

LaRA: Layer-wise Representation Analysis for Detecting Data Contamination in RL Post-Training

Minju Gwak¹ Minseo Kwak¹ Dongseok Lee¹

Guijin Son² Alan Ritter³ Jaehyung Kim¹

¹Yonsei University ²Seoul National University ³Georgia Institute of Technology
mjgwak@yonsei.ac.kr, jaehyungk@yonsei.ac.kr

Abstract

Reinforcement learning (RL) post-training has shown to improve reasoning in large language models (LLMs). However, there has been little exploration on the problem of data contamination in RL post-training, potentially undermining generalization and evaluation reliability of the training process itself. Existing detection methods primarily rely on output-level signals such as likelihood or entropy, which become unreliable for RL-trained models since RL shapes behavior through trajectory-level rewards rather than token likelihoods. We propose LaRA, a layer-wise representation analysis framework for detecting contamination in RL post-trained LLMs. LaRA introduces three complementary metrics, measuring perturbation sensitivity, directional collapse, and local representation rigidity under controlled perturbations. We find that contamination produces progressive geometric deviations across layers, including amplified perturbation sensitivity, stronger directional collapse, and enhanced local rigidity. Based on our findings, we also develop a contamination detection protocol that aggregates representation-level deviations across layers and metrics. Experiments on RL-trained reasoning models show that our protocol outperforms existing output-level baselines for contamination detection.¹

1 Introduction

Reinforcement learning (RL) has shown its effectiveness in training Large Language Models (LLMs) for complex reasoning tasks (Guo et al., 2025; Guha et al., 2025; Li et al., 2025b; Hochlehnert et al., 2025). However, it also raises a critical but underexplored issue of *data contamination* in RL post-training (Tao et al., 2025; Wang et al., 2025; Wu et al., 2026), the inclusion of evaluation or benchmark samples within the RL training data.

¹The code is available at <https://github.com/talzoomanzoo/lara>.




Figure 1: **Output vs. Representation-level detection.** Output-level signals are sensitive to output miscalibration, such as overconfidence of superficially plausible tokens. In contrast, representation geometry provides more robust and interpretable contamination signals.

Contaminated samples can induce reward-driven overfitting and implicit memorization, undermining generalization and evaluation reliability.

Prior work on data contamination in LLMs has mostly focused on pre-training or supervised fine-tuning (SFT) stages (Zhang et al., 2024; Shi et al., 2023; Xie et al., 2024), where memorization is typically characterized by higher token likelihoods or lower entropy (Gonen et al., 2023). Consequently, existing approaches primarily rely on output-level signals derived from model likelihoods or generation statistics. Recent work has extended this paradigm to reasoning trajectories for detecting data contamination in RL, using entropy or behavioral divergence across generation stages as contamination signals (Tao et al., 2025).

However, such output-level statistics can be unreliable due to the poor calibration of LLM output distributions, as shown in Figure 2 (Leng et al., 2025; Xiao et al., 2025). Moreover, unlike pre-training or SFT, RL optimizes expected reward over entire reasoning trajectories rather than token-wise likelihoods, making likelihood-based behavioral signals less directly aligned with the underlying training objective. These challenges motivate a shift toward representation-level analysis, where memorization can be probed directly in the model’s internal geom-

etry, bypassing the calibration issues and objective mismatch that confound output-level signals.

We propose **LaRA**, a **L**ayerwise **R**epresentation **A**nalysis framework for detecting data contamination in RL post-training. Our key hypothesis is that RL-induced memorization produces abnormal representation responses under controlled perturbations: memorized samples become overly stable to semantically equivalent variations, yet exhibit disproportionately large representation shifts when memorized information is removed. To test this, we construct structural control groups of semantically similar questions, apply consistent information masking, and analyze layer-wise representation dynamics across perturbations.

Specifically, our metrics are motivated by three complementary questions about memorized representations. First, does removing key information perturb memorized representations more strongly than non-memorized ones? To examine this question *Representation Shift Magnitude (RSM)* measures how strongly representations change when important information is removed, capturing perturbation sensitivity. Second, do the resulting changes become concentrated along shared dominant directions? *Directional Collapse (DC)* measures whether representation changes collapse toward shared dominant directions, indicating reduced representational diversity. Third, do memorized representations become locally rigid across semantically equivalent variations? *Representation Stability Index (RSI)* quantifies how invariant representations remain across semantically similar variants, capturing local rigidity under meaning-preserving perturbations. Together, these metrics characterize distinct geometric signatures of RL-induced memorization.

Across multiple RL-trained models, we empirically show that contaminated samples exhibit consistent geometric abnormalities compared to non-trained samples. In particular, contaminated samples exhibit abnormal directional collapse, higher local representational rigidity, and greater sensitivity to information removal. Furthermore, our LaRA-based contamination detection score consistently outperforms output-level baselines, suggesting that representation geometry provides a more reliable signal of RL-induced memorization.

In summary, our contributions are as follows:

- We are the first to propose a representation-level framework as well as a training and evaluation

setup for detecting contamination in RL post-training via stiffness and rigidity.

- We introduce a contamination-detection protocol that consistently outperforms output-level baselines across RL-trained models, achieving up to +9.6% AUC improvement and $3.5\times$ higher TPR@FPR=5% compared to the strongest prior output-level method.
- We provide empirical insights into how RL training affects representation geometry across layers.

2 Related Work

Data Contamination. Data contamination detection (Golchin and Surdeanu, 2024, 2025; Deng et al., 2024) is commonly formulated as a membership inference attack (MIA) problem (Wu and Cao, 2025), where contaminated samples are identified through behavioral differences between training and non-training data. Existing methods primarily exploit output-level statistics (Gonen et al., 2023; Xie et al., 2024; Zhang et al., 2024; Shi et al., 2023; Kwak and Kim, 2026; Tao et al., 2025). While these signals are strong indicators of memorization under likelihood-maximization training (*i.e.*, pre-training and SFT), they become unreliable for RL-trained models, since RL optimizes models through reward-driven exploration of reasoning trajectories rather than token-level likelihoods. Contamination detection specifically for RL post-training, however, remains underexplored: existing attempts largely transfer the same output-level signals, *e.g.*, entropy-based detection (Tao et al., 2025). Consequently, they inherit the limitations above, often compounded by exploration dynamics.

Representation Dynamics in LLMs. Recent work has increasingly leveraged representation dynamics in LLMs to study behaviors beyond outputs (Kang et al., 2025; Lee et al., 2025; Gwak et al., 2025; Zhao et al., 2025). One line of work analyzes internal states and their evolution across layers to characterize properties emerging during post-training (Bi et al., 2026; Wang et al., 2024; Hao et al., 2024; Li et al., 2025a). Another line shows that semantic and behavioral attributes are encoded in hidden representations, where specific directions can be exploited to steer, detect, or modulate model behavior (Turner et al., 2023; Lee et al., 2024; Li et al.; Roh et al., 2026; Wurgajt et al., 2026). Closer to our setting, internal representations have also been used for contamination

analysis: Kernel Divergence Score (Choi et al., 2025) quantifies contamination by measuring how fine-tuning on a benchmark dataset changes the similarity structure of sample embeddings. Neural Breadcrumbs (Makhija et al., 2025) extracts features from hidden states and attention patterns, while G-Drift MIA (Ranjan et al., 2026) measures representation changes induced by a targeted gradient update. Both methods demonstrate that internal model states contain informative membership signals, but train supervised classifiers using labeled member and non-member examples. In contrast, LaRA directly derives an instance-level contamination score from layer-wise representation geometry without training an auxiliary classifier or updating the target model. This training-free design isolates the informativeness of the representation signal itself and enables direct application to RL post-trained LLMs.

3 LaRA: Layer-wise Representation Analysis to Detect RL Contamination

We frame the problem of detecting data contamination during RL post-training as Membership Inference Attack (MIA). Given an RL-trained model $\mathcal{M}$ and a candidate sample x , our goal is to determine membership $\mathcal{F}(\mathcal{M}, x) \in \{0, 1\}$, where $\mathcal{F}(\mathcal{M}, x) = 1$ indicates that x was a member of the training dataset and therefore indicates contamination, while 0 indicates otherwise. The central question motivating our analyses is: *do layer-wise representation signals behave differently between member and non-member samples?* To answer this, we introduce three complementary metrics.

3.1 Contamination Dataset Construction

To explore MIA in RL training setting, we construct controlled contamination benchmarks that support a two-stage analysis: (i) detecting contamination in released open-source RL checkpoints based on their known training data, and (ii) tracking how detection signals evolve under additional RL training that we perform on a controlled corpus.

Evaluation set. We construct a contamination evaluation set from the publicly open dataset of the three open-source RL-trained models (EURUS-2-7B-PRIME (Eurus) (Cui et al., 2025), LIMR (Li et al., 2025b), and Olmo-3.1-7B-RL-Zero-Math (Olmo) (Olmo et al., 2025)). For each model, we sample 30 Olympiad-level mathematics problems from its

own RL training set as members, and pair them with 30 problems from AIME 2026 (Balunović et al., 2025) as non-members. This yields a balanced 60-sample evaluation set per model. Non-member split is shared across all three models, while the member split is model-specific.

Training set. To study how contamination signals vary during continued RL training, we re-use each model’s 30 member samples as deliberate contamination targets and augment them with 970 Olympiad-level problems drawn from the RL-MIA (Tao et al., 2025) Math dataset, yielding a 1,000-sample training corpus per model. Using this data, we resume RL training on each open-source checkpoint and track how member vs. non-member signals diverge during RL post-training. Further details on datasets are provided in Appendix A.

