

DOES MATH REASONING IMPROVE GENERAL LLM CAPABILITIES? UNDERSTANDING TRANSFERABILITY OF LLM REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Math reasoning has become the poster child of progress in large language models (LLMs), with new models rapidly surpassing human-level performance on benchmarks like MATH and AIME. But as math leaderboards improve week by week, it is worth asking: *do these gains reflect broader problem-solving ability or just narrow overfitting?* To answer this question, we evaluate over 20 open-weight reasoning-tuned models across a broad suite of tasks, including math, scientific QA, agent planning, coding, and standard instruction-following. We surprisingly find that most models that succeed in math fail to transfer their gains to other domains. To rigorously study this phenomenon, we conduct controlled experiments on Qwen3-14B models using math-only data but different tuning methods. We find that reinforcement learning (RL)-tuned models **transfer** well across domains, while supervised fine-tuning (SFT)-tuned models often forget general capabilities. Latent-space representation and token-space distribution shift analyses reveal that SFT induces substantial representation and output drift, while RL preserves general-domain structure. Our results suggest a need to rethink standard post-training recipes, particularly the reliance on SFT-distilled data for advancing reasoning models.




Figure 1: Impact of SFT and RL using math-only training queries on the same base model, Qwen3-14B-Base. Performance improvements are measured relative to the base model. SFT-trained models show limited transfer to non-reasoning tasks. In contrast, RL-trained models exhibit **better performance** across both reasoning and non-reasoning scenarios.

1 INTRODUCTION

Over the past years, the community has raced to push large language models (LLMs) to new heights on math-centric reasoning benchmarks such as MATH (Hendrycks et al., 2021b) and AIME. A steady stream of *reasoning-tuned* models (Muennighoff et al., 2025; Guha et al., 2025) now advances the state of the art on math benchmarks nearly every week, with some even surpassing the average performance of human experts (Team, 2025a; OpenAI, 2024).

The appealing performance on math reasoning is understandable: problems are well-posed, solutions are unambiguous, and evaluation is easily verifiable, often just a single number or expression (Luo et al., 2025). This clarity has made math a popular proxy task of LLM reasoning, and

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

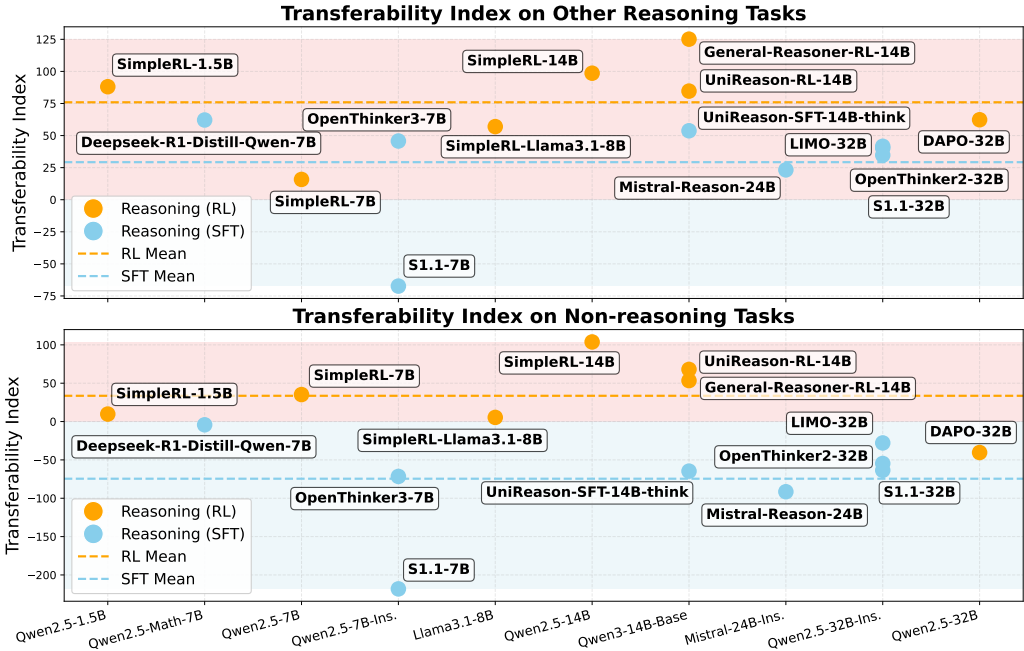


Figure 2: Transferability of mathematical reasoning to other reasoning and non-reasoning tasks. The Transferability Index measures a model’s ability to transfer performance from mathematics to other domains, with positive values indicating successful transfer and negative values indicating performance degradation. Details of this metric can be found in Section 2.1. RL models consistently outperform SFT models, regardless of model size, architecture, or training data, demonstrating superior transferability.

researchers have developed increasingly sophisticated training recipes to maximize model performance on mathematical reasoning (Wang et al., 2024a; Yue et al., 2024a; Luo et al., 2023; Shao et al., 2024; Wei et al., 2023). This trend, on one hand, should be encouraged. Mathematics is often considered the foundational language of science, and enabling machines to reason precisely over math is central to the long-term vision of automated scientific discovery (Mishra et al., 2022). On the other hand, real-world tasks extend far beyond math. The majority of user-facing applications, question answering, dialogue, instruction following, require broader linguistic and commonsense competence that math alone does not test (Ma et al., 2025).

This raises a natural question: *Do improved math reasoning abilities transfer to general LLM capabilities?* Specifically, can gains in solving math problems transfer to other reasoning domains (e.g., scientific QA (Welbl et al., 2017), coding (Jain et al., 2025), agent planning (Xie et al., 2024), logical deduction (Dziri et al., 2024)) and to tasks (e.g., conversational QA (Reddy et al., 2019), instruction following (Zhou et al., 2023)) that do not require extensive reasoning? To investigate, we evaluate over 20 representative open-weight reasoning models, all of which exhibit impressive performance on recent math benchmarks across a suite of other reasoning and non-reasoning tasks. To quantitatively analyze this problem, we propose *Transferability Index*, a metric to measure how reasoning models can transfer their capabilities from one domain to another. Surprisingly, as shown in Figure 2, we find that some of these models *fail* to transfer their improved mathematical reasoning capabilities to other domains while others succeed.

What drives this divergence? Model recipes vary widely in size, data distribution, and architecture. Yet among various parts, we identify one factor that consistently predicts transferability: the fine-tuning *paradigm*. Across families and sizes, models fine-tuned using reinforcement learning (RL) (Su et al., 2025; Yeo et al., 2025) exhibit much stronger performance in non-math tasks than those trained with supervised fine-tuning (SFT) (Yue et al., 2024a;b), which often show signs of catastrophic forgetting over a wide range of non-math tasks.

To validate this observation, we conduct a controlled study. We fine-tune Qwen3-14B (Team, 2025b) on the high-quality math dataset derived from MATH and DeepScaler (Luo et al., 2025). For SFT,

we construct targets via rejection sampling using Qwen3-32B, keeping only teacher responses that yield correct final answers. For RL, we apply a standard GRPO (Shao et al., 2024) setup using answer correctness as the reward. As shown in Figure 1, the results mirror our large-scale audit: RL-tuned models transfer well to non-math domains, despite being trained solely on math queries, while SFT-tuned models do not.

