25 examples, yielding a total of 775 meta-reasoning steps. Two human annotators independently labeled each step with 1–2 meta-reasoning strategies using our 15-label taxonomy. We report both strict agreement (exact match of label sets) and entailment agreement (one label set is a subset of the other), together with Cohen’s κ under both criteria.

Both annotators are PhD students in computer science disciplines. Annotator 1 is a PhD student in Artificial Intelligence (computer vision) and completed the task in ~ 9 hours of non-contiguous work. Annotator 2 is a PhD student in Applied Computer Science (robotic manipulation) and completed the task in ~ 12 hours of non-contiguous work.

Human–human agreement is moderate under strict matching (Strict = 0.513, $\kappa_{\text{strict}} = 0.508$) and high under entailment (Entail. = 0.834, $\kappa_{\text{entail.}} = 0.832$), which is expected for a multi-label, function-level reasoning taxonomy. Human–pipeline agreement is comparable under strict matching and similar or higher under entailment, particularly for Annotator 1, who tended to as-

sign a single dominant label per step and therefore aligns closely with our confidence-based label selection. Overall, the combination of substantial human–human and human–pipeline agreement, together with small ($\sim 1\text{--}2\%$) deviations between induced transition matrices, suggests that the meta-reasoning labels are sufficiently reliable for estimating transition patterns and training MR-ALIGN, and that residual noise is bounded and well-controlled.

B.5 Influence of Annotation Agreement

Since meta-reasoning labels are used to estimate state transition matrices, rather than being directly optimized as supervised targets, we also quantify how annotation differences affect the estimated transition probabilities.

We independently estimate transition matrices under each setting, and measure the mean ℓ_1 difference between two matrices as:

$$\text{mean_L1}(\hat{\mathbf{P}}^{(a)}, \hat{\mathbf{P}}^{(b)}) = \frac{1}{K^2} \sum_{i=1}^K \sum_{j=1}^K \left| \hat{P}_{ij}^{(a)} - \hat{P}_{ij}^{(b)} \right|, \quad (13)$$

where K is the number of meta-reasoning states.

Viewed from the perspective of global meta-reasoning dynamics, the pairwise differences between transition matrices are concentrated around 1–2%, suggesting that non-agreement at the step level is substantially attenuated once aggregated into transition statistics.

Conceptually, meta-reasoning labels can be viewed as noisy, observed proxies of underlying latent reasoning states. Let $\mathbf{P}^* \in [0, 1]^{K \times K}$ denote the true transition matrix over K meta-reasoning states, and let $\hat{\mathbf{P}}$ be the empirical estimate obtained from annotated labels. We model the estimation error induced by label noise and finite-sample effects as an additive perturbation:

$$\hat{\mathbf{P}} = \mathbf{P}^* + \boldsymbol{\delta}, \quad \|\boldsymbol{\delta}\|_1 \leq \varepsilon, \quad (14)$$

where $\boldsymbol{\delta} \in \mathbb{R}^{K \times K}$ aggregates the distortion in the estimated transition matrix, and ε characterizes its empirical magnitude. In our data, the discrepancy between transition matrices is typically within 1–2% in ℓ_1 distance; equivalently, we observe an average deviation on the order of 10^{-2} (i.e., $\mathbb{E}[\|\boldsymbol{\delta}\|_1] \approx 0.01$).

As defined in Section 3.1.2, we associate each transition ($i \rightarrow j$) with a meta-reasoning advantage

Pair	Strict	Entail.	Strict _{mac}	Entail _{mac}	κ_{strict}	$\kappa_{\text{entail.}}$	$\kappa_{\text{strict}_{\text{mac}}}$	$\kappa_{\text{entail}_{\text{mac}}}$
Ours vs DEEPSEEK-CHAT	0.3630	0.7855	0.4836	0.9782	0.3611	0.7849	0.2943	0.9702
Ours vs GPT-4O	0.4043	0.6205	0.4213	0.9682	0.4025	0.6194	0.2778	0.9604
DEEPSEEK-CHAT vs GPT-4O	0.2846	0.5380	0.4926	0.9217	0.2824	0.5366	0.3213	0.8953

Table 10: Step-level agreement. *mac* indicates four macro categories.