3.2 Three Metrics for Analysis

Metric 1: Representation Shift Magnitude. To quantify how strongly a model’s internal representation responds to the removal of important information, we introduce Representation Shift Magnitude (RSM). Given an original question q_0 , we construct a set of semantically similar questions

$$\mathcal{Q} = \{q_0, q_1, \dots, q_K\},$$

where K denotes the number of generated semantic neighbors excluding the original question. For each question $q_i \in \mathcal{Q}$, we apply an importance-based blanking operator BLANKIMPORTANT that removes key information spans while preserving the overall question structure:

$$q_i^- \leftarrow \text{BLANKIMPORTANT}(q_i, k),$$

where k denotes the number of inserted [BLANK] tokens. Refer to Appendix J for details of the perturbation construction process. Let $h_\ell(\cdot)$ denote the mean-pooled hidden representation extracted from transformer layer ℓ , where $\ell \in \mathcal{L} = \{0, 1, \dots, L-1\}$. For each ℓ , we extract hidden representations $u_i = h_\ell(q_i)$ and $w_i = h_\ell(q_i^-)$, where $u_i, w_i \in \mathbb{R}^d$ and d is the hidden representation dimension. We then compute the perturbation-induced representation shift Δ_i and define its magnitude S_i as:

$$S_i = \|\Delta_i\|_2, \quad \Delta_i = u_i - w_i$$

where $\|\cdot\|_2$ denotes the Euclidean norm. To capture how anomalously the original question responds to




Figure 2: **Overview of LaRA, the proposed layer-wise representation geometry analysis framework.** Given an original question q_0 , we generate semantically similar questions and remove shared key information to construct perturbed variants. Hidden representations are extracted across all transformer layers for original, perturbed, and paraphrased inputs. We then compute three complementary geometric metrics: Representation Shift Magnitude (RSM), Directional Collapse (DC), and Representation Stability Index (RSI), which characterize perturbation sensitivity, directional organization, and local representation variability under controlled perturbations.

perturbation relative to its semantic neighbors, we standardize its shift magnitude using the mean and standard deviation of the similar-question set:

$$RSM_\ell = \frac{S_0 - \mu_S}{\sigma_S + \epsilon},$$

where

$$\mu_S = \frac{1}{K} \sum_{i=1}^K S_i, \quad \sigma_S = \sqrt{\frac{1}{K-1} \sum_{i=1}^K (S_i - \mu_S)^2},$$

and $\epsilon > 0$ is a numerical stability constant. A high RSM_ℓ indicates that the original question exhibits a larger representation shift under information removal compared to semantically similar questions, suggesting stronger perturbation sensitivity from memorization and risk of contamination.

Metric 2: Directional Collapse. We introduce Directional Collapse (DC) to characterize the directional organization of perturbation-induced representation changes. We first compute the average perturbation direction across similar questions:

$$\bar{s}_\ell = \frac{1}{K} \sum_{i=1}^K \Delta_i,$$

where $\bar{s}_\ell \in \mathbb{R}^d$ represents the average perturbation direction shared across the semantic group. DC is

then defined as:

$$DC_\ell = \frac{\Delta_0^\top \bar{s}_\ell}{(\|\Delta_0\|_2 + \epsilon)(\|\bar{s}_\ell\|_2 + \epsilon)}.$$

This quantity measures the cosine alignment between the original perturbation direction and the average perturbation direction of semantically similar questions. High DC_ℓ values indicate that perturbation responses are strongly aligned along a shared low-dimensional direction, whereas lower values indicate more distributed or heterogeneous perturbation dynamics.

Metric 3: Representation Stability Index. Finally, we measure local representation stability under semantically preserving perturbations through the Representation Stability Index (RSI). For each perturbed question q_i^- , we generate M paraphrastic variants while preserving the blank positions:

$$\{v_{i,1}, \dots, v_{i,M}\} \sim \text{VARIANTGEN}(q_i^-).$$

We then extract their hidden representations:

$$\phi_{i,m} = h_\ell(v_{i,m}),$$

where $\phi_{i,m} \in \mathbb{R}^d$. Next, we compute the local representation centroid:

$$\bar{\phi}_i = \frac{1}{M} \sum_{m=1}^M \phi_{i,m},$$




Figure 3: **Layer-wise representation geometry patterns of RL post-trained model.** We compare layer-wise representation geometry between contaminated and clean samples under input perturbations. Contaminated samples consistently exhibit deviating perturbation sensitivity (**RSM**), abnormal directional concentration dynamics (**DC**), and altered local representation variability patterns (**RSI**) across transformer layers, indicating that memorized samples form distinct and less robust internal representation structures compared to clean samples.

and define its average deviation:

$$R_i = \frac{1}{M} \sum_{m=1}^M \|\phi_{i,m} - \bar{\phi}_i\|_2.$$

We then standardize the original question’s local variability relative to its semantic neighbors:

$$RSI_\ell = \frac{R_0 - \mu_R}{\sigma_R + \epsilon},$$

where

$$\mu_R = \frac{1}{K} \sum_{i=1}^K R_i, \sigma_R = \sqrt{\frac{1}{K-1} \sum_{i=1}^K (R_i - \mu_R)^2}.$$

A high RSI_ℓ indicates that the original question exhibits larger local representation variability relative to semantically similar questions under paraphrastic perturbations, while lower values indicate more locally stable representation behavior.

3.3 Layer-wise Analysis with Three Metrics

Figure 3 shows the representation geometry patterns measured by the three metrics in Section 3.2. Contaminated samples consistently exhibit larger perturbation-induced representation shifts (RSM) than clean samples across most layers, while clean samples remain near zero throughout depth. In particular, contaminated samples sharply deviate around layers 7–9, indicating substantially higher sensitivity to targeted information removal and stronger dependence on memorized information. DC results further show that contaminated samples

exhibit distinct directional concentration dynamics compared to clean samples. RSI results show that contaminated samples exhibit lower local representation variability, particularly in early layers, indicating more rigid and invariant local representation geometry under paraphrastic perturbations. Additional results are provided in Appendix E.

4 Contamination Detection Protocol

Motivated by Section 3.3, we formulate contamination detection as a layer-aware representation anomaly detection problem. We find that contaminated samples exhibit distinct representation profiles across depth, including amplified perturbation sensitivity, abnormal directional concentration dynamics, and local variability under controlled perturbations. Consequently, contamination should be characterized through *deviation from clean geometric profiles* across multiple metrics and layers, rather than from isolated layer-wise statistics.

Step 1: Clean-reference Robust Standardization. Let $\mathcal{M} = \{\text{RSM}, \text{DC}, \text{RSI}\}$ denote the set of representation geometry metrics, $\mathcal{L}$ denote the set of probed transformer layers, and $m_\ell(x)$ denote the value of metric m at layer ℓ for sample x . The three metrics span several orders of magnitude in raw form, so we first apply a sign-preserving compression to tame their heavy-tailed regime while leaving values near zero unchanged:

$$\tilde{m}_\ell(x) = \text{sign}(m_\ell(x)) \log(1 + |m_\ell(x)|). \quad (1)$$

For each (m, ℓ) , we estimate the clean reference *center* and *scale* from non-contaminated validation

samples $\mathcal{D}^{\text{clean}}$:

$$\mu_{m,\ell}^{\text{clean}} = \text{median}(\tilde{m}_\ell(x) : x \in \mathcal{D}^{\text{clean}}), \quad (2)$$

$$\sigma_{m,\ell}^{\text{clean}} = 1.4826 \cdot \text{MAD}(\tilde{m}_\ell(x) : x \in \mathcal{D}^{\text{clean}}), \quad (3)$$

where the factor 1.4826 is the standard scaling to make median absolute deviation (MAD) a consistent estimator of the standard deviation under Gaussian noise (see Appendix G). The standardized geometric deviation of sample x at (m, ℓ) is then:

$$z_{m,\ell}(x) = \frac{\tilde{m}_\ell(x) - \mu_{m,\ell}^{\text{clean}}}{\sigma_{m,\ell}^{\text{clean}} + \epsilon}, \quad (4)$$

with ϵ a small numerical constant. This formulation preserves the relative magnitude of geometric deviations while preventing a small number of extreme contaminated samples from inflating the clean-reference scale and washing out the signal for the rest of the population.

Step 2: Metric-specific Anomaly Alignment.

Our analyses show that contamination affects each metric through a different geometric mechanism. Contaminated samples tend to exhibit elevated perturbation sensitivity in RSM, directional concentration shifts in DC, and lower local representation variability in RSI. Importantly, the raw RSI metric introduced in Section 3 measures local representation variability, such that larger values indicate greater variability. As shown in Section 3.3, contaminated samples generally have lower RSI values than clean samples, reflecting greater local representational rigidity. Therefore, before aggregating the metrics, we align their directions so that a larger aligned score consistently indicates stronger evidence of contamination:

$$\hat{z}_{m,\ell}(x) = \begin{cases} z_{m,\ell}(x), & m = \text{RSM}, \\ z_{m,\ell}(x), & m = \text{DC}, \\ -z_{m,\ell}(x), & m = \text{RSI}. \end{cases} \quad (5)$$

For RSM and DC, the original sign is preserved because their positive standardized deviations are aligned with the corresponding contamination-associated patterns. For RSI, we negate the standardized score because lower raw RSI values indicate the increased local rigidity associated with contamination. This sign alignment is applied only when constructing the aggregated detection score and does not change the definition or interpretation of the underlying RSI metric.

Step 3: Layer-wise Aggregation. We aggregate the aligned deviations $\hat{z}_{m,\ell}(x)$ across the metric set $\mathcal{M}$ and layer set $\mathcal{L}$ to obtain a single per-sample score, where larger values indicate stronger overall deviation from the clean geometric profile.

$$S_{\text{LaRA}}(x) = \frac{1}{|\mathcal{M}| |\mathcal{L}|} \sum_{m \in \mathcal{M}} \sum_{\ell \in \mathcal{L}} \hat{z}_{m,\ell}(x). \quad (6)$$

Because all (m, ℓ) contributions are standardized onto the same robust z-scale before aggregation, abnormalities arising from different layers and metrics can be consistently compared and combined within $S_{\text{LaRA}}(x)$. Accordingly, their combination is intended not merely to maximize detection performance, but to offer a richer analysis of representation geometry under RL-induced memorization.