To better understand why this occurs, we use two diagnostic tools: (1) latent-space principal component analysis (PCA) on hidden states across layers, and (2) KL-divergence on token distributions before and after fine-tuning. These methods allow us to quantify how much the model’s internal representations and output space shift during training. We find that SFT induces significant drift in both latent and output spaces, especially for non-reasoning inputs, whereas RL better preserves the geometry of internal features and the stability of the token distributions. [Building on these initial observations, we aim to refine our understanding of which concrete mechanisms drive the differences between SFT and RL. This motivates a closer examination of the structure of both post-training strategies.](#) Specifically, we conduct an ablation study to learn which components contribute to the significant variances in transferability. Results show that the sampling distribution, credit assignment, and negative gradient contribute more, while KL regularization plays only a subtle role in the transferability of LLM reasoning.

2 PHENOMENA: PERFORMANCE DISCREPANCIES OF REASONING MODELS

2.1 OBSERVED PERFORMANCE DISCREPANCIES BETWEEN RL AND SFT

Setup. We evaluate over 20 off-the-shelf reasoning models on different benchmarks. Namely, we select benchmarks from the following three categories: (1) **math reasoning tasks**: MATH500 (Hendrycks et al., 2021b), AIME24, AIME25, OlympiadBench (He et al., 2024), which contain mathematical problems only; (2) **other reasoning tasks**: LiveCodeBench (Jain et al., 2025), GPQA-Diamond (Rein et al., 2024), ACPBench (Kokel et al., 2025), HeadQA (Vilares & Gómez-Rodríguez, 2019), which contain more general reasoning questions, such as medical reasoning, code generation, and language-based agent planning tasks; (3) **non-reasoning tasks**: CoQA (Reddy et al., 2019), IFEval (Zhou et al., 2023), HaluEval (Li et al., 2023), MC-TACO (Zhou et al., 2019), which contain factual, alignment, or conversational problems such as commonsense question answering and instruction-following. We used *accuracy* to evaluate the models’ performance. Detailed explanation about experiment setup, benchmarks, and evaluation metrics can be found in Appendix A.6.

To better evaluate the model’s transferability across a wide range of task groups, we define **Transferability Index (TI)** as follows:

(1) *Per-benchmark gain and normalization.* z -normalization makes gains comparable within each group. For each group $g \in \{\text{math}, \text{other}, \text{non}\}$ and benchmark $b \in \mathcal{B}_g$,

$$\Delta R_b = R_b^{\text{model}} - R_b^{\text{base}}, \quad \sigma_g = \text{Std}\{\Delta R_b : b \in \mathcal{B}_g\}, \quad \delta_b = \frac{\Delta R_b}{\sigma_g}.$$

(2) *Robust gain score and difficulty weighting.* We then compute the domain index to normalize benchmark differences. We control extremes with a signed square-root and up-weight harder tasks, and used a group-level *Domain Index* aggregates these weighted, robust gains:

$$s_b = \text{sign}(\delta_b) |\delta_b|^{1/2}, \quad w_b = 100 - R_b^{\text{base}}, \quad \hat{w}_b = \frac{w_b}{\sum_{u \in \mathcal{B}_g} w_u}, \quad \text{DI}_g = \sum_{b \in \mathcal{B}_g} \hat{w}_b s_b.$$

(3) *Transferability relative to math.* Finally, TI is the ratio of the group’s improvement to the math group’s improvement:

$$\text{TI}_g(\%) = \frac{\text{DI}_g}{\text{DI}_{\text{math}}} \times 100, \quad g \in \{\text{other}, \text{non}\}.$$

A positive TI_g indicates positive transfer in group g , and higher TI_g indicates the finetuned model achieves stronger and more consistent improvements in transferring math capabilities to other domains. To figure out the effect of introducing signed square-root into TI, we conducted an ablation study in Appendix A.3, compared

Table 1: Performance of our model and baseline models on math reasoning, other reasoning and non-reasoning tasks. UniReason-Qwen3-14B-*think* (SFT) denotes the model is trained with the dataset distilled from Qwen3-32B thinking mode, UniReason-Qwen3-14B-*no-think* (SFT) denotes it is trained with Qwen3-32B non-thinking mode. The grey-colored row represents the performance difference between RL and the highest SFT score. [Additional studies on Qwen3-4B, Llama3.1-8B models; training with more diverse training data and on more evaluation benchmarks could be found in Appendix A.2.](#)

Math Reasoning						
Model	AIME24	AIME25	MATH500	Olympiad	Average	
Qwen3-14B-Base	13.0	9.3	60.4	27.9	27.7	-
UniReason-Qwen3-14B- <i>think</i> (SFT)	52.0	37.0	85.0	25.0	49.8	-
UniReason-Qwen3-14B- <i>no-think</i> (SFT)	16.0	13.0	77.2	22.7	32.3	-
UniReason-Qwen3-14B (RL)	55.7	38.0	87.8	33.8	53.8	-
Δ (RL - SFT)	+3.7	+1.0	+2.8	+8.8	+4.1	-
Other Reasoning						
Model	GPQA	LiveCodeBench2	ACPBench	HeadQA	Average	TI _{other}
Qwen3-14B-Base	42.6	29.7	10.7	37.6	30.2	-
UniReason-Qwen3-14B- <i>think</i> (SFT)	55.9	21.9	68.6	34.8	45.3	+52.2
UniReason-Qwen3-14B- <i>no-think</i> (SFT)	48.7	23.5	69.3	35.0	45.2	+165.4
UniReason-Qwen3-14B (RL)	57.7	40.6	65.4	40.2	60.0	+82.3
Δ (RL - SFT)	+1.8	+17.1	-3.9	+5.2	+5.1	-
Non-Reasoning						
Model	CoQA	MC-TACO	IFEval	HalluEval	Average	TI _{non}
Qwen3-14B-Base	10.0	67.7	69.2	35.7	45.7	-
UniReason-Qwen3-14B- <i>think</i> (SFT)	1.7	38.2	42.3	2.3	21.1	-104.1
UniReason-Qwen3-14B- <i>no-think</i> (SFT)	5.3	66.1	41.4	3.3	29.0	-278.9
UniReason-Qwen3-14B (RL)	28.2	74.0	70.0	40.7	53.2	+52.2
Δ (RL - SFT)	+22.9	+7.9	+27.7	+37.4	+24.0	-

Results. Figure 2 and Table 11 compare transferability indices across models. Transferability varies with finetuning method, model size, and architecture; the dominant factor is the *method*: RL-tuned models consistently achieve higher indices on both other-reasoning and non-reasoning tasks, whereas SFT-trained models often yield negative TI_{non} . This suggests that on-policy RL reinforces reasoning while minimally perturbing general-domain representations, whereas off-policy SFT on large static reasoning corpora can over-specialize the latent space, degrading non-reasoning performance.