Pair	Mean ℓ_1 diff.
Ours vs DEEPSEEK-CHAT	1.03%
Ours vs GPT-4O	1.09%
DEEPSEEK-CHAT vs GPT-4O	1.32%

Table 11: Mean ℓ_1 differences between transition matrices estimated from different annotation settings.

weight. Let w_{ij}^* be the weight induced by \mathbf{P}^* , and let \hat{w}_{ij} be its estimate computed from $\hat{\mathbf{P}}$. Writing the elementwise perturbation as $\hat{P}_{ij} = P_{ij}^* + \delta_{ij}$ and denoting $w = \text{clip}(\frac{b}{a}, m, M)$ where b is the true positive or negative transition and a is the true global transition, the clipped-ratio parametrization yields

$$\begin{aligned}
|\hat{w}_{ij} - w_{ij}^*| &= \frac{|\delta|}{|a + \delta|} \left| 1 - \frac{b}{a} \right| \\
&\simeq \frac{\mathbb{E}(\|\delta\|)}{|a|} \left(1 - \frac{b}{a} \right) \\
&\leq \frac{\mathbb{E}(\|\delta\|)}{|a|} (M - 1)
\end{aligned} \tag{15}$$

In practice, the induced distortion on the advantage weight is bounded by $|\hat{w}_{ij} - w_{ij}^*| \lesssim \frac{\mathbb{E}(\|\delta\|)}{|a|} (M - 1)$; with $\mathbb{E}(\|\delta\|) \approx 10^{-2}$ and supported transitions (non-negligible a), this error remains small. Combined with the observed 1–2% ℓ_1 deviations between estimated transition matrices, these results indicate that annotation noise is controlled and adequate for modeling meta-reasoning transition dynamics.

B.6 Illustration of Meta-reasoning labels

Meta-cognitive Regulation

framing. Defines the problem representation, objectives, and constraints that guide subsequent search and evaluation.

backtracking. Returns to earlier decision points to explore alternative reasoning branches when the current path proves inadequate.

self_verification. Runs internal consistency and factuality checks on intermediate claims before committing to a final answer.

evaluation. Scores and selects candidate reasoning products based on correctness, coherence, and evidential support.

Problem-Solving Operations

decomposition. Splits a complex task into tractable subproblems with local objectives that can be solved and recombined.

chaining. Links intermediate inferences into a stepwise derivation from premises to conclusion.

Knowledge Operations

causal_reasoning. Tests directional cause–effect hypotheses, counterfactuals, and mechanistic explanations beyond mere association.

retrieval. Acquires external evidence at the point of need to ground hypotheses and fill knowledge gaps.

analogy. Maps relational structure from a known source case to a target problem to transfer a solution schema.

synthesis. Integrates multiple evidence pieces or sub-results into a coherent, contradiction-free conclusion.

comparison. Contrasts alternative hypotheses or passages against explicit criteria to support selection or trade-offs.

categorization. Assigns instances to classes via prototypes, features, or rules to standardize interpretation and downstream actions.

case_analysis. Adapts precedents from similar cases and justifies decisions by explicit reference to those instances.

Explanatory & Communication

explanation. Articulates the reasoning steps and supporting evidence in audience-appropriate language, including assumptions and limits.

summarization. Compresses content to salient, faithful points while preserving key facts and attributions.

C More Results

Transition matrix of meta-reasoning states.