5 Experiments

Setups. We evaluate contamination detection performance using standard metrics for MIA (Zhang et al., 2024; Tao et al., 2025; Kwak and Kim, 2026). *ROC-AUC* (AUC) measures the model’s ability to distinguish between member and non-member samples across all possible decision thresholds. *TPR@FPR=5%* reports the true positive rate (*i.e.*, correctly identified members) when the false positive rate (*i.e.*, non-members incorrectly flagged as members) is fixed at 5%. We consider six representative baselines, Recall, CDD, Min-K%, Min-K%++, PPL, and Self-Critique (SC). Refer to Appendix H for further details.

5.1 Main Results

Table 1 shows that our proposed representation-based membership score, S_{LaRA} , consistently achieves strong and stable detection performance across different RL model families and training checkpoints. In the initial checkpoints, S_{LaRA} attains the best overall performance on LIMR with an AUC of 0.80 and *TPR@FPR=5%* of 0.46, substantially outperforming standard baselines such as Recall, Min-K%, and SC. We also explore combining S_{LaRA} with SC, the sota output-level detection method, to see the complementarity of the two detection regimes. Combining S_{LaRA} with SC ($\text{SC} + S_{\text{LaRA}}$) achieves the strongest overall performance on Euris, reaching an AUC of 0.73 and *TPR@FPR=5%* of 0.31 at initialization, while also maintaining competitive performance throughout

Method	Eurus		LIMR		OLMO	
	AUC	TPR	AUC	TPR	AUC	TPR
Recall	0.59	0.07	0.54	0.07	0.42	0.03
CDD	0.38	0.00	0.64	0.00	<u>0.50</u>	0.00
Min-K%	0.32	0.07	0.26	0.03	<u>0.50</u>	0.03
Min-K%++	0.29	0.03	0.60	0.07	0.46	<u>0.07</u>
PPL	0.68	0.23	<u>0.73</u>	<u>0.13</u>	0.49	0.07
SC	<u>0.70</u>	<u>0.27</u>	0.44	0.00	0.42	0.03
S_{LaRA}	0.63	0.19	0.80	0.46	0.54	0.04
SC + S_{LaRA}	0.73	0.31	0.47	0.04	<u>0.50</u>	0.08

(a) Initial checkpoint results across models.

Method	Eurus						LIMR					
	Init		E1		E2		Init		E1		E2	
	AUC	TPR	AUC	TPR	AUC	TPR	AUC	TPR	AUC	TPR	AUC	TPR
Recall	0.59	0.07	0.56	0.13	0.52	0.00	0.54	0.07	0.47	0.03	0.47	0.00
CDD	0.38	0.00	0.38	0.00	0.38	0.00	0.64	0.00	0.61	0.00	0.59	0.00
Min-K%	0.32	0.07	0.31	0.07	0.33	0.07	0.26	0.03	0.26	0.00	0.26	0.00
Min-K%++	0.29	0.03	0.26	0.03	0.20	0.03	0.60	0.07	0.58	0.07	0.63	0.00
PPL	0.68	0.23	0.69	0.23	0.67	0.17	<u>0.73</u>	<u>0.13</u>	0.74	0.13	<u>0.74</u>	<u>0.17</u>
SC	<u>0.70</u>	<u>0.27</u>	0.58	0.07	<u>0.78</u>	<u>0.30</u>	0.44	0.00	0.40	0.00	0.46	0.00
S_{LaRA}	0.63	0.19	<u>0.68</u>	0.19	0.70	0.15	0.80	0.46	<u>0.72</u>	0.23	0.81	0.20
SC + S_{LaRA}	0.73	0.31	0.65	0.35	0.79	0.38	0.47	0.04	0.46	0.04	0.54	0.04

(b) Results across Eurus and LIMR training checkpoints.

Table 1: **Membership inference performance across different RL checkpoints and model families.** (a) compares initial checkpoints across Eurus, LIMR, and OLMO, while (b) analyzes performance evolution during RL training for Eurus and LIMR. TPR denotes TPR@FPR=5%. Details are in C.

RSM	DC	RSI	Init		E1		E2	
			AUC	TPR	AUC	TPR	AUC	TPR
$\times$	$\times$	$\checkmark$	0.60	0.19	0.51	0.15	0.56	0.16
$\times$	$\checkmark$	$\times$	0.71	0.42	0.66	0.08	0.76	0.24
$\checkmark$	$\times$	$\times$	0.75	0.12	0.69	0.08	0.71	0.04
$\times$	$\checkmark$	$\checkmark$	0.76	0.46	0.69	0.15	0.81	0.32
$\checkmark$	$\times$	$\checkmark$	0.74	0.27	0.68	0.08	0.71	0.12
$\checkmark$	$\checkmark$	$\times$	0.78	0.31	0.71	0.19	0.79	0.12
$\checkmark$	$\checkmark$	$\checkmark$	0.80	0.46	0.72	0.23	0.81	0.20

Table 2: **Metric ablations across RL training epochs.**

RL training. Across Eurus checkpoints, the combined score steadily improves from (0.73, 0.31) to (0.79, 0.38) in terms of (AUC, TPR@FPR=5%), suggesting that representation-level contamination signals become increasingly separable during RL optimization. LIMR exhibits a similar trend for S_{LaRA} , where performance remains consistently high across checkpoints, peaking at (0.81, 0.20) at epoch2. Although PPL occasionally shows relatively high AUC values, their TPR@FPR=5% remains substantially lower and less stable than the proposed methods. PPL often relies on superficial token likelihood differences that can fluctuate across model families and RL stages, whereas S_{LaRA} and SC + S_{LaRA} capture deeper geometric inconsistencies in hidden representations. Therefore, our approach leads to more reliable detection under strict low-FPR operating regimes critical for realistic settings.

5.2 Additional Analyses

Metric Ablations. Table 2 shows that combining all three components (RSM, DC, and RSI) consistently achieves the best overall performance across RL checkpoints, improving both AUC and robustness to training-stage shifts. While DC provides the strongest standalone discrimination signal, its per-

formance varies more across epochs, particularly in TPR@FPR=5%, indicating reduced robustness when used alone. In contrast, **RSM** and **RSI** individually produce weaker detection performance but contribute to improving generalization under RL post-training. Removing an individual component from the full metric consistently degrades performance, showing that the final score benefits from jointly modeling perturbation sensitivity, directional representation geometry, and local invariance. Overall, the results suggest that robust contamination detection requires integrating multiple representation-level signals rather than relying on a single geometric statistic.

Beta Sweep over S_{LaRA} and SC Mix. We sweep the mixture weight β in $\text{mix} = \beta(\text{SC}) + (1 - \beta)(S_{LaRA})$ on the member-detection benchmark ($\beta \in \{0, 0.25, 0.5, 0.65, 0.75, 1\}$). The optimal balance between SC and S_{LaRA} is strongly model-dependent. For Eurus, performance improves as more weight is assigned to SC, with AUC peaking at the default $\beta = 0.65$ and TPR@FPR=5% at $\beta = 0.75$. In contrast, LIMR performs best with only the S_{LaRA} , with performance degrading as β increases. For OLMO, AUC is highest at $\beta = 0$, while TPR@FPR=5% peaks at $\beta = 0.65$. Overall, these results suggest that no single mixture weight is universally optimal; however, despite not being tuned per model, the shared default $\beta = 0.65$ still achieves competitive overall performance, consistent with the main results.

Correlation with Output-level Metrics. Figure 5(a) shows the correlation between the proposed S_{LaRA} and several output-level metrics. Higher S_{LaRA} values are negatively correlated with SC ($\rho = -0.16$) and PPL ($\rho = -0.29$), while exhibiting a weak positive correlation with Min-

data_source	prompt	member	RSM	DC	RSI
numina_amc_aime	Let the set $S = \{P_1, P_2, \dots, P_{12}\}$ consist of the twelve vertices of a regular 12-gon. A subset Q of S is called "communal" if there is a circle such that all points of Q are inside the circle, and all points of S not in Q are outside of the circle. How many communal subsets are there? (Note that the empty set is a communal subset.) Present the answer in LaTeX format: <code>\boxed{Your answer}</code>	1	0.151	0.423	0.310
aime26	Let triangle ABC have side lengths AB = 13, BC = 14, and CA = 15. Triangle A'B'C' is obtained by rotating triangle ABC about its circumcenter so that AC is perpendicular BC, with A' and B not on the same side of line B'C'. Find the integer closest to the area of hexagon AA'CC'BB'. Present the answer in LaTeX format: <code>\boxed{Your answer}</code>	0	1.191	9.765	0.370

Table 3: **Error analysis.** (1) member sample often undetected, and (2) non-member sample often wrongly detected.




Figure 4: **Beta sweep over score mix.** Default $\beta = 0.65$.

K%++ ($\rho = 0.17$). Additionally, samples with low S_{LaRA} display substantially larger variability across all metrics, whereas high- S_{LaRA} samples are concentrated within narrower output regimes. Correlations suggest that stronger contamination-related geometric deviations are associated with increasingly confident, less reflective, and more behaviorally concentrated generations. In particular, high- S_{LaRA} samples occupy narrower output regimes characterized by reduced variability across output-level metrics.

Analysis on Number of Perturbations. The perturbation-count sweep on Euruss (Figure 5(b)) shows that S_{LaRA} remains relatively stable across different numbers of blanks. Although AUC improves from the default $k = 1$ (0.62) to a peak at $k = 3$ (0.77), TPR@FPR=5% only substantially increases at $k = 4$ (0.32), where AUC slightly declines (0.69). This indicates that increasing the number of blanks can sometimes slightly strengthen detection, but the gains depend on the evaluation metric and operating regime. Overall, the method is robust to the choice of k , and the default setting $k = 1$ already provides competitive performance without requiring additional perturbation variants.