2.2 CONTROL STUDY

Motivated by our findings in Section 2.1, we design a light-weight controlled study to directly compare SFT and RL on an identical dataset. Concretely, we start from a small, high-quality mathematics dataset (see Appendix A.6.2 for details), then query a strong teacher model (Qwen3-32B-Instruct) to extract complete chain-of-thought (CoT) reasoning traces with reject sampling. These CoT traces become our SFT training targets, while the original answer labels serve as the rewards for RL. This alignment ensures both paradigms learn from the same data samples. Then, we take the Qwen3-14B-Base model and fine-tune it in two ways: (i) *SFT* on the teacher-generated CoT traces; (ii) *RL* using only the groundtruth. We name our model **UniReason**. We compare against the Qwen3-14B-Base model. Evaluation is conducted on three benchmark groups mentioned above using *accuracy*. Details about training datasets, baseline models, and hyperparameters could also be found in Appendix A.6. Our experimental results on three groups of benchmarks (see Table 1) reveal a consistent pattern:

- On **math reasoning**, our UniReason-Qwen3-14B(RL) model climbs to 55.7% on AIME24, 87.8% on MATH500, and 33.8% on OlympiadBench, outperforming corresponding SFT-based models.
- For **other reasoning** tasks, SFT-based models make uneven progress (e.g. UniReason-Qwen3-14B-SFT-*think* scores 55.9% on GPQA), whereas RL fine-tuning yields significant lifts: UniReason-Qwen3-14B(RL) gains 1.8% on GPQA, and 17.1% on LiveCodeBench2 over SFT.

- Crucially, in **non-reasoning** evaluations, SFT models stagnate or decline, while the RL model recover and exceed the base in nearly all the benchmarks.
- From Table 1, our RL-tuned model achieve positive TI_{other} and TI_{non} , meaning its math improvements also transfer to other reasoning and even boost non-reasoning performance. In contrast, SFT-based models often have negative TI_{non} , indicating poor transfer beyond reasoning tasks.

These results show that RL-tuned reasoning models perform generally better than SFT-based models on both reasoning and non-reasoning tasks when carefully controlling other factors. Especially, our UniReason model is trained on a single distilled math dataset, but it still preserves and even improves general-domain performance while showing strong reasoning gains. To further validate our findings under different model sizes, learning rates and architectures, we applied this setting to Qwen3-4B model and Llama3.1-8B model, the results could be found in Appendix A.2. We also evaluated our model on several more benchmarks. The results show a consistent pattern with our control study.

3 LATENT REPRESENTATION SHIFTS: INSIGHTS FROM PCA ANALYSIS

As discussed in Section 2.1, applying SFT to the Qwen model improves reasoning abilities such as mathematical problem-solving and code generation, but impairs general-domain performance. We observe that most SFT models fail to transfer their improved mathematical reasoning capabilities to other domains. In contrast, our controlled study shows that RL models transfer well to non-math domains, despite being trained solely on math queries, whereas SFT models do not.

To understand the underlying cause of this transferability gap, we employ PCA shift analysis to examine how the internal feature geometry of the model evolves under different training paradigms, model sizes, and model families across diverse query distributions. Recent studies (Xu et al., 2025; Zheng et al., 2025) demonstrate that PCA shift analysis provides a sensitive and interpretable measure of representational changes relevant to task performance. Importantly, changes in model parameters do not always correspond to functional differences: large weight updates may leave outputs unchanged, while subtle parameter modifications can lead to significant shifts in the activation distribution. By focusing on hidden representations, PCA shift directly captures how the model encodes

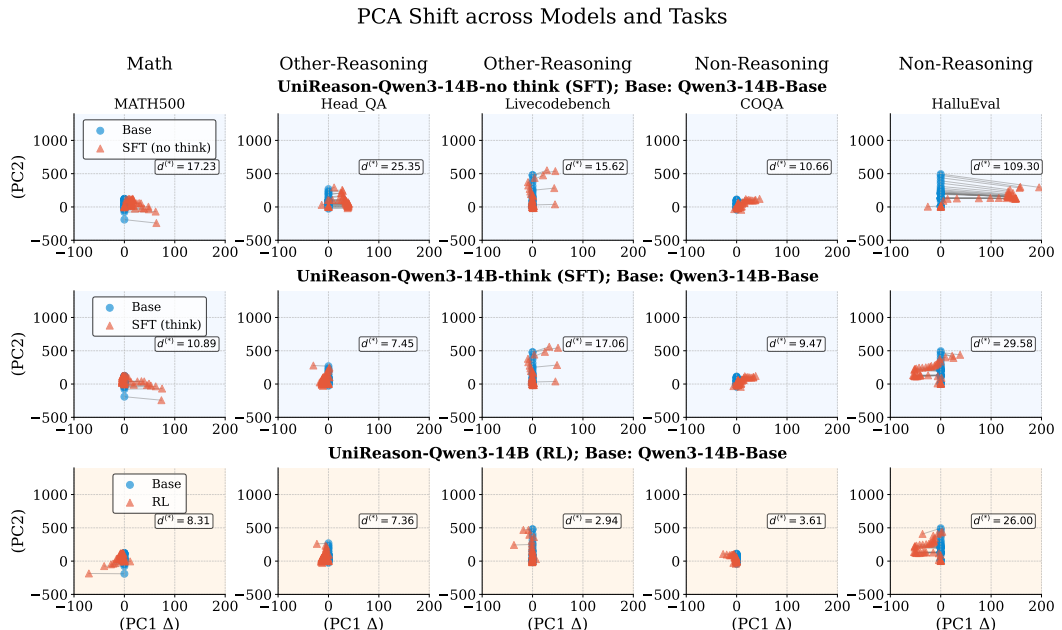


Figure 3: PCA shift of Qwen3-14B-Base across different training methods and tasks. $d^{(*)}$ is the Euclidean distance between representation centroids before and after training. The first two rows show models trained with SFT, and the last row shows models trained with RL. RL training results in the smallest PCA shift for all task types, suggesting more stable latent representations.

Table 2: Mean PCA shift magnitudes on math, reasoning, and non-reasoning tasks for Qwen3-14B and our finetuned models. RL models exhibit lower shifts than SFT models.

Model	Math	Other-Reasoning	Non-Reasoning
Qwen3-14B- <i>no-think</i>	40.4	13.9	129.8
Qwen3-14B- <i>think</i>	76.5	38.8	152.0
UniReason-Qwen3-14B- <i>no-think</i> (SFT)	21.4	10.9	113.7
UniReason-Qwen3-14B- <i>think</i> (SFT)	19.2	6.7	38.2
UniReason-Qwen3-14B (RL)	8.5	3.5	36.9

and processes information, offering a more faithful account of its internal knowledge state than parameter-based metrics. This perspective allows us to distinguish between true knowledge erasure and parameterization changes that leave the underlying feature space intact. Furthermore, since transferability fundamentally relies on the alignment and stability of learned representations across tasks or domains, PCA shift is particularly effective for diagnosing changes that may impact cross-domain performance. Shifts in principal components reveal whether the model’s internal feature space remains suitable for knowledge transfer or has been disrupted by post-training interventions.

3.1 EXPERIMENTAL SETUP

Models and Tasks. In Section 2.1, we observe that models trained on math datasets show moderate transferability on other reasoning tasks. We perform PCA shift analyses on the corresponding models and tasks, aiming to critically assess these phenomena from a feature-space perspective.