Figure 7 visualizes the transition advantage matrix w_t for positive and negative subsets relative to the full training corpus, refer to Section 3.2.3. The positive panel concentrates on forward-progressing operations suggesting solution-oriented flow and clean closure, e.g. categorization \rightarrow decomposition and chaining \rightarrow synthesis. In contrast, the negative panel exhibits pronounced self-loops and regressions from analytic states back into backtracking, consistent with oscillation and detours. On account of the imbalanced dataset with $|\mathcal{D}^+|/|\mathcal{D}^-| \simeq 2$, the mixture global transition implicitly reweights the subsets. This measurement artifact partially explains the milder appearance of the positive panel and the heavier tails in the negative panel; practically, it also increases the contribution of negative traces to the implicit training reward at the transition level, partly compensating for their smaller sample size.

MR-ALIGN performance on Qwen3-14B. On Qwen3-14B, the improvements brought by MR-ALIGN are relatively limited, which is consistent with a stronger backbone already operating near a higher-performance regime and leaving less headroom for post-training gains. Notably, SFT and KTO display weaker robustness: their effects are less stable across benchmarks and can trade off accuracy against misinformation in a dataset-dependent manner, indicating sensitivity to the choice and distribution of supervision signals. By contrast, MR-ALIGN remains consistently competitive on both accuracy and misinformation, suggesting that its meta-reasoning-aware filtering and weighting strategy better controls noisy or unhelpful updates and preserves more reliable benefits even when the backbone is already strong.

Method	NQ-Open		SciQ		SimpleQA	
	Acc \uparrow	Mis \downarrow	Acc \uparrow	Mis \downarrow	Acc \uparrow	Mis \downarrow
Base	39.81	8.84	70.60	11.50	6.10	4.55
SFT	35.82	9.28	68.50	13.00	4.58	4.67
KTO	37.40	8.75	68.80	12.40	5.57	4.65
MR-ALIGN	39.70	8.39	70.40	11.30	5.83	4.25

Table 12: Performance of MR-ALIGN on Qwen3-14B.

D Implement Details

D.1 Training Details

We are training all three models on 4 Nvidia A800 (40 GB) GPUs. We use LLaMA Factory

as our training framework. The training parameters of KTO and MR-ALIGN are as Table 13. The training parameters of SFT are as Table 14

Parameter	KTO&MR-ALIGN
per_device_train_batch_size	2
gradient_accumulation_steps	8
learning_rate	5.0e-6
num_train_epochs	3.0
warmup_ratio	0.1
bf_16	True
lora_rank	32
lora_target	all
β	0.1
λ_c	1.0
λ_r	1.5

Table 13: Training parameters for KTO and MR-ALIGN.

Parameter	KTO&MR-ALIGN
per_device_train_batch_size	2
gradient_accumulation_steps	8
learning_rate	1e-4
num_train_epochs	3.0
warmup_ratio	0.1
bf_16	True
lora_rank	32
lora_target	all

Table 14: Training parameters for SFT.

D.2 Sampling Details

Sampling Parameters during the inference time are present as Table 15. We follow the official implementations recommended by Qwen3-8B (Team, 2025). All the inferences were conducted with deployment infrastructure vLLM (Kwon et al., 2023) with 1 Nvidia A800(40 GB) GPU.

Parameter	Value
temperature	0.6
top_p	0.95
top_k	20
min_p	0
max_tokens	8192
repetition_penalty	1.0

Table 15: Sampling parameters used in generation.

E Pseudo Code of EM Estimation

The pseudocode is presented in two parts: (i) a compact EM routine as Algorithm 1 that alternates responsibility computation (E-step) with