Analysis on Perturbation Types. The perturbation analysis in Figure 6 shows that LaRA remains relatively robust across different perturbation types, with all variants achieving comparable AUC values between 0.56 and 0.69. Although Distractor Insert. achieves the highest AUC (0.69) and Num

Replace. obtains the best TPR@FPR=5% (0.17), the default setting based on Info Rem., the most naive method in making perturbations, still maintains reasonable detection performance without requiring semantically invasive modifications. In contrast to perturbations that directly alter numerical or semantic reasoning components, the default perturbation preserves the original problem structure more conservatively while still producing distinguishable contamination signals. Overall, results suggest that LaRA does not rely on a single perturbation type and exhibits stable detection behavior across diverse perturbation strategies.

Error analysis. Failure cases presented in Table 3 suggest that S_{LaRA} misses member examples whose representation geometry remains close to the non-member manifold. The false-negative member sample exhibits uniformly low component scores, with RSM, DC, and RSI values of 0.151, 0.423, and 0.310, respectively, producing a low aggregate score of 0.295, well below the detection threshold. This indicates that the sample does not induce strong perturbation sensitivity, directional concentration, or local invariance, causing its layer-wise representation trajectory to resemble that of a clean example. Conversely, the non-member example that is falsely detected exhibits an abnormally high S_{LaRA} score, primarily driven by an unusually large DC value. Although the sample is not a member, its representation geometry deviates substantially from the typical clean distribution, leading the detector to classify it as contaminated.

6 Conclusion

We introduce LaRA, a layer-wise representation analysis framework for contamination detection in RL-trained models. Unlike prior methods, LaRA detects contamination through perturbation-induced representation geometry across layers, measured with three proposed metrics. Based on layer-wise analysis, we propose a detection framework that aggregates geometric deviations across metrics and layers. Experiments across various




Figure 5: (a) Correlation plot of output-level signals and S_{LaRA} . (b) Analysis on the # of perturbations.




Figure 6: **Perturbation type analysis.** Results show that S_{LaRA} is robust to perturbation types.

RL-trained models demonstrate that representation-level signals complement output-level methods and improve detection performance. Overall, our results show that contamination in RL-trained LLMs is strongly reflected in internal representation geometry, highlighting the effectiveness of representation-level auditing.

Limitations

Despite its new findings and effectiveness, our work has several limitations. First, LaRA requires access to intermediate hidden representations and is therefore currently limited to white-box settings. This requirement reflects the primary objective of our study: to investigate whether representation geometry provides a more reliable contamination signal than output-level statistics under RL post-training, rather than to develop a universally deployable detector. Nevertheless, this constraint limits the applicability of LaRA to closed-source models accessible only through black-box APIs. An important direction for future work is to translate our representation-level insights into practical black-box detectors, for example, by learning output-level proxies that approximate hidden-state geometry.

Second, LaRA relies on representation extraction under multiple semantic perturbations and layer-wise hidden-state analyses, introducing additional computational overhead compared with lightweight output-level approaches. In particular, the framework requires generating perturbed vari-

ants, extracting intermediate representations across transformer layers, and aggregating multiple geometric statistics, making inference more expensive than methods operating solely on final outputs or token probabilities.

Third, although the proposed framework consistently improves contamination detection over existing baselines, detection remains imperfect for challenging examples whose representation geometry closely overlaps with the clean distribution. This suggests that some memorized samples may not produce sufficiently distinctive internal signatures for reliable separation. Moreover, while our analyses reveal consistent trends across models and training checkpoints, the precise relationship between RL post-training dynamics and representation-level memorization remains only partially understood. Future work may explore more computationally efficient perturbation strategies, stronger representation aggregation methods, and deeper investigations into the causal mechanisms underlying memorization and contamination in reasoning models.

Broader Impact and Ethical Implications

This work contributes to improving transparency and reliability in the evaluation of RL-trained LLMs by introducing a representation-level framework for contamination detection. Stronger contamination auditing can help researchers better assess benchmark integrity, reduce hidden memorization effects, and improve the trustworthiness of reported reasoning capabilities. At the same time, contamination detection methods may potentially be adapted for membership inference or unauthorized dataset auditing, which could raise privacy or data governance concerns when applied to sensitive or proprietary data. Our experiments are conducted exclusively on publicly available open-source models and public mathematical datasets, and we do not use private or personally identifiable information. Overall, we believe representation-level auditing should be used responsibly as one component of broader efforts toward reliable, transparent, and

reproducible evaluation of large language models.

Acknowledgments

Minju Gwak, Minseo Kwak, Dongseok Lee, and Jaehyung Kim are affiliated with the Department of Artificial Intelligence at Yonsei University. This research was supported in part by Institute for Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2020-II201361, Artificial Intelligence Graduate School Program (Yonsei University); No. RS-2025-02215344, Development of AI Technology with Robust and Flexible Resilience Against Risk Factors).

References

- Mislav Balunović, Jasper Dekoninck, Ivo Petrov, Nikola Jovanović, and Martin Vechev. 2025. Matharena: Evaluating llms on uncontaminated math competitions.
- Jinhe Bi, Danqi Yan, Yifan Wang, Wenke Huang, Haokun Chen, Guancheng Wan, Mang Ye, Xun Xiao, Hinrich Schuetze, Volker Tresp, and Yunpu Ma. 2026. Reasoning self-evaluation via trajectory dynamics modeling.
- Hyeong Kyu Choi, Maxim Khanov, Hongxin Wei, and Yixuan Li. 2025. How contaminated is your benchmark? quantifying dataset leakage in large language models with kernel divergence. In *International Conference on Machine Learning*.
- Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Yuchen Zhang, Jiacheng Chen, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, et al. 2025. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*.
- Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gestein, and Arman Cohan. 2024. Investigating data contamination in modern benchmarks for large language models. 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 8706–8719.
- Yihong Dong, Xue Jiang, Huanyu Liu, Zhi Jin, Bin Gu, Mengfei Yang, and Ge Li. 2024. Generalization or memorization: Data contamination and trustworthy evaluation for large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12039–12050.
- Shahriar Golchin and Mihai Surdeanu. 2024. Time travel in llms: Tracing data contamination in large language models. In *International Conference on Learning Representations*, volume 2024, pages 43008–43029.
- Shahriar Golchin and Mihai Surdeanu. 2025. Data contamination quiz: A tool to detect and estimate contamination in large language models. *Transactions of the Association for Computational Linguistics*, 13:809–830.
- Hila Gonen, Srinu Iyer, Terra Blevins, Noah A Smith, and Luke Zettlemoyer. 2023. Demystifying prompts in language models via perplexity estimation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10136–10148.
- Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, et al. 2025. Openthoughts: Data recipes for reasoning models. In *International Conference on Learning Representations*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Minju Gwak, Guijin Son, and Jaehyung Kim. 2025. Revisiting the uniform information density hypothesis in llm reasoning traces. *arXiv preprint arXiv:2510.06953*.
- Frank R. Hampel, Elvezio M. Ronchetti, Peter J. Rousseeuw, and Werner A. Stahel. 1986. *Robust Statistics: The Approach Based on Influence Functions*. Wiley.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. Training large language models to reason in a continuous latent space. In *Conference on Language Modeling*.
- Andreas Hochlehnert, Hardik Bhatnagar, Vishal Udaydarao, Samuel Albanie, Ameya Prabhu, and Matthias Bethge. 2025. A sober look at progress in language model reasoning: Pitfalls and paths to reproducibility. In *Conference on Language Modeling*.
- Peter J. Huber and Elvezio M. Ronchetti. 2009. *Robust Statistics*, 2 edition. Wiley.
- Zhewei Kang, Xuandong Zhao, and Dawn Song. 2025. Scalable best-of-n selection for large language models via self-certainty. In *Advances in Neural Information Processing Systems*.
- Minseo Kwak and Jaehyung Kim. 2026. Gap-k%: Measuring top-1 prediction gap for detecting pretraining data. *arXiv preprint arXiv:2601.19936*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.