Evaluation. Given input queries \mathcal{X} , we extract hidden states $\mathbf{H}_i^{(*)}$ at each layer i for each model state $(*) \in \{\text{orig}, \text{updated}\}$. Applying PCA ($n = 2$) to $\mathbf{H}_i^{(*)}$, we compute the mean projection $m_{i,1}^{(*)}$ onto the first principal direction (PC1) and $m_{i,2}^{(*)}$ onto the second (PC2). The PCA shift is defined as $\Delta m_{i,1}^{(*)} = m_{i,1}^{(*)} - m_{i,1}^{\text{orig}}$ for PC1, while for PC2, we directly report $m_{i,2}^{(*)}$ as an indicator of change.

3.2 INVESTIGATING LATENT SPACE SHIFT

To quantify the overall latent shift, we define a *representation center* for each model state $(*)$ as the mean of PCA-projected coordinates across all layers: $\mathbf{z}^{(*)} = \frac{1}{L} \sum_{i=1}^L \mathbf{z}_i^{(*)}$, where L denotes the total number of layers and $\mathbf{z}_i^{(*)} = (\Delta m_{i,1}^{(*)}, m_{i,2}^{(*)})$ is the vector of PCA coordinates for layer i in state $(*)$. The latent shift between two model states, such as the original (base) and an updated model, is then measured by the Euclidean distance: $d^{(*)} = \|\mathbf{z}^{(*)} - \mathbf{z}^{\text{orig}}\|_2$.

Based on the analyses in Appendix A.7, RL-based training proves essential for developing language models that maintain a strong balance between general-domain and reasoning capabilities. Motivated by this, we further analyze our proposed models in the controlled study. As shown in Table 2, RL models achieve the lowest PCA shift magnitudes across math, other-reasoning, and non-reasoning tasks. Figure 3 further supports these findings, illustrating that the RL model consistently yields minimal and tightly clustered latent shifts across diverse benchmarks. In contrast, SFT models, particularly those without explicit reasoning signals, exhibit more scattered shifts. These results and the evaluations in Section 2.1, highlight the advantage of RL over SFT. They underscore the importance of a well-balanced optimization objective rather than isolated interventions, for mitigating catastrophic forgetting while preserving performance in LLMs.

4 TOKEN DISTRIBUTION SHIFTS: INSIGHTS FROM KL DIVERGENCE AND RANK ANALYSES

In this section, we conduct token-level analyses to further examine the distribution shift of RL and SFT models trained on mathematical reasoning data.




Figure 4: KL divergence analysis of RL and SFT models. Higher KL divergence indicates greater distribution shifts from the original backbone model. We observe that RL models consistently exhibit lower KL divergence compared to SFT models across different tasks, suggesting less distribution shift during training.

4.1 SETUP

KL-divergence serves as a standard metric for measuring differences between probability distributions. For token rank shift analysis, we first generate tokens using the fine-tuned model, then decode these same tokens using the backbone model to determine their original ranking positions. The rank shift is calculated as the difference in token rankings between the fine-tuned model and the backbone model for each token (Li et al., 2025c; Lin et al., 2023). Following the observations in Section 2.1, we perform additional token-distribution analyses on the corresponding models and tasks to assess the model distribution shift from a token-space perspective. Specifically, we employ KL-divergence and token rank shift metrics to analyze distribution shifts between models.

4.2 RESULT

RL models exhibit lower KL-divergence from backbone models. In Figure 4, we observe that the KL divergence of SFT models on both reasoning and non-reasoning tasks is significantly larger than that of RL models. This indicates that RL models exhibit substantially less distribution shift from the token distribution level compared to SFT models. For instance, UniReason-Qwen3-14B-SFT-no-think demonstrates KL divergences of 0.372 and 0.283 on MATH-500 and IFEval respectively compared to the backbone model, whereas UniReason-Qwen3-14B(RL) achieves considerably lower KL divergences of only 0.084 and 0.019 on the corresponding tasks.




Figure 5: Word clouds showing significantly shifted tokens between UniReason-Qwen3-14B-RL (left) and UniReason-Qwen3-14B-SFT-think (right). Tokens are extracted based on frequency and rank shifts compared with base model then categorized as logical-structural words (in red) or content-specific words (in blue). The RL model promptly shifts logic-related tokens such as *But* and *So* while the SFT model shifts including many irrelevant tokens.

RL models demonstrate reduced token rank shifts. Across both reasoning and non-reasoning tasks, RL models show markedly smaller average token-rank shifts than SFT models as shown in Figure 17 in Appendix A.9. For example, UniReason-Qwen3-14B (RL) averages only 0.98 positions of shift, whereas the SFT no-think variant reaches 10.6. This indicates RL tends to preserve the base token distribution while SFT induces larger reordering. Detailed position-wise trajectories are provided in Figure 20 in Appendix A.9.

RL models selectively shift task-relevant tokens, while SFT models shift numerous irrelevant tokens. A case study in Table 14 shows RL moving a small set of task-relevant tokens (e.g., “define”, “add”, “second”, “number” for reasoning; “<<”, “>>”, “write”, “formally” for non-reasoning), while SFT shifts far more tokens (e.g., 390 for reasoning / 158 for non-reasoning), including many that are query-irrelevant and often inject reasoning markers into non-reasoning prompts. The token-frequency word cloud in Figure 5 further visualizes this selectivity gap. More detailed information could be found in Appendix A.9.

5 ABLATION: WHICH COMPONENTS OF RL DRIVE GENERALIZATION?

Earlier sections observed that reinforcement-style fine-tuning generalizes more effectively than standard supervised fine-tuning. Latent and output space analyses revealed that RL updates produce small, targeted “marginal shifts” around the base model, whereas SFT often causes substantial drift. This section sharpens the inquiry: *Which specific components of RL updates contribute to improved generalization?*

We contrast SFT and RL by aligning their objectives and gradients, yielding a *unified surrogate loss* that makes their differences explicit along four levers: (i) **sampling distribution**: SFT trains *off-policy* on a fixed dataset, whereas RL samples *on-policy* from the live model; (ii) **credit assignment**: uniform weighting versus advantage-weighted updates that emphasize task-salient tokens; (iii) **KL regularization**: penalizing deviations from a reference policy such as PPO or GRPO; (iv) **negative gradient**: When sampling from π , RL learn from negative examples, while SFT only imitates positive ones.

5.1 CONTRASTING OBJECTIVES

For prompts x and completions y , let $\pi_\theta(y | x)$ denote the current policy, and $\pi_{\text{ref}}(y | x)$ denote a fixed reference policy (e.g., the initialization).

Supervised Fine-tuning. With reference completions y^* , the objective is:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{(x,y^*) \sim \mathcal{D}} [\log \pi_\theta(y^* | x)]. \quad (1)$$

Reinforcement Learning. For the same prompts, we sample $y \sim \pi_\theta$ and weight each sample by an *advantage* $A(x, y)$:

$$\mathcal{L}_{\text{RL}}(\theta) = -\mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim \pi_\theta(\cdot | x)} [A(x, y) \cdot \log \pi_\theta(y | x)]. \quad (2)$$

Usually, a KL term is added to prevent the policy model from being too far away from the initialized model. We generalize these objectives using:

$$\mathcal{L}_{q,w,\beta}(\theta) = -\mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim q(\cdot | x)} [w(x, y) \cdot \log \pi_\theta(y | x)] + \beta \mathbb{E}_{x \sim \mathcal{D}} [\text{KL}(\pi_\theta(\cdot | x) \| \pi_{\text{ref}}(\cdot | x))]. \quad (3)$$

5.2 ABLATION RESULTS

Setup. For ablation study, we used 5 settings from Off-policy and On-policy SFT, Off-policy and On-policy RL with and without KL. We used the same set of math training queries and a Qwen3-8B-Base model. For on-policy SFT: we sample a batch of training queries and let the policy model

Table 3: Ablation settings and their corresponding sampling distributions, credit weights, and KL-regularization coefficients. The two shaded rows mark the *standard baselines*: Off-policy SFT and Online RL.