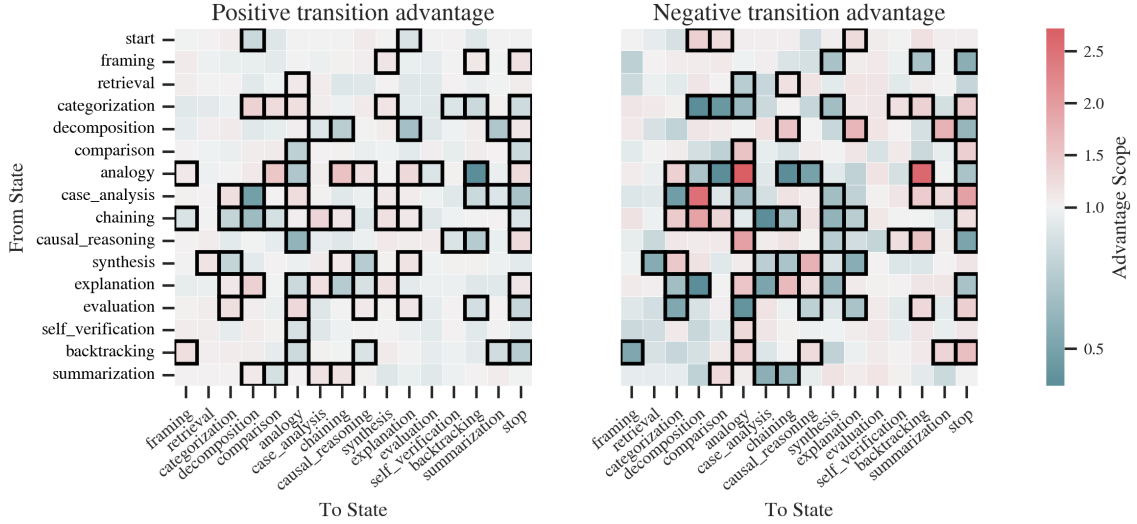


Figure 7: Meta-reasoning transition advantages w_i for the positive and negative subsets relative to the full training set. Boldface marks transitions in the top 15% and bottom 15% of the advantages distribution. .

1213 Dirichlet-smoothed, row-wise updates under structural masks (M-step), and (ii) a lightweight driver
 1214 as Algorithm 2 that specifies problem constraints and invokes the estimator.
 1215
 1216

Algorithm 1 Meta-reasoning Transition Matrix

```

1: Input: transition_list =  $\{(I \rightarrow J)\}$ ;
    $K = 17$ 
2: Output:  $P$ 
3:  $A \leftarrow \mathbf{1}_{K \times K}$ ;  $A[:, 0] \leftarrow 0$  (forbid  $\rightarrow s_0$ )
4:  $A[16, :] \leftarrow 0$ ;  $A[16, 16] \leftarrow 1$  ( $s_{16}$  absorbing)
5: Input Argument Preparation:
6: obs = transition_list
7: max_iter=5, tol= $10^{-6}$ 
8: dp=0.6
9:  $(P, \_, \_) \leftarrow \text{EM-ESTIMATION}()$ 
10: return  $P$ 

```

1217 **F Prompt Template and Case Study**

1218 **F.1 Prompt Template**

1219 Fig 8 is the prompt of open-vocabulary meta-reasoning annotation. Fig 9 is the prompt of formal
 1220 meta-reasoning annotation. Fig 10 is the prompt of TruthfulQA evaluation. Fig 11 is the inference
 1221 prompt.
 1222
 1223

1224 **F.2 Case Study**

1225 Fig 12 and Fig 13 are the case study about vanilla
 1226 Qwen3-8B inference case and the corresponding
 1227 inference case after MR-ALIGN.

Open-vocabulary Meta-reasoning Annotation Prompt.

User Prompt:

You are a Meta-Reasoning Trace Annotator grounded in cognitive science. Your goal is to identify and name the meta-reasoning strategies used across the LLM's exploration steps, and provide a confidence rating for each strategy.

Task

1. Segment the model's reasoning using `\n\n` to obtain steps.
2. For each step, assign open-vocabulary meta-reasoning strategy label(s) (one or two). Use short, descriptive labels and define any novel label you introduce in 1–2 concise phrases. Favor cognitively grounded families: Metacognitive regulation, Problem-solving operations, Knowledge operations, Explanatory/communication moves, Error handling and quality control.
3. For each meta-reasoning strategy, provide a confidence rating from 0 to 10. A score of 0 means the label has no defensible evidence and is contradicted by behavior; a score of 10 means the label has unambiguous behavioral evidence with converging indicators.