- Bruce W Lee, Inkit Padhi, Karthikeyan Natesan Ramamurthy, Erik Miehl, Pierre Dognin, Manish Nagireddy, and Amit Dhurandhar. 2024. Programming refusal with conditional activation steering. In *International Conference on Learning Representations*.
- Dongseok Lee, Jimyung Hong, Dongyoung Kim, and Jaehyung Kim. 2025. Training-free llm verification via recycling few-shot examples. *arXiv preprint arXiv:2506.17251*.
- Jixuan Leng, Chengsong Huang, Banghua Zhu, and Jiaxin Huang. 2025. Taming overconfidence in llms: Reward calibration in rlhf. In *International Conference on Learning Representations*, volume 2025.
- Kenneth Li, Oam Patel, Fernanda Viégas, Hanspeter Pfister, and Martin Wattenberg. Inference-time intervention: Eliciting truthful answers from a language model. In *Advances in Neural Information Processing Systems*.
- Melody Zixuan Li, Kumar Krishna Agrawal, Arna Ghosh, Komal Kumar Teru, Adam Santoro, Guillaume Lajoie, and Blake A Richards. 2025a. Tracing the representation geometry of language models from pretraining to post-training. *arXiv preprint arXiv:2509.23024*.
- Xuefeng Li, Haoyang Zou, and Pengfei Liu. 2025b. Limr: Less is more for rl scaling. *arXiv preprint arXiv:2502.11886*.
- Disha Makhija, Manoj Ghuman Arivazhagan, Vinayshekhar Bannihatti Kumar, and Rashmi Gangadharaiyah. 2025. *Neural breadcrumbs: Membership inference attacks on llms through hidden state and attention pattern analysis*. *Preprint*, arXiv:2509.05449.
- Ricardo A. Maronna, R. Douglas Martin, Victor J. Yohai, and Matias Salibián-Barrera. 2019. *Robust Statistics: Theory and Methods (with R)*, 2 edition. Wiley.
- Team Olmo, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, et al. 2025. Olmo 3. *arXiv preprint arXiv:2512.13961*.
- OpenAI. 2024. Gpt-4o mini: advancing cost-efficient intelligence.
- OpenRouter. 2024. Openrouter: Unified api for AI models.
- Ravi Ranjan, Utkarsh Grover, Xiaomin Lin, and Agoritsa Polyzou. 2026. *G-drift mia: Membership inference via gradient-induced feature drift in llms*. *Preprint*, arXiv:2604.00419.
- Youngji Roh, Hyunjin Cho, and Jaehyung Kim. 2026. Embracing anisotropy: Turning massive activations into interpretable control knobs for large language models. *arXiv preprint arXiv:2603.00029*.
- Peter J. Rousseeuw and Christophe Croux. 1993. Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424):1273–1283.
- Peter J. Rousseeuw and Annick M. Leroy. 1987. *Robust Regression and Outlier Detection*. Wiley.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2024. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv:2409.19256*.
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. 2023. Detecting pretraining data from large language models. In *International Conference on Learning Representations*.
- Yongding Tao, Tian Wang, Yihong Dong, Huanyu Liu, Kechi Zhang, Xiaolong Hu, and Ge Li. 2025. Detecting data contamination from reinforcement learning post-training for large language models. In *International Conference on Learning Representations*.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J Vazquez, Ulisse Mini, and Monte MacDiarmid. 2023. Steering language models with activation engineering. *arXiv preprint arXiv:2308.10248*.
- Han Wang, Haoyu Li, Brian Ko, and Huan Zhang. 2025. On the fragility of benchmark contamination detection in reasoning models. In *International Conference on Learning Representations*.
- Yiming Wang, Pei Zhang, Baosong Yang, Derek F Wong, and Rui Wang. 2024. Latent space chain-of-embedding enables output-free llm self-evaluation. In *International Conference on Learning Representations*.
- Hengyu Wu and Yang Cao. 2025. Membership inference attacks on large-scale models: A survey. In *Advances in Neural Information Processing Systems*.
- Mingqi Wu, Zhihao Zhang, Qiaole Dong, Zhiheng Xi, Jun Zhao, Senjie Jin, Xiaoran Fan, Yuhao Zhou, Huijie Lv, Ming Zhang, et al. 2026. Reasoning or memorization? unreliable results of reinforcement learning due to data contamination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 33944–33952.
- Daniel Wurgajt, Can Rager, Matthew Kowal, Vasudev Shyam, Sheridan Feucht, Usha Bhalla, Tal Haklay, Eric Bigelow, Raphael Sarfati, Thomas McGrath, Owen Lewis, Jack Merullo, Noah Goodman, Thomas Fel, Atticus Geiger, and Ekdeep Singh Lubana. 2026. *Manifold steering reveals the shared geometry of neural network representation and behavior*. *Preprint*, arXiv:2605.05115.

- Jiancong Xiao, Bojian Hou, Zhanliang Wang, Ruochen Jin, Qi Long, Weijie J Su, and Li Shen. 2025. Restoring calibration for aligned large language models: A calibration-aware fine-tuning approach. In *International Conference on Machine Learning*.
- Roy Xie, Junlin Wang, Ruomin Huang, Minxing Zhang, Rong Ge, Jian Pei, Neil Zhenqiang Gong, and Bhuwan Dhingra. 2024. Recall: Membership inference via relative conditional log-likelihoods. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8671–8689.
- Jingyang Zhang, Jingwei Sun, Eric Yeats, Yang Ouyang, Martin Kuo, Jianyi Zhang, Hao Frank Yang, and Hai Li. 2024. Min-k%++: Improved baseline for detecting pre-training data from large language models. In *International Conference on Learning Representations*.
- Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. 2025. Learning to reason without external rewards. In *Advances in Neural Information Processing Systems*.

A Details of Curated Datasets

We provide examples of the curated evaluation and training dataset in Table 4 and Table 5.

B Implementation Details

B.1 Training Details

We fine-tune the base model using Group Relative Policy Optimization (GRPO) within the VeRL (Sheng et al., 2024) training framework. Training is conducted for 2 epochs with a learning rate of 1×10^{-6} , train batch size 128, validation batch size 512, maximum prompt length 1024, and maximum response length 4096. We enable gradient checkpointing and dynamic batch sizing during optimization, with a per-GPU token budget of 16384 tokens. For rollout generation, we use vLLM (Kwon et al., 2023) with 4 sampled responses per prompt at temperature 1.0, while validation uses temperature 0.6. We do not apply explicit KL regularization during training. All training experiments are performed on $8 \times$ NVIDIA A6000 GPUs.

B.2 Inference Details

We conduct two types of evaluation: reasoning evaluation and contamination evaluation.

Reasoning Evaluation. We conduct reasoning evaluation in Section B.2 to test training robustness. Here, we generate 5 sampled responses per example and report pass@5, where a prediction is considered correct if at least one sampled response matches the ground-truth answer after extracting the final boxed or numeric answer. In addition, we compute the mean length-normalized token log-probability of the gold answer under the model:

$$\frac{1}{T} \sum_{t=1}^T \log p(a_t \mid \text{prompt}, a_{<t}),$$

where T denotes the answer length. Log-probabilities are computed using either a standard Transformers forward pass or vLLM-based prompt log-probability scoring.

Contamination Evaluation. Contamination evaluation follows a two-stage pipeline. First, we generate model responses together with token-level statistics such as log-probabilities and entropies. Second, we compute contamination detection scores using both output-based and representation-based methods.

The evaluated baselines include Min-K%, PPL, CDD, Recall, and Self-Critique. For representation-based methods, LaRA constructs semantically related and perturbed variants of each question using gpt-4o-mini-generated (via OpenRouter API) (OpenAI, 2024; OpenRouter, 2024) paraphrases and incomplete-question variants, from which representation-level signals such as RSI/RSM and directional collapse are derived. We evaluate member versus non-member separability using AUROC, TPR at fixed FPR (0.05). All inference and evaluation experiments are performed on $4 \times$ NVIDIA RTX 3090 GPUs.

C Main Table with Confidence Intervals

To quantify the statistical uncertainty arising from the limited evaluation size, we report bootstrap confidence intervals for both AUC and TPR@5% FPR. Due to computational constraints, we conduct this analysis using EURUS as a representative testbed. Each evaluation set contains 60 examples, consisting of 30 member and 30 non-member examples. We perform 10,000 stratified bootstrap resamples with replacement, preserving the number of member and non-member examples in each resample, using random seed 42. For each resample, we recompute the detection threshold and evaluate both metrics. We report 95% percentile bootstrap confidence intervals, while the point estimate corresponds to the metric computed on the full evaluation set.

Overall, the bootstrap analysis confirms that statistical uncertainty is non-negligible at the current evaluation scale, particularly for TPR@5% FPR. Nevertheless, the principal AUC-based observations remain unchanged. In particular, at the Epoch 2 checkpoint, Self-Critique combined with S_LaRA achieves the highest AUC, while S_LaRA remains competitive across checkpoints. Given the instability of low-FPR estimates at this sample size, however, small differences in TPR@5% FPR should not be interpreted as statistically meaningful.

D Evaluation on Same-Source Member and Non-Member Sets

A potential confounding factor in contamination detection is a distributional difference between the member and non-member sets. In our primary evaluation, we mitigate this concern by restricting member examples to Olympiad-level mathematics

Split	Source	Composition		Purpose
		Mem.	Non-Mem.	
Evaluation	Open-source + AIME'26	30	30	Evaluation
Training	30 members + 970 RL-MIA	30	970	RL Analysis

Table 4: Overview of the curated datasets.

data_source	prompt	answer	member	metadata
Contamination evaluation samples				
aime26	[{role: system, content: When tackling complex reasoning tasks, you have access to actions such as ASSESS, ADVANCE, VERIFY, SIMPLIFY, SYNTHESIZE, PIVOT, and OUTPUT.}, {role: user, content: Patrick started walking at a constant rate from school to the park. Tanya ran 2 miles per hour faster than Patrick, Jose bicycled 7 miles per hour faster than Tanya, and all three arrived at the same time. Find $m+n$.}]	277	0	–
numina_amc_aime	[{role: system, content: When tackling complex reasoning tasks, you have access to actions such as ASSESS, ADVANCE, VERIFY, SIMPLIFY, SYNTHESIZE, PIVOT, and OUTPUT.}, {role: user, content: The highest price is \$8.50 and the lowest price is \$5.50. Calculate the percent by which the highest price is more than the lowest price.}]	70%	1	–
RL training sample				
olympiads	[{role: system, content: Your task is to follow a systematic, thorough reasoning process before providing the final solution. Structure your response into two sections: Thought and Solution. In the Thought section, present your reasoning using <code><think>{thoughts}</think></code> . In the Solution section, provide the final logical answer, optionally in <code>\boxed{}</code> format.}, {role: user, content: Determine the largest even positive integer which cannot be expressed as the sum of two composite odd positive integers.}]	38	–	{style: "rule"}

Table 5: Samples from the curated contamination evaluation and RL training datasets. Each instance contains the data source, structured conversational prompt, ground-truth answer, membership label when applicable, and metadata annotations.

problems comparable to AIME 2026, rather than sampling indiscriminately from the model’s entire training corpus. Nevertheless, the two sets may still differ in their collection sources.

To more directly control for this factor, we conduct an additional evaluation following the same-source protocol of Tao et al. (2025). Specifically, we use the KK and SAT benchmarks from Tao et al. (2025), for which both member and non-member examples are drawn from the same underlying dataset and therefore share the same collection procedure, problem domain, and overall distribution. Within each benchmark, the only intended distinction is whether an example was included in the model’s RL training set. We evaluate checkpoints after the first and second RL training epochs for both EURUS and LIMR.

Tables 8 and 9 report ROC AUC and TPR at 5% FPR. For EURUS, LARA and its combination with Self-Critique consistently achieve above-random ROC AUC across both benchmarks and training checkpoints. On KK, the combined method obtains ROC AUCs of 0.739 and 0.699 after the first

and second epochs, respectively. On SAT, LARA achieves ROC AUCs of 0.622 and 0.689, while the combined method obtains 0.627 and 0.670. These results also indicate that the relative effectiveness of the methods can change over the course of RL training.