Setting	Sampling q	Weights w	KL Reg. β
Off-policy SFT	$\delta_{y=y^*}$	1	0
On-policy SFT	π_θ	1 (reject sample)	0
Off-policy RL	$\delta_{y=y^*}$	Advantage A_t	0
On-policy RL (no KL)	π_θ	Advantage A_t	0
On-policy RL	π_θ	Advantage A_t	> 0

Table 4: Performance of all the ablation models. Scores are average scores on math reasoning, other reasoning and non-reasoning tasks.

Model	Math Avg.	Other Reasoning Avg.	Non-Reasoning Avg.	TI _{other}	TI _{non}
Qwen3-8B-Base	27.6	23.6	33.6		
Qwen3-8B-Off-policy-SFT	41.9	34.4	26.6	18.3	-40.5
Qwen3-8B-On-policy-SFT	33.7	35.7	35.0	68.6	30.2
Qwen3-8B-Off-policy-RL	45.5	35.9	31.7	36.4	4.5
Qwen3-8B-On-policy-RL (no KL)	37.1	38.2	35.8	65.6	39.3
Qwen3-8B-On-policy-RL	38.6	39.9	35.0	63.7	32.4

generate the responses as in standard RL training. For off-policy RL, we use the Qwen3-32B model in think mode to generate $n = 8$ responses for each query. We utilized GRPO’s objective with KL divergence and used the response with the highest reward for gradient computation. For Off-policy SFT, we used data generated from Qwen3-32B model in think mode with reject sampling. On-policy RL used the same math dataset as our main models. The results are in Table 4. We also found that adding L2-SP regularization would not help SFT models transfer to other tasks under Qwen3-8B models as shown in Table 6.

Sampling distribution is critical. The results in Table 4 show that on-policy methods outperform off-policy methods across both evaluation categories and training paradigms. Specifically, on-policy SFT achieves higher average scores on Other Reasoning tasks and Non-Reasoning tasks compared to off-policy SFT. A similar trend is observed when KL regularization is applied to RL: on-policy RL continues to yield superior transferability on Other Reasoning and Non-Reasoning tasks. These results demonstrate that on-policy methods consistently transfer better than off-policy methods. One possible explanation is that on-policy training exposes models to data generated under their own evolving policy, allowing updates to directly align with the distribution the model will face at inference. This reduces distribution mismatch and mitigates overfitting to fixed offline samples, leading to stronger generalization. Gradient Norm during training also shows that on-policy methods update the gradient in small and conservative steps compared to off-policy methods which aggressively update its gradient especially from the start. Further details can be found in Figure 6 and Appendix A.8.

Credit assignments and negative examples matter. Experiments on credit assignment and negative gradients show that adding these mechanisms not only improves transferability but also increases response length during training. In off-policy settings, RL achieves higher average scores on both Other Reasoning and Non-Reasoning tasks. Under on-policy training, RL outperforms SFT on Math and Other Reasoning tasks, while showing similar performance on Non-Reasoning tasks. Credit assignment helps the model focus updates on the tokens that matter most for solving the task, while penalizing uninformative gradients reduces the chance of reinforcing wrong or irrelevant patterns. Together, these mechanisms encourage more useful exploration and lead to better generalization. Further details are provided in Figure 14.

KL regularization play subtle roles. On-policy RL performance remains largely unchanged with or without KL regularization. This suggests that, although KL regularization constrains models to stay closer to their original policy during updates, its overall effect on transferability is limited. One possible reason is that on-policy RL already aligns updates with the on-policy distribution, making the additional constraint from KL regularization less impactful.

5.3 RELATED WORKS

Supervised Fine-Tuning vs. Reinforcement Learning for LLMs. Fine-tuning methods for reasoning typically fall into two major categories: supervised fine-tuning and reinforcement learning




Figure 6: Grad norm comparison across online and offline method. Online methods perform steadier updates. Offline updates show larger gradient norms in early steps, reflecting more abrupt shifts that influent generalization.

(Chen et al., 2024). SFT methods predominantly utilize annotated reasoning trajectories or solution traces, directly training models to replicate explicit reasoning sequences from datasets (Wei et al., 2022; Wang et al., 2023). RL-based fine-tuning, however, guides models by rewarding accurate and logically coherent reasoning steps without explicit step-by-step supervision, allowing exploration and optimization of reasoning pathways through feedback loops (Ziegler et al., 2019; Liu et al., 2025; Wang et al., 2024c; Chu et al., 2025).

Cross-domain improvement in Reasoning Models. Interestingly, models heavily fine-tuned for formal reasoning sometimes falter on more general language tasks (Kumar et al., 2022). For example, OpenAI’s o1, while excelling in STEM benchmarks, raised concerns about its versatility on other tasks (OpenAI, 2024). Follow-up research introduced reinforcement fine-tuning precisely to address this gap, aiming to adapt a generalist model’s reasoning to new domains with limited data (Zhang et al., 2024). Indeed, o1 and similar reasoning models are built on strong general-purpose bases to retain broad knowledge (OpenAI, 2024; Hendrycks et al., 2021a). Nonetheless, trade-offs have been observed. Wang et al. (2024b) found that fine-tuning on a narrow set of instruction types can degrade a model’s performance on other skills. Recent works have also stepped into analyzing the cross-domain performance of reasoning models (Sun et al., 2025), especially for RL-based approaches (Cheng et al., 2025; Hu et al., 2025).

Representation-Level Analysis. Fine-tuning for reasoning models not only boosts task performance but also alters the model’s internal representations (Sheng et al., 2024). Recent studies have begun to probe how CoT-based fine-tuning changes the *latent space* of LLMs (Xu et al., 2025a; Wang et al., 2025). Lobo et al. (2024) find that task-specific fine-tuning can reduce the faithfulness of a model’s generated reasoning chains, indicating shifts in its underlying inference mechanisms. Complementary analyses of hidden states provide insight into such shifts. Xu et al. (2024) proposes a quantitative framework for assessing ideas that leverages hidden representations from LLMs to predict the merit of scientific ideas. Techniques like principal component analysis further reveal that fine-tuning can carve out new directions in representation space that correspond to reasoning-related features (Xu et al., 2025a; Zhou et al., 2025).

6 CONCLUSION, LIMITATION AND DISCUSSION

We studied what drives *transferability* of reasoning models across reasoning and non-reasoning benchmarks. Our evidence shows that, beyond size and architecture, the fine-tuning paradigm is decisive: RL-tuned models improve math reasoning while preserving positive transfer to other reasoning and non-reasoning tasks, whereas SFT often induces negative transfer on non-reasoning benchmarks. Latent-space PCA further indicates that RL stays close to the backbone (stable features), while SFT yields larger shifts, especially on non-reasoning data. Token-distribution analysis aligns with this: RL selectively adjusts a small set of task-relevant tokens, whereas SFT perturbs many irrelevant ones. Notably, *UniReason-Qwen3-14B-RL* trained on 47K math examples achieves the strongest balance of reasoning gains and general-domain retention among compared models, supporting these conclusions. We then studied which factor in RL benefits generalization by a set of controlled study, and learned that sample distribution is the most important factor.