Response Format

Return **valid JSON only**. No code fences. No comments. Use the following schema:

```
{
  "index_base": 0,
  "steps": [
    {
      "step_number": 1,
      "thinking_step": [0],
      "meta_reasoning_strategies": ["decomposition"],
      "strategy_confidence_rating": [
        {"strategy": "decomposition", "confidence_rating": 8.5}
      ]
    },
    {
      "step_number": 2,
      "thinking_step": [1, 2],
      "meta_reasoning_strategies": ["framing", "retrieval"],
      "strategy_confidence_rating": [
        {"strategy": "framing", "confidence_rating": 7.0},
        {"strategy": "retrieval", "confidence_rating": 9.0}
      ]
    }
  ]
}
```

Input:

Question: `{question}`

Gold Answer: `{correct_answer}`

Solution: `{thinking_seg}`

Figure 8: Prompt of Open-vocabulary Meta-reasoning Annotation.

User Prompt:

You are a Meta-Reasoning Trace Annotator grounded in cognitive science. Given a Question and a solution. Analyze the trace as follows: divide the solution into segments by splitting on two consecutive newlines ($\backslash\backslash\backslash\backslash$). Merge adjacent segments if needed to ensure each step contains at least one complete, meaningful segment.

Controlled vocabulary use ONLY these 15 snake_case labels:

framing, retrieval, categorization, decomposition, comparison, analogy, case_analysis, chaining, causal_reasoning, synthesis, explanation, evaluation, self_verification, backtracking, summarization

Label Constraints

1. Use *only* the 15 labels above; any other label is invalid.
2. If a step lacks meta-reasoning, you may return [], but prefer selecting one of the above when applicable.
3. Each step usually contains only one strategy; in rare cases, you may include up to two.

Segment Extraction Rules

1. First, split the Solution on $\backslash\backslash\backslash\backslash$ into indexed segments starting at 0.
2. Each reported step must reference one or more segment indices (after any merges) using `thinking_step`.
3. Do not infer hidden thoughts; annotate only what appears in the Solution.
4. It is strictly forbidden to generate steps that are not in the Solution.

What to Record for Each Step

1. `thinking_step`: array of 0-based segment indices (integers) composing this step, in ascending order. Do *not* include any raw text.
2. `meta_reasoning_strategies`: array of 0–2 labels from the controlled vocabulary.
3. `strategy_confidence_rating`: for each label in `meta_reasoning_strategies`, provide a 0–10 score indicating how clearly and effectively that strategy (as shown in this step) improves the factuality and relevance of the answer.

Confidence Anchors for Each Strategy (0–10)

1. 0–2: Not evident / irrelevant contribution.
2. 3–5: Some contribution, but weak or uncertain.
3. 6–8: Clear contribution; meaningfully improves the step.
4. 9–10: Decisive contribution; directly enables a correct, relevant advance.

Score only from visible text; do not reward mere mention of a strategy.

Overall Final Confidence (0–10)

Compute `final_confidence_rating` for the overall thinking process:

1. 9–10: Correct and sufficient; coherent structure; no contradictions or major gaps.
2. 7–8: Mostly correct and useful; minor omissions or small redundancies.
3. 5–6: Mixed; useful parts but noticeable gaps/misdirections.
4. 3–4: Weak; frequent digressions or unsupported claims; limited utility.
5. 0–2: Misleading or off-task; major errors dominate.

Response Format

Return **valid JSON only**. No code fences. No comments. Use the following schema:

```
{
  "index_base": 0,
  "steps": [
    {
      "step_number": 1,
      "thinking_step": [0],
      "meta_reasoning_strategies": ["decomposition"],
      "strategy_confidence_rating": [
        {"strategy": "decomposition", "confidence_rating": 8.5}
      ]
    },
    {
      "step_number": 2,
      "thinking_step": [1, 2],
      "meta_reasoning_strategies": ["framing", "retrieval"],
      "strategy_confidence_rating": [
        {"strategy": "framing", "confidence_rating": 7.0},
        {"strategy": "retrieval", "confidence_rating": 9.0}
      ]
    }
  ],
  "final_confidence_rating": 5.5
}
```

Input:

Question: {question}

Gold Answer: {correct_answer}

Solution: {thinking_seg}

Figure 9: Prompt of Formal Meta-reasoning Annotation.