The results are particularly clear for LIMR. Across all four benchmark-checkpoint combinations, LARA achieves ROC AUCs between 0.781 and 0.821, whereas Self-Critique remains close to random performance. Combining the two signals does not consistently improve upon LARA alone for this model, suggesting that the output-level Self-Critique signal contributes limited complementary information in this setting.

Overall, LARA continues to distinguish member from non-member examples when dataset source, collection procedure, and problem domain are controlled. These findings reduce the likelihood that its performance is driven primarily by dataset-specific artifacts or superficial domain shifts, and instead support the conclusion that its representation-level signals capture whether examples were encoun-

Model	Overall Pass@5	Pass@5 (%)		LogP(answer prompt)	
		Member	Non-member	Member	Non-member
Qwen2.5-Math-7B (Base)	15.0 (9/60)	26.7 (8/30)	3.3 (1/30)	-2.08	-2.44
Eurus-2-7B-PRIME (Initial)	18.3 (11/60)	33.3 (10/30)	3.3 (1/30)	-8.54	-7.88
Eurus-2-7B-PRIME (Epoch 1)	18.3 (11/60)	36.7 (11/30)	0.0 (0/30)	-9.06	-8.50
Eurus-2-7B-PRIME (Epoch 2)	18.3 (11/60)	36.7 (11/30)	0.0 (0/30)	-8.85	-9.12
Qwen2.5-Math-7B (Base)	10.0 (6/60)	20.0 (6/30)	0.0 (0/30)	-2.07	-2.44
LIMR (Initial)	8.0 (5/60)	17.0 (5/30)	0.0 (0/30)	-2.48	-2.92
LIMR (Epoch 1)	13.3 (8/60)	23.3 (7/30)	3.0 (1/30)	-2.44	-2.90
LIMR (Epoch 2)	10.0 (6/60)	20.0 (6/30)	0.0 (0/30)	-2.43	-2.87

Table 6: Comparison between the RL-trained open-source model and its base model on member and non-member samples.

Checkpoint	Method	AUC (95% CI)	TPR@5% FPR (95% CI)
Initial	S_LaRA	0.628 [0.461, 0.784]	0.192 [0.033, 0.444]
Initial	Self-Critique + S_LaRA	0.682 [0.521, 0.829]	0.154 [0.000, 0.654]
Epoch 1	Self-Critique	0.586 [0.436, 0.730]	0.233 [0.054, 0.417]
Epoch 1	S_LaRA	0.679 [0.519, 0.825]	0.192 [0.040, 0.409]
Epoch 1	Self-Critique + S_LaRA	0.671 [0.512, 0.813]	0.269 [0.077, 0.519]
Epoch 2	Self-Critique	0.820 [0.706, 0.915]	0.367 [0.179, 0.645]
Epoch 2	S_LaRA	0.704 [0.546, 0.849]	0.154 [0.000, 0.591]
Epoch 2	Self-Critique + S_LaRA	0.858 [0.732, 0.961]	0.231 [0.000, 0.920]

Table 7: Bootstrap point estimates and 95% percentile confidence intervals on EURUS, computed using 10,000 stratified bootstrap resamples. Confidence intervals are generally wider for TPR@5% FPR because its threshold is determined by only a small number of non-member examples.

tered during RL training.

E Additional Results on Figure Geometry

Figure 7 further shows that the layer-wise representation geometry patterns vary across RL-trained models while still exhibiting contamination-associated deviations. In Eurus-2-7B-PRIME, contaminated samples exhibit persistently elevated RSM values throughout depth with substantially larger magnitudes than clean samples, indicating strong perturbation sensitivity under information removal. Regarding DC, while the specific directional profiles vary across models and checkpoints, contaminated samples consistently exhibit deviations from the stable directional organization observed in clean samples, indicating disrupted and lower-dimensional perturbation dynamics. Exhibit comparatively lower RSI values than clean samples, suggesting induced local representation flexibility under paraphrastic perturbations.

In contrast, Olmo-3.1-7B-RL-Zero-Math exhibits weaker RSM separation, with both contaminated and clean samples remaining near zero across

Run	Method	ROC AUC	TPR@5% FPR
KK epoch 1	Self-Critique	0.664	0.233
	LARA	0.641	0.192
	Self-Critique + LARA	0.739	0.385
KK epoch 2	Self-Critique	0.619	0.200
	LARA	0.724	0.269
	Self-Critique + LARA	0.699	0.269
SAT epoch 1	Self-Critique	0.579	0.167
	LARA	0.622	0.154
	Self-Critique + LARA	0.627	0.231
SAT epoch 2	Self-Critique	0.604	0.100
	LARA	0.689	0.154
	Self-Critique + LARA	0.670	0.231

Table 8: Same-source contamination detection results for EURUS. Member and non-member examples are drawn from the same underlying benchmark. The best result within each run is shown in bold.

most layers. However, contamination-related deviations remain observable in DC and RSI, where contaminated samples maintain consistently higher directional concentration and smoother local variability patterns than clean samples. Overall, these results suggest that while the precise geometric behaviors vary across models, contamination consistently alters perturbation sensitivity, directional organization, and local representation variability.

F Further Analyses

F.1 Analyses of Layer-window Trends on Different Models

Analysis on Eurus-2-7B-Prime. Figure 9 shows contaminated-vs-clean separation (Cohen’s d) across early (0–8), mid (9–17), and late (18–27) layers over three RL checkpoints. Three consistent trends emerge. First, **RSM remains nearly unchanged across both depth and training.** All checkpoints show stable positive separation ($d \approx 0.27$ – 0.35), indicating that perturbation sensitivity is largely preserved throughout RL

Run	Method	ROC AUC	TPR@5% FPR
KK epoch 1	Self-Critique	0.456	0.000
	LARA	0.798	0.435
	Self-Critique + LARA	0.567	0.044
KK epoch 2	Self-Critique	0.524	0.067
	LARA	0.798	0.480
	Self-Critique + LARA	0.624	0.080
SAT epoch 1	Self-Critique	0.528	0.033
	LARA	0.821	0.520
	Self-Critique + LARA	0.656	0.240
SAT epoch 2	Self-Critique	0.476	0.000
	LARA	0.781	0.160
	Self-Critique + LARA	0.598	0.080

Table 9: **Same-source contamination detection results for LIMR.** Member and non-member examples are drawn from the same underlying benchmark. The best result within each run is shown in bold.




Figure 7: **Representation Geometry of Additional Models.**

fine-tuning. Second, **DC becomes increasingly negative as RL training progresses**, especially in deeper layers. While the initial checkpoint shows weak or near-zero separation, Epoch 1 and Epoch 2 exhibit progressively larger negative effects, with the strongest separation appearing in late layers ($d \approx -0.65$ at Epoch 2). This suggests that RL training amplifies directional-collapse behavior on contaminated samples, particularly in deeper representations. Third, **RSI is strongest in early and mid layers but weakens in late layers after training**. Early-layer separation remains consistently negative across checkpoints, whereas the late-layer signal gradually diminishes and nearly disappears by Epoch 2. This indicates that local-invariance differences are primarily concentrated in shallower representations. Overall, the figure reveals a clear layer-conditioned structure: RSM is stable across training, DC is progressively amplified by RL in deeper layers, and RSI is concentrated in earlier layers. These

complementary trends motivate combining all three metrics in LaRA rather than relying on a single layer-wise signal.

Analysis on LIMR. Figure 10 shows contaminated-vs-clean separation (Cohen’s d) across early (0–8), mid (9–17), and late (18–27) layers over three RL checkpoints. Several distinct trends emerge. First, **RSM remains consistently positive across all layer groups**, indicating that contaminated samples exhibit larger representation shifts under perturbation throughout training. However, unlike the relatively stable behavior observed in Eurus-2-7B-Prime, the magnitude of separation decreases after RL training, particularly in deeper layers. Early and mid layers initially show strong separation ($d \approx 0.43$), but this gradually weakens by Epoch 2, especially in the late layers where the effect becomes nearly negligible. This suggests that RL training partially suppresses perturbation-sensitive geometry in LIMR. Second, **DC exhibits the strongest and most stable separation across all checkpoints**. All layer groups consistently show large negative effects ($d \approx -0.85$ to -1.25), with mid and late layers displaying the strongest separation. Moreover, RL training slightly amplifies this negative separation in Epoch 1 and Epoch 2, indicating that contaminated representations become increasingly directionally collapsed during RL optimization. Compared to the other metrics, DC provides the clearest and most persistent distinction between contaminated and clean samples across depth. Third, **RSI shows progressively stronger positive separation as RL training proceeds**, particularly in mid and late layers. While the initial checkpoint exhibits only moderate separation, Epoch 2 produces substantially larger effects ($d \approx 0.40$ – 0.45), especially beyond the middle layers. This indicates that RL training amplifies local-invariance differences between contaminated and clean samples, causing contaminated representations to become increasingly invariant under local perturbations. Overall, LIMR exhibits a complementary layer-wise structure distinct from Eurus-2-7B-Prime: RSM weakens during RL training, DC remains consistently dominant across all depths, and RSI becomes progressively amplified in deeper layers. These trends suggest that RL optimization in LIMR increasingly concentrates contaminated representations into geometrically collapsed yet locally invariant structures, motivating the joint

use of RSM, DC, and RSI in LaRA for robust contamination detection.

F.2 Aggregation over Different Layer-windows.