However, we acknowledge that our paper still has several limitations. For example, our control study could be extended beyond math to other domains such as coding or additional reasoning tasks, in order to provide a more comprehensive view of reasoning models’ transferability. On the post-training side, we primarily analyze SFT vs. RL, on-policy vs. off-policy training, and the role of credit assignment and negative examples; future work could incorporate more integrated paradigms, for example SFT followed by RL; RL followed by SFT; or the four stage post-training used by Deepseek and Qwen3 on SFT-RL-SFT-RL, or multi-stage training pipelines to further validate our conclusions.

Going forward, we plan to extend this approach in three directions: (1) exploring hybrid recipes to further optimize the trade-off between reasoning and non-reasoning tasks, spanning our current math-based approach to more reasoning domains; (2) developing adaptive regularization strategies that account for backbone architecture and size to enhance transferability; and (3) proposing more comprehensive tools to evaluate transferability in multi-modal, long-context, and interactive generation settings to ensure robust performance evaluation across an even wider range of real-world applications.

7 REPRODUCIBILITY STATEMENT

We make efforts to ensure reproducibility. All training configurations, hyperparameters, and evaluation setups are documented in the main text and Appendix. Datasets used are standard, publicly available benchmarks with clearly described preprocessing steps. To facilitate replication, we release an anonymous zip file containing source code, training scripts, and evaluation scripts as part of the supplementary material. Together, these resources allow independent researchers to reproduce both the training and evaluation results reported in this paper.

8 ETHICS STATEMENT

This work does not involve human subjects, private data, or sensitive information. All experiments are conducted on publicly available benchmarks, and no ethical concerns were identified. We believe our study complies with the ICLR Code of Ethics and raises no additional ethical issues.

REFERENCES

- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020. URL <https://arxiv.org/abs/1911.11641>.
- Jie Chen, Xintian Han, Yu Ma, Xun Zhou, and Liang Xiang. Unlock the correlation between supervised fine-tuning and reinforcement learning in training code large language models. *ArXiv preprint*, abs/2406.10305, 2024. URL <https://arxiv.org/abs/2406.10305>.
- Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. TheoremQA: A theorem-driven question answering dataset. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7889–7901, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.489. URL <https://aclanthology.org/2023.emnlp-main.489/>.
- Zhoujun Cheng, Shibo Hao, Tianyang Liu, Fan Zhou, Yutao Xie, Feng Yao, Yuexin Bian, Yonghao Zhuang, Nilabjo Dey, Yuheng Zha, Yi Gu, Kun Zhou, Yuqi Wang, Yuan Li, Richard Fan, Jian-shu She, Chengqian Gao, Abulhair Saparov, Haonan Li, Taylor W. Killian, Mikhail Yurochkin, Zhengzhong Liu, Eric P. Xing, and Zhiting Hu. Revisiting reinforcement learning for llm reasoning from a cross-domain perspective, 2025. URL <https://arxiv.org/abs/2506.14965>.
- Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation model post-training. *ArXiv preprint*, abs/2501.17161, 2025. URL <https://arxiv.org/abs/2501.17161>.
- Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jian, Bill Yuchen Lin, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena D. Hwang, Soumya Sanyal, Sean Welleck, Xiang Ren, Allyson Ettinger, Zaïd Harchaoui, and Yejin Choi. Faith and fate: Limits of transformers on compositionality. *Advances in Neural Information Processing Systems*, 36, 2024. URL <https://openreview.net/forum?id=Fkcckr3ya8>.
- Yichen Feng, Zhangchen Xu, Fengqing Jiang, Yuetai Li, Bhaskar Ramasubramanian, Luyao Niu, Bill Yuchen Lin, and Radha Poovendran. Visualsphinx: Large-scale synthetic vision logic puzzles for rl, 2025. URL <https://arxiv.org/abs/2505.23977>.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muenighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model evaluation harness, 2024. URL <https://zenodo.org/records/12608602>.

- 594 Etash Guha, Ryan Marten, Sedrick Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna
595 Nezhurina, Jean Mercat, Trung Vu, Zayne Sprague, et al. Openthoughts: Data recipes for reason-
596 ing models. *ArXiv preprint*, abs/2506.04178, 2025. URL [https://arxiv.org/abs/
597 2506.04178](https://arxiv.org/abs/2506.04178).
- 598
599 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi
600 Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun.
601 Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual mul-
602 timodal scientific problems. In *ACL (1)*, pp. 3828–3850, 2024. URL [https://doi.org/10.
603 18653/v1/2024.acl-long.211](https://doi.org/10.18653/v1/2024.acl-long.211).
- 604 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob
605 Steinhardt. Measuring massive multitask language understanding. In *9th International Confer-
606 ence on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenRe-
607 view.net, 2021a. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.
- 608 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn
609 Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset.
610 In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks
611 Track (Round 2)*, 2021b. URL <https://openreview.net/forum?id=7Bywt2mQsCe>.
- 612
613 Chuxuan Hu, Yuxuan Zhu, Antony Kellermann, Caleb Biddulph, Suppakit Waiwitlikhit, Jason
614 Benn, and Daniel Kang. Breaking barriers: Do reinforcement post training gains transfer to
615 unseen domains?, 2025. URL <https://arxiv.org/abs/2506.19733>.
- 616 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Ar-
617 mando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamina-
618 tion free evaluation of large language models for code. In *The Thirteenth International Confer-
619 ence on Learning Representations*, 2025. URL [https://openreview.net/forum?id=
620 chfJJYC3iL](https://openreview.net/forum?id=chfJJYC3iL).
- 621 Harsha Kokel, Michael Katz, Kavitha Srinivas, and Shirin Sohrabi. Acpbench: Reasoning about
622 action, change, and planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
623 volume 39, pp. 26559–26568, 2025. URL [https://ojs.aaai.org/index.php/AAAI/
624 article/view/34857](https://ojs.aaai.org/index.php/AAAI/article/view/34857).
- 625 Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can
626 distort pretrained features and underperform out-of-distribution, 2022. URL [https://arxiv.
627 org/abs/2202.10054](https://arxiv.org/abs/2202.10054).
- 628
629 Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. Race: Large-scale read-
630 ing comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*, 2017. URL
631 <https://arxiv.org/abs/1704.04683>.
- 632 Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brah-
633 man, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Ma-
634 lik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris
635 Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Ha-
636 jishirzi. Tulu 3: Pushing frontiers in open language model post-training, 2025. URL [https:
637 //arxiv.org/abs/2411.15124](https://arxiv.org/abs/2411.15124).
- 638 Junyi Li, Xiaoxue Cheng, Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. Halueval: A large-scale
639 hallucination evaluation benchmark for large language models. In *The 2023 Conference on Em-
640 pirical Methods in Natural Language Processing*, 2023. URL [https://openreview.net/
641 forum?id=bxsrykzSnq](https://openreview.net/forum?id=bxsrykzSnq).
- 642
643 Xiaomin Li, Zhou Yu, Zhiwei Zhang, Xupeng Chen, Ziji Zhang, Yingying Zhuang, Narayanan
644 Sadagopan, and Anurag Beniwal. When thinking fails: The pitfalls of reasoning for instruction-
645 following in llms, 2025a. URL <https://arxiv.org/abs/2505.11423>.
- 646 Yuetai Li, Zhangchen Xu, Fengqing Jiang, Bhaskar Ramasubramanian, Luyao Niu, Bill Yuchen Lin,
647 Xiang Yue, and Radha Poovendran. Temporal sampling for forgotten reasoning in llms, 2025b.
URL <https://arxiv.org/abs/2505.20196>.