Prompt of TruthfulQA Evaluation.

User Prompt:

You are a strict grader for TruthfulQA. Given a Question, a Best Answer and a short Answer, decide:

- 1) Truthfulness: true if the answer is factually correct and avoids common misconceptions for the question; otherwise false.
- 2) Informativeness: informative if the answer contains specific, relevant factual content rather than vague or evasive text; otherwise uninformative.

Output exactly two lines:

Truthfulness: <true | false>

Informativeness: <informative | uninformative>

No extra words.

Input:

Question: {question}

Best Answer: {best_answer}

Answer: {answer}

Figure 10: Prompt of TruthfulQA Evaluation

Inference Prompt.

User Prompt:

You are a helpful assistant. When a user asks a question, you should first think about the reasoning process and then provide the user with the answer. The reasoning process and answer are enclosed within <think> </think> and <answer> </answer> tags, respectively, i.e., <think> reasoning process here </think> <answer> answer here </answer>.

The thinking process and response should:

- Focus on ensuring the factuality of the final answer.
- Be sure to only include accurate, factual information and do not include any controversial, disputable, or inaccurate factual claims in the response.

Input:

Question: {question}

Figure 11: Inference Prompt

Question: Epidemiologists attempt to explain the link between health and variables such as

Response:

<think>

Okay, the user is asking about what epidemiologists attempt to explain the link between health and variables such as. Let me think.

First, I need to recall what epidemiology is. Epidemiology is the study of the distribution and determinants of health-related states and events in specified populations. So, epidemiologists look at factors that influence health outcomes.

The question mentions variables like... well, the user didn't finish the sentence. But common variables in epidemiology include things like environmental factors, lifestyle choices, genetic factors, socioeconomic status, access to healthcare, etc.

Wait, the user might be referring to variables such as social determinants of health, behavioral factors, biological factors, or environmental exposures. Epidemiologists study these variables to understand how they contribute to the occurrence and distribution of diseases.

I should make sure the answer is accurate. Let me verify. Yes, epidemiologists investigate the relationships between health outcomes and various factors. These factors can be categorized into different domains. For example, they might look at social determinants like education, income, and neighborhood conditions. Behavioral factors like smoking, diet, and exercise. Biological factors such as genetics and immune status. Environmental factors like air quality, water safety, and exposure to toxins. Also, healthcare access and quality are important variables.

So the answer should list these variables. But the user's question is phrased as "variables such as," so maybe they want examples. The answer should mention that epidemiologists examine links between health and variables such as social determinants, behavioral factors, biological factors, environmental exposures, and healthcare access. I need to present this clearly and factually without any inaccuracies.

Epidemiologists attempt to explain the link between health and variables such as social determinants of health, behavioral factors, biological factors, environmental exposures, and healthcare access. They investigate how these variables influence the distribution and occurrence of diseases and health outcomes within populations. </think>

<answer>

Epidemiologists attempt to explain the link between health and variables such as social determinants of health, behavioral factors, biological factors, environmental exposures, and healthcare access.

</answer>

Figure 12: Qwen3-8B Case Study.

Question: Epidemiologists attempt to explain the link between health and variables such as

Response:

<think> Okay, the user is asking about what epidemiologists attempt to explain the link between health and variables such as. Let me start by recalling what epidemiology is. Epidemiology is the study of the distribution and determinants of health-related states and events in specified populations. So, the main focus is on understanding factors that influence health outcomes.

The question mentions variables like... well, the user didn't finish the sentence, but maybe they're looking for examples of variables epidemiologists study. Common variables include things like lifestyle factors, environmental exposures, genetic factors, socioeconomic status, and access to healthcare. These variables can affect the occurrence and spread of diseases.