Figure 11 compares contamination detection performance when aggregating representation statistics over different layer windows (Early, Mid, and Late) as well as over all layers jointly. We report both AUC and TPR@FPR=5% across three RL-trained models. A clear and consistent trend emerges: **the relative contamination-separation behavior remains largely stable across different layer regions.**

Across all models, the performance ordering is highly consistent regardless of which layer window is used. LIMR consistently achieves the strongest separation across every window, obtaining nearly identical AUC values (≈ 0.8) for Early, Mid, and Late layers, along with the highest TPR values throughout. EURUS exhibits moderate but stable performance across all windows, while OLMO consistently shows weaker separation. Importantly, no single layer region dominates the results; instead, contamination-related geometric signals persist throughout the network depth. The results demonstrate that contamination-related representation geometry is remarkably stable across layer depth. Rather than arising from isolated layers, the separation signal persists throughout the network, motivating the use of layer-window aggregation as a robust and model-agnostic strategy for contamination detection.

F.3 Metric Ablations on Additional Model

We also conduct an additional ablation study on another model, Eurur-2-7B-Prime in Table 10. The results show that S_{LaRA} is relatively robust to ablations, as shown in LIMR results in the main ablation studies.

F.4 Analysis on the Evolution of Representation Geometry across RL Training Stages

Figure 8 illustrates how RL training progressively reshapes the layer-wise representation geometry of contaminated samples relative to clean samples in Eurur-2-7B-PRIME. Across all three metrics, the separation between clean and contaminated trajectories becomes increasingly pronounced as training advances from the initial checkpoint to Epochs 1 and 2. In particular, **RSM** exhibits the strongest




Figure 8: **Evolution of layer-wise representation geometry across RL training stages.** RL progressively amplifies geometric deviations in each metric (**RSM**, **DC**, and **RSI**) between clean and contaminated samples over epochs.

RSM	DC	RSI	Init		E1		E2	
			AUC	TPR	AUC	TPR	AUC	TPR
✗	✗	✓	0.59	0.19	0.49	0.00	0.54	0.08
✗	✓	✗	0.68	0.12	0.74	0.35	0.62	0.04
✓	✗	✗	0.59	0.27	0.53	0.23	0.52	0.27
✗	✓	✓	0.68	0.19	0.69	0.35	0.61	0.08
✓	✗	✓	0.62	0.23	0.54	0.23	0.56	0.23
✓	✓	✗	0.67	0.23	0.73	0.15	0.62	0.19
✓	✓	✓	0.68	0.19	0.70	0.15	0.63	0.19

Table 10: **Ablations on Eurur across RL epochs.**

divergence, where contaminated samples show dramatically amplified sensitivity beginning in early-middle layers and remaining consistently elevated throughout deeper layers, while clean samples stay nearly flat across all stages. **DC** further reveals that RL training induces increasingly distinct directional structure: contaminated samples transition from negative or near-zero values in middle layers to strongly positive values in deeper layers, whereas clean samples maintain smoother and more stable trajectories. Similarly, **RSI** demonstrates that contaminated samples undergo substantial geometric instability during RL training, especially in shallow layers where sharp spikes emerge after training, while clean samples retain comparatively moderate and stable behavior. Importantly, these trends remain consistent across layer depth and training stages, suggesting that RL fine-tuning systematically amplifies hidden-state geometric irregularities associated with memorized or contaminated data rather than producing isolated layer-specific effects.




Figure 9: **Effect-size separation across early, mid, and late layer windows of Eurur-2-7B-PRIME.** Contaminated samples are consistently elevated on RSM but progressively lower on DC and RSI in mid-to-late layers, with the widened gap indicating that RL training amplifies a layer-selective representational signature of contamination.




Figure 10: **Effect-size separation across early, mid, and late layer windows of LIMR.** Contaminated samples exhibit consistently lower DC across all layer windows, while RSI becomes increasingly positive in mid-to-late layers as RL training progresses. In contrast to Eurur-2-7B-PRIME, the separation dynamics in LIMR are dominated by strong directional concentration differences together with progressively amplified RSI gaps, indicating that RL training induces a distinct layer-dependent representation geometry signature of contamination in LIMR.

G Justifications of the Scaling Factor Metric in Contamination Detection Protocol

The factor $1.4826 = 1/\Phi^{-1}(0.75)$ is the standard Fisher-consistency constant used in robust statistics to make the median absolute deviation (MAD) a consistent estimator of the standard deviation under a Gaussian reference (Hampel et al., 1986; Huber and Ronchetti, 2009; Rousseeuw and Croux, 1993). Specifically, if $X \sim \mathcal{N}(\mu, \sigma^2)$, then $\text{median}|X - \mu| = \sigma \Phi^{-1}(0.75)$, so multiplying the empirical MAD by $1/\Phi^{-1}(0.75) \approx 1.4826$ recovers σ in the large-sample limit. Two properties of this scaling are particularly important for our protocol. First, it places the robust scale estimate on the same numerical units as the ordinary sample standard deviation, so the standardized deviations $z_{m,\ell}(x)$, the metric-specific alignments $\hat{z}_{m,\ell}(x)$, and the aggregated score $S_{\text{LaRA}}(x)$ retain their interpretation as approximate Gaussian-style z -scores. Consequently, replacing the standard deviation with a robust scale estimator does not implicitly retune downstream thresholds or alter the semantic interpretation of the score.




Figure 11: **Results over different layer windows on three models.**

Second, unlike the sample standard deviation, which has a 0% breakdown point and can be driven arbitrarily large by a single extreme outlier, the MAD achieves a 50% breakdown point and therefore remains stable even when the clean-reference pool $\mathcal{D}^{\text{clean}}$ contains a substantial fraction of atypical samples (Hampel et al., 1986; Rousseeuw and Leroy, 1987). This is particularly relevant for representation-geometry metrics, whose raw distributions are often heavy-tailed even after signed-log($1 + |\cdot|$) compression. The resulting robustness-efficiency trade-off is well established in the robust statistics literature: while the MAD is less asymptotically efficient than the sample standard deviation under perfectly Gaussian noise, it provides substantially improved stability under contamination and heavy-tailed deviations (Huber and Ronchetti, 2009; Maronna et al., 2019).

H Metrics and Baselines

H.1 Metrics

ROC-AUC measures the model’s ability to distinguish between member and non-member samples across all possible decision thresholds. It captures the overall separability of the two classes and is threshold-independent, making it a robust indicator of detection quality. **TPR@FPR=5%** reports the true positive rate (i.e., correctly identified members) when the false positive rate (i.e., non-members incorrectly flagged as members) is fixed at 5%. This reflects performance in the low false-positive regime, which is critical in contamination detection where incorrectly labeling clean samples as contaminated is costly. These metrics provide a comprehensive evaluation of global discrimination ability (ROC-AUC) and practical operating performance under strict error constraints (TPR@FPR=5%).

H.2 Baselines

We use six baselines to compare against our contamination detection protocol. (1) Recall (Xie et al., 2024), which probes memorization by measuring the model’s ability to regenerate ground-truth answers under controlled prompting; (2) CDD (Dong et al., 2024), which detects contamination via discrepancies in model predictions under input or prompt perturbations, based on the intuition that memorized samples are less sensitive to such changes; (3) Min-K% Prob (Shi et al., 2023), a likelihood-based metric that averages the log-probability over the lowest-probability tokens in a sequence, assuming memorized samples exhibit fewer low-confidence tokens; (4) Min-K%++ (Zhang et al., 2024), which extends Min-K% with improved normalization and calibration for greater robustness across settings; (5) PPL (Gonen et al., 2023), which measures sequence-level likelihood via perplexity, where unusually low values indicate potential memorization; and (6) Self-Critique (Tao et al., 2025), which leverages the model’s own reflective reasoning to assess contamination based on the confidence and consistency of its self-evaluation.

I Algorithm

Details of the LaRA algorithm is in Appendix I.

J Details of Generating Similar and Perturbed Questions

J.1 Prompts

To analyze representation dynamics under controlled perturbations, we use a three-stage prompt pipeline consisting of: (1) generating structurally similar questions, (2) identifying removable key information, and (3) generating paraphrased perturbation variants.

As shown in Table 11, we first generate semantically similar math problems that preserve the same reasoning structure and difficulty while modifying numerical values. This produces structurally matched control groups for representation comparison. Table 12 and 13 illustrates the prompt used to generate paraphrased variants of perturbed questions while preserving the exact position and semantic role of the [BLANK] placeholder. These variants enable measurement of local representation variability for computing RSI. Other prompts used to make the variants are in Table 14, 15, and 16.

J.2 Examples of the Generated Questions

We present the examples of the original and generated questions to qualitatively show their generations in Table 17. Other questions generated by other perturbation techniques are in Table 18, 19, and 20.

K Usage of AI Assistants

In preparing this work, we used AI-based writing assistants to improve sentence structure, correct grammatical errors, and enhance overall readability. These tools were employed solely for language refinement and did not contribute to the development of technical content, research methodology, or experimental analysis. All scientific ideas, results, and conclusions presented in the paper were conceived and authored entirely by the researchers. Use of AI assistance was restricted to editorial purposes and did not affect the originality or intellectual contributions of the work.