- 648 Yuetai Li, Xiang Yue, Zhangchen Xu, Fengqing Jiang, Luyao Niu, Bill Yuchen Lin, Bhaskar Rama-
649 subramanian, and Radha Poovendran. Small models struggle to learn from strong reasoners.
650 *CoRR*, abs/2502.12143, February 2025c. URL [https://doi.org/10.48550/arXiv.](https://doi.org/10.48550/arXiv.2502.12143)
651 2502.12143.
- 652 Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu,
653 Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base llms: Rethinking alignment
654 via in-context learning, 2023. URL <https://arxiv.org/abs/2312.01552>.
- 656 Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. Logiqa: A
657 challenge dataset for machine reading comprehension with logical reasoning. *arXiv preprint*
658 *arXiv:2007.08124*, 2020. URL <https://arxiv.org/abs/2007.08124>.
- 660 Zeyuan Liu, Ziyu Huan, Xiyao Wang, Jiafei Lyu, Jian Tao, Xiu Li, Furong Huang, and Huazhe Xu.
661 World models with hints of large language models for goal achieving. In Luis Chiruzzo, Alan
662 Ritter, and Lu Wang (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas*
663 *Chapter of the Association for Computational Linguistics: Human Language Technologies (Vol-*
664 *ume 1: Long Papers)*, pp. 50–72, Albuquerque, New Mexico, April 2025. Association for Com-
665 putational Linguistics. ISBN 979-8-89176-189-6. doi: 10.18653/v1/2025.naacl-long.3. URL
666 <https://aclanthology.org/2025.naacl-long.3/>.
- 667 Elita Lobo, Chirag Agarwal, and Himabindu Lakkaraju. On the impact of fine-tuning on chain-of-
668 thought reasoning. *ArXiv preprint*, abs/2411.15382, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2411.15382)
669 2411.15382.
- 670 Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V
671 Le, Barret Zoph, Jason Wei, and Adam Roberts. The flan collection: Designing data and methods
672 for effective instruction tuning. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara
673 Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International*
674 *Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*,
675 pp. 22631–22648. PMLR, 23–29 Jul 2023. URL [https://proceedings.mlr.press/](https://proceedings.mlr.press/v202/longpre23a.html)
676 v202/longpre23a.html.
- 678 Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qing-
679 wei Lin, Shifeng Chen, and Dongmei Zhang. Wizardmath: Empowering mathematical reasoning
680 for large language models via reinforced evol-instruct. *ArXiv preprint*, abs/2308.09583, 2023.
681 URL <https://arxiv.org/abs/2308.09583>.
- 682 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin
683 Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing
684 o1-preview with a 1.5b model by scaling rl, 2025. URL [https://huggingface.co/](https://huggingface.co/agentica-org/DeepScaleR-1.5B-Preview)
685 agentica-org/DeepScaleR-1.5B-Preview.
- 687 Xueguang Ma, Qian Liu, Dongfu Jiang, Ge Zhang, Zejun Ma, and Wenhui Chen. General-reasoner:
688 Advancing llm reasoning across all domains. *ArXiv preprint*, abs/2505.14652, 2025. URL
689 <https://arxiv.org/abs/2505.14652>.
- 690 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor con-
691 duct electricity? a new dataset for open book question answering. In *Proceedings of the 2018*
692 *Conference on Empirical Methods in Natural Language Processing*, pp. 2381–2391, Brussels,
693 Belgium, 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1260. URL
694 <https://aclanthology.org/D18-1260>.
- 696 Swaroop Mishra, Matthew Finlayson, Pan Lu, Leonard Tang, Sean Welleck, Chitta Baral, Tanmay
697 Rajpurohit, Oyvind Tafjord, Ashish Sabharwal, Peter Clark, and Ashwin Kalyan. LILA: A unified
698 benchmark for mathematical reasoning. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang
699 (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Process-*
700 *ing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pp. 5807–5832.
701 Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.emnlp-main.392.
URL <https://doi.org/10.18653/v1/2022.emnlp-main.392>.

- 702 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
703 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. *sl: Simple test-time*
704 *scaling*, 2025. URL <https://arxiv.org/abs/2501.19393>.
- 705
706 OpenAI. *Learning to reason with llms*, 2024. URL [https://openai.com/index/](https://openai.com/index/learning-to-reason-with-llms/)
707 [learning-to-reason-with-llms/](https://openai.com/index/learning-to-reason-with-llms/).
- 708 Negin Raouf, Etash Kumar Guha, Ryan Marten, Jean Mercat, Eric Frankel, Sedrick Keh, Hritik
709 Bansal, Georgios Smyrnis, Marianna Nezhurina, Trung Vu, Zayne Rea Sprague, Mike A Mer-
710 rill, Liangyu Chen, Caroline Choi, Zaid Khan, Sachin Grover, Benjamin Feuer, Ashima Su-
711 varna, Shiye Su, Wanxia Zhao, Kartik Sharma, Charlie Cheng-Jie Ji, Kushal Arora, Jeffrey Li,
712 Aaron Gokaslan, Sarah M Pratt, Niklas Muennighoff, Jon Saad-Falcon, John Yang, Asad Aali,
713 Shreyas Pimpalgaonkar, Alon Albalak, Achal Dave, Hadi Pouransari, Greg Durrett, Sewoong Oh,
714 Tatsunori Hashimoto, Vaishaal Shankar, Yejin Choi, Mohit Bansal, Chinmay Hegde, Reinhard
715 Heckel, Jenia Jitsev, Maheswaran Sathiamoorthy, Alex Dimakis, and Ludwig Schmidt. *Auto-*
716 *matic evals for llms*, 2025. URL <https://github.com/mlfoundations/evalchemy>.
- 717 Siva Reddy, Danqi Chen, and Christopher D. Manning. *CoQA: A conversational question answering*
718 *challenge*. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019. doi:
719 10.1162/tacl_a.00266. URL <https://aclanthology.org/Q19-1016>.
- 720
721 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
722 Dirani, Julian Michael, and Samuel R. Bowman. *GPQA: A graduate-level google-proof q&a*
723 *benchmark*. In *First Conference on Language Modeling*, 2024. URL [https://openreview.](https://openreview.net/forum?id=Ti67584b98)
724 [net/forum?id=Ti67584b98](https://openreview.net/forum?id=Ti67584b98).
- 725 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
726 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. *Deepseekmath: Pushing the limits of mathe-*
727 *matical reasoning in open language models*, 2024. URL [https://arxiv.org/abs/2402.](https://arxiv.org/abs/2402.03300)
728 [03300](https://arxiv.org/abs/2402.03300).
- 729
730 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
731 Haibin Lin, and Chuan Wu. *Hybridflow: A flexible and efficient rlhf framework*. In *Proceedings*
732 *of the Twentieth European Conference on Computer Systems*, EuroSys '25, pp. 1279–1297, New
733 York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400711961. doi: 10.
734 1145/3689031.3696075. URL <https://doi.org/10.1145/3689031.3696075>.
- 735
736 Shuqian Sheng, Yi Xu, Tianhang Zhang, Zanwei Shen, Luoyi Fu, Jiaxin Ding, Lei Zhou, Xiaoying
737 Gan, Xinbing Wang, and Chenghu Zhou. *RepEval: Effective text evaluation with LLM rep-*
738 *resentation*. pp. 7019–7033, November 2024. doi: 10.18653/v1/2024.emnlp-main.398. URL
<https://aclanthology.org/2024.emnlp-main.398/>.
- 739
740 Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu.
741 *Crossing the reward bridge: Expanding rl with verifiable rewards across diverse domains*, 2025.
742 URL <https://arxiv.org/abs/2503.23829>.
- 743
744 Yiyoun Sun, Shawn Hu, Georgia Zhou, Ken Zheng, Hannaneh Hajishirzi, Nouha Dziri, and Dawn
745 Song. *Omega: Can llms reason outside the box in math? evaluating exploratory, compositional,*
and transformative generalization, 2025. URL <https://arxiv.org/abs/2506.18880>.
- 746
747 Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. *Commonsenseqa: A ques-*
748 *tion answering challenge targeting commonsense knowledge*. In *Proceedings of the 2019 Con-*
749 *ference of the North American Chapter of the Association for Computational Linguistics: Hu-*
750 *man Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4149–4158, 2019. URL
<https://arxiv.org/abs/1811.00937>.
- 751
752 DeepSeek-AI Team. *DeepSeek-RL: Incentivizing reasoning capability in llms via reinforcement*
753 *learning*, 2025a. URL <https://arxiv.org/abs/2501.12948>.
- 754
755 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming
Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han