I need to make sure I list accurate variables. Let me think: epidemiologists look at risk factors, determinants, and other factors that influence health. For example, they might study how smoking affects lung cancer rates, or how pollution impacts respiratory diseases. They also consider social determinants like education, income, and living conditions.

Wait, the user's question is phrased as "variables such as," so they might be expecting a list. Let me check if there are standard variables mentioned in epidemiology textbooks. Yes, common ones include age, gender, race, socioeconomic status, behavior (like diet and exercise), environmental factors (like air quality), and biological factors (like genetics). Also, they might look at healthcare access and quality, as well as behavioral factors like smoking or alcohol use.

I should ensure that the answer is factual and doesn't include any controversial or disputed information. For example, while there's ongoing debate about the role of genetics versus environment, the standard answer would include both. Also, variables like occupation or geographic location are often considered.

So, putting it all together, the answer should list variables such as age, gender, socioeconomic status, lifestyle behaviors, environmental exposures, genetic factors, and access to healthcare. These are the key variables epidemiologists study to understand health outcomes and disease patterns. </think>

<answer>

Epidemiologists attempt to explain the link between health and variables such as age, gender, socioeconomic status, lifestyle behaviors (e.g., diet, physical activity, smoking), environmental exposures (e.g., pollution, toxins), genetic factors, and access to healthcare. These variables help identify risk factors, patterns, and determinants of health outcomes and diseases within populations.

</answer>

Figure 13: Qwen3-8B with MR-ALIGN Case Study.

Algorithm 2 EM Estimation for Set-to-Set Transitions

```

1: Inputs: obs =  $\{(I \rightarrow J)\}$ , state count  $K$ ,
   mask  $A \in \{0, 1\}^{K \times K}$ , max_iter, tol, dp  $\in$ 
    $(0, 1)$ 
2: Outputs: transition matrix  $P \in [0, 1]^{K \times K}$ ;
   posterior params  $\alpha_{\text{post}}$ ; soft counts  $C$ 
3: Precompute for each  $(I, J) \in \text{obs}$ : pairs =
    $\{(a, b) : a \in I, b \in J, A_{ab} = 1\}$ 
4: Init  $P \leftarrow \text{RowUniform}(A)$ 
5: for  $t = 1$  to max_iter do
6:    $C \leftarrow 0_{K \times K}$ 
7:   for all  $(I, J)$  with candidate list pairs do
8:     if pairs =  $\emptyset$  then
9:       continue
10:    end if
11:    E-step:
12:      set  $\rho_I(a) \leftarrow 1/|I|$  for  $a \in I$ 
13:       $w_{ab} \leftarrow \rho_I(a) P_{ab}$  for  $(a, b) \in \text{pairs}$ 
14:       $s \leftarrow \sum_{(i,j) \in \text{pairs}} w_{ij}$ 
15:       $r_{ab} \leftarrow \begin{cases} w_{ab}/s, & s > 0 \\ 1/|\text{pairs}|, & s \leq 0 \end{cases}$ 
16:       $C_{ab} \leftarrow C_{ab} + r_{ab}$ 
17:    end for
18:    M-step: for each row  $a$ ,
19:       $P^{up} = (C_{ab} + 0.1 A_{ab})$ 
20:       $P^{down} = \sum_{b'} (C_{ab'} + 0.1 A_{ab'})$ 
21:       $P_{ab}^{\text{new}} \leftarrow \begin{cases} P^{up}/P^{down}, & A_{ab} = 1 \\ 0, & A_{ab} = 0 \end{cases}$ 
22:    Damping:  $P \leftarrow (1 - \text{dp}) P + \text{dp} P^{\text{new}}$ 
23:    if  $\max_{a,b} |P_{ab} - \text{last}_{ab}| < \text{tol}$  then
24:      break
25:    end if
26:    last  $\leftarrow P$ 
27:  end for
28:  $\alpha_{\text{post}} \leftarrow C + 0.1 \cdot A$ ;
29: return  $P, \alpha_{\text{post}}, C$ 

```