Algorithm 1: LaRA: Per-sample Layer-wise Representation Geometry Extraction

Input: Original question q_0 ; similar-question generator SIMILARGEN; importance-based blanking operator BLANKIMPORTANT producing k [BLANK] tokens; LLM paraphrase generator VARIANTGEN that preserves [BLANK] positions; mean-pooled hidden-state extractor $h_\ell(\cdot)$; layer set $\mathcal{L} = \{0, 1, \dots, L-1\}$ (every transformer layer); number of similar questions K ; number of paraphrase variants M ; number of [BLANK] tokens k ; numerical floor ϵ

Output: Per-layer geometric scores $\{\text{RSM}_\ell, \text{DC}_\ell, \text{RSI}_\ell\}_{\ell \in \mathcal{L}}$

```
{q1, ..., qK} ← SIMILARGEN(q0);
Q ← {q0, q1, ..., qK};
foreach qi ∈ Q do
    qi- ← BLANKIMPORTANT(qi, k);
    {vi,1, ..., vi,M} ← VARIANTGEN(qi-);
end
foreach ℓ ∈ L do
    /* (1) Representation Shift Magnitude → RSM */
    for i = 0 to K do
        ui ← hℓ(qi);
        wi ← hℓ(qi-);
        Δi ← ui - wi;
        Si ← ||Δi||2;
    end
    μS ←  $\frac{1}{K} \sum_{i=1}^K S_i$ ; // similars only
    σS ←  $\sqrt{\frac{1}{K-1} \sum_{i=1}^K (S_i - \mu_S)^2}$ ; // sample std
    RSMℓ ←  $\frac{S_0 - \mu_S}{\sigma_S + \epsilon}$ ;

    /* (2) Directional Collapse → DC */
    s̄ℓ ←  $\frac{1}{K} \sum_{i=1}^K \Delta_i$ ; // mean shift over similars
    DCℓ ←  $\frac{\Delta_0^\top \bar{s}_\ell}{(\|\Delta_0\|_2 + \epsilon)(\|\bar{s}_\ell\|_2 + \epsilon)}$ ; // cosine similarity

    /* (3) Representation Stability Index → RSI */
    for i = 0 to K do
        for m = 1 to M do
            φi,m ← hℓ(vi,m);
        end
        φ̄i ←  $\frac{1}{M} \sum_{m=1}^M \phi_{i,m}$ ;
        Ri ←  $\frac{1}{M} \sum_{m=1}^M \|\phi_{i,m} - \bar{\phi}_i\|_2$ ; // mean L2 distance
    end
    μR ←  $\frac{1}{K} \sum_{i=1}^K R_i$ ;
    σR ←  $\sqrt{\frac{1}{K-1} \sum_{i=1}^K (R_i - \mu_R)^2}$ ;
    RSIℓ ←  $\frac{R_0 - \mu_R}{\sigma_R + \epsilon}$ ;
end
return {RSMℓ, DCℓ, RSIℓ}ℓ ∈ L;
```

Settings	Content
Similar Questions	You are a math problem generator. Given an original math problem, create {num_questions} similar problems that: 1. Follow the EXACT same structure and solution method as the original 2. Use DIFFERENT numerical values (change all numbers to make the problem unique) 3. Maintain the same difficulty level 4. Have the same type of solution approach 5. Are valid, solvable problems Original Problem: {original_question} Generate {num_questions} similar problems. For each problem: - Change ALL numerical values to create unique scenarios - Keep the problem structure and mathematical concepts identical - Ensure the problem remains solvable and realistic - Make sure the new numbers create valid mathematical relationships Output ONLY a JSON array of {num_questions} similar problems, where each element is a string containing the full problem text. Do not include solutions or explanations, only the problems. Format your response as: [<code>"Problem 1 text here..."</code>, <code>"Problem 2 text here..."</code>]

Table 11: Prompt used for generating similar math questions.

Settings	Content
Perturbed Questions	You are a question editor that identifies key information to remove from math problems. Given a math problem, identify ONE key piece of information that should be removed. Describe this information in a way that can be consistently applied to similar problems. For example: - "the total number of residents/people" - "the initial quantity" - "the final result value" - "the time duration" - "the distance measurement" Original Problem: {original_question} Output ONLY a short description of what information type should be removed (e.g., "the total number of residents"). Do not include the actual value or explain why, just describe the information type in 5-10 words.

Table 12: Prompt used for identifying removable information in math questions.

Settings	Content
Perturbed Variants	You are a text rewriter that creates paraphrased versions of math problems. Given an incomplete math problem with [BLANK] placeholders, create {num_variants} paraphrased versions that: 1. Preserve the EXACT position and meaning of [BLANK] - do NOT move or change [BLANK] 2. Use different wording and phrasing while maintaining the same mathematical meaning 3. Keep the same structure and logical flow 4. Do NOT reveal what the blank should be 5. Maintain all mathematical relationships and constraints Incomplete Problem: {incomplete_question} Output ONLY a JSON array of {num_variants} strings with the paraphrased versions. Do not include explanations or notes.

Table 13: Prompt used for generating perturbed paraphrased variants.

Setting	Content
Perturbation Analysis (Variable Renaming)	You edit math problems for robustness testing. For each problem, replace EXACTLY ONE single-letter variable (e.g. x, y, n, t) with a random common English noun (not a single letter). Example: replace "x" with "widget" everywhere that variable appears in the problem, or only once if the variable appears once – pick one occurrence if multiple variables exist. - Do NOT use [BLANK]. - Keep all numbers and the problem structure unchanged except for that substitution. Problems: {numbered} Output ONLY a JSON array of strings with the same length and order as the input.

Table 14: Prompt used for perturbation analysis through renaming variable.

Setting	Content
Perturbation Analysis (Number Replacement)	You edit math problems for robustness testing. For each problem, change EXACTLY ONE numeric literal to a different numeric literal. - Do NOT use [BLANK]. - Do NOT change anything else (wording, variables, structure). - The new number must differ from the original. Problems: numbered Output ONLY a JSON array of strings with the same length and order as the input.

Table 15: Prompt used for perturbation analysis through number replacement.

Setting	Content
Perturbation Analysis (Insert Distractor)	You edit math problems for robustness testing. For each problem, insert EXACTLY ONE short distractor sentence (one clause) that is irrelevant or misleading but plausible-sounding. Place it naturally in the problem text. - Do NOT use [BLANK]. - Do not change the original numbers or the core question being asked. Problems: numbered Output ONLY a JSON array of strings with the same length and order as the input.

Table 16: Prompt used for perturbation analysis through inserting distractor.

Settings	Content
Original Question	A plane contains points A and B with $AB = 1$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{4}{3}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$. Present the answer in LaTeX format: Your answer
Similar Question	A plane contains points A and B with $AB = 2$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{5}{4}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$.
Perturbed Question	A plane contains points A and B with $AB = [BLANK]$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{4}{3}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$.
Perturbed Variant	In a plane, there are points A and B such that $AB = [BLANK]$. Point A is rotated counterclockwise around point B through an acute angle θ, resulting in point A'. Next, point B is rotated clockwise around point A' through the same angle θ, leading to point B'. It is given that $AB' = \frac{9}{8}$. The value of $\cos \theta$ can be expressed as $\frac{m}{n}$, where m and n are positive integers with no common factors. Determine the sum $m + n$.

Table 17: Examples of Target, Similar, and Perturbed Questions

Settings	Content
Original Question	A plane contains points A and B with $AB = 1$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point $A^{/prime}$. Then B is rotated in the plane clockwise through angle θ around point $A^{/prime}$ to point $B^{/prime}$. Suppose that $AB^{/prime} = \frac{4}{3}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$. Present the answer in LaTeX format: Your answer
Similar Question	A plane contains points A and B with $AB = 2$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = 3$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$. Present the answer in LaTeX format: Your answer
Perturbed Question	A plane contains points A and B with $AB = 1$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{4}{3}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$.
Perturbed Variant	In a plane, there are points A and B such that $AB = 1$. Point A is rotated around point B counterclockwise through an acute angle θ to reach point A'. Next, point B is rotated clockwise around point A' through the same angle θ, resulting in point B'. Given that $AB' = \frac{4}{3}$, express the value of $\cos \theta$ as $\frac{m}{n}$, where m and n are coprime positive integers. Determine $m + n$.

Table 18: Examples of Target, Similar, and Perturbed Questions in Variable Renaming Perturbation.

Settings	Content
Original Question	A plane contains points A and B with $AB = 1$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{4}{3}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$. Present the answer in LaTeX format: Your answer
Similar Question	A plane contains points A and B with $AB = 2$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{5}{4}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$.
Perturbed Question	A plane contains points A and B with $AB = 3$. Point A is rotated in the plane counterclockwise through an acute angle θ around point B to point A'. Then B is rotated in the plane clockwise through angle θ around point A' to point B'. Suppose that $AB' = \frac{7}{3}$. The value of $\cos \theta$ can be written as $\frac{m}{n}$, where m and n are relatively prime positive integers. Find $m + n$.
Perturbed Variant	In a plane, there are points A and B such that the distance AB equals 3. Point A is rotated counterclockwise by an acute angle θ around point B to reach point A'. Next, point B is rotated clockwise by angle θ around point A' to arrive at point B'. If the distance AB' is $\frac{7}{3}$, then the value of $\cos \theta$ can be expressed as $\frac{m}{n}$, where m and n are relatively prime positive integers. Calculate $m + n$.

Table 19: Examples of Target, Similar, and Perturbed Questions in Number Replacement Perturbation.

Settings	Content
Original Question	Isosceles triangle $\triangle ABC$ has $AB = BC$. Let I be the incenter of $\triangle ABC$. The perimeters of $\triangle ABC$ and $\triangle AIC$ are in the ratio $125 : 6$, and all the sides of both triangles have integer lengths. Find the minimum possible value of AB . Present the answer in LaTeX format: Your answer
Similar Question	Isosceles triangle $\triangle DEF$ has $DE = EF$. Let J be the incenter of $\triangle DEF$. The perimeters of $\triangle DEF$ and $\triangle DEJ$ are in the ratio $100 : 5$, and all the sides of both triangles have integer lengths. Find the minimum possible value of DE .
Perturbed Question	Isosceles triangle $\triangle DEF$ has $DE = EF$. Let J be the incenter of $\triangle DEF$. The perimeters of $\triangle DEF$ and $\triangle DEJ$ are in the ratio $100 : 5$, and all the sides of both triangles have integer lengths, as is common in many geometric problems . Find the minimum possible value of DE .
Perturbed Variant	In isosceles triangle $\triangle DEF$, the sides satisfy $DE = EF$. Denote the incenter of $\triangle DEF$ as point J . The ratio of the perimeters of $\triangle DEF$ to $\triangle DEJ$ is $100 : 5$, and all sides of both triangles are integers, which is typical in geometric scenarios. Determine the smallest possible value for DE .

Table 20: Examples of target, similar, and perturbed questions in distractor insertion perturbation.