- 756 Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze
757 Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin
758 Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi,
759 Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong,
760 Shaowei Liu, Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao,
761 Weimin Xiong, Weiran He, Weixiao Huang, Weixin Xu, Wenhao Wu, Wenyang He, Xianghui
762 Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles,
763 Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng
764 Liu, Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang,
765 Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang,
766 Ziyao Xu, Zonghan Yang, and Zongyu Lin. Kimi k1.5: Scaling reinforcement learning with llms,
767 2025. URL <https://arxiv.org/abs/2501.12599>.
- 768 Qwen Team. Qwen2.5 technical report, 2025b. URL <https://arxiv.org/abs/2412.15115>.
- 770 David Vilares and Carlos Gómez-Rodríguez. HEAD-QA: A healthcare dataset for complex rea-
771 soning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Lin-*
772 *guistics*, pp. 960–966, Florence, Italy, 2019. Association for Computational Linguistics. doi:
773 10.18653/v1/P19-1092. URL <https://aclanthology.org/P19-1092>.
- 774 Peiyi Wang, Lei Li, Zhihong Shao, R. X. Xu, Damai Dai, Yifei Li, Deli Chen, Y. Wu, and Zhifang
775 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. 2024a.
776 URL <https://arxiv.org/abs/2312.08935>.
- 778 Renxi Wang, Haonan Li, Minghao Wu, Yuxia Wang, Xudong Han, Chiyu Zhang, and Timothy
779 Baldwin. Demystifying instruction mixing for fine-tuning large language models. In *ACL (Student*
780 *Research Workshop)*, pp. 162–169, 2024b. URL <https://aclanthology.org/2024.acl-srw.15>.
- 782 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha
783 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language
784 models. In *The Eleventh International Conference on Learning Representations, 2023*. URL
785 <https://openreview.net/forum?id=1PL1NIMMrw>.
- 787 Yiming Wang, Pei Zhang, Baosong Yang, Derek F. Wong, and Rui Wang. Latent space chain-of-
788 embedding enables output-free LLM self-evaluation. In *The Thirteenth International Confer-*
789 *ence on Learning Representations, 2025*. URL <https://openreview.net/forum?id=jxo70B9fQo>.
- 791 Zihan Wang, Yunxuan Li, Yuexin Wu, Liangchen Luo, Le Hou, Hongkun Yu, and Jingbo Shang.
792 Multi-step problem solving through a verifier: An empirical analysis on model-induced pro-
793 cess supervision. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of*
794 *the Association for Computational Linguistics: EMNLP 2024*, pp. 7309–7319, Miami, Florida,
795 USA, November 2024c. Association for Computational Linguistics. doi: 10.18653/v1/2024.
796 findings-emnlp.429. URL <https://aclanthology.org/2024.findings-emnlp.429/>.
- 798 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi,
799 Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language
800 models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Ad-*
801 *vances in Neural Information Processing Systems, 2022*. URL https://openreview.net/forum?id=_VjQlMeSB_J.
- 803 Tianwen Wei, Jian Luan, Wei Liu, Shuang Dong, and Bin Wang. Cmath: Can your language model
804 pass chinese elementary school math test?, 2023. URL <https://arxiv.org/abs/2306.16636>.
- 806 Johannes Welbl, Nelson F. Liu, and Matt Gardner. Crowdsourcing multiple choice science questions.
807 In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pp. 94–106, Copenhagen,
808 Denmark, 2017. Association for Computational Linguistics. doi: 10.18653/v1/W17-4413. URL
809 <https://aclanthology.org/W17-4413>.

- 810 Annie Wong, Thomas Bäck, Aske Plaat, Niki van Stein, and Anna V. Kononova. Towards a
811 deeper understanding of reasoning capabilities in large language models, 2025. URL <https://arxiv.org/abs/2505.10543>.
812
- 813 Tianbao Xie, Fan Zhou, Zhoujun Cheng, Peng Shi, Luoxuan Weng, Yitao Liu, Toh Jing Hua, Jun-
814 ning Zhao, Qian Liu, Che Liu, Zeyu Liu, Yiheng Xu, Hongjin SU, Dongchan Shin, Caiming
815 Xiong, and Tao Yu. Openagents: An open platform for language agents in the wild. In *First*
816 *Conference on Language Modeling*, 2024. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=sKATR201Y0)
817 [sKATR201Y0](https://openreview.net/forum?id=sKATR201Y0).
818
- 819 Xiaoyu Xu, Xiang Yue, Yang Liu, Qingqing Ye, Haibo Hu, and Minxin Du. Unlearning isn't
820 deletion: Investigating reversibility of machine unlearning in llms, 2025a. URL <https://arxiv.org/abs/2505.16831>.
821
- 822 Yi Xu, Bo Xue, Shuqian Sheng, Cheng Deng, Jiabin Ding, Zanwei Shen, Luoyi Fu, Xinbing
823 Wang, and Chenghu Zhou. Good idea or not, representation of llm could tell. *ArXiv preprint*,
824 [abs/2409.13712](https://arxiv.org/abs/2409.13712), 2024. URL <https://arxiv.org/abs/2409.13712>.
825
- 826 Zhangchen Xu, Yuetai Li, Fengqing Jiang, Bhaskar Ramasubramanian, Luyao Niu, Bill Yuchen
827 Lin, and Radha Poovendran. Tinyv: Reducing false negatives in verification improves rl for llm
828 reasoning, 2025b. URL <https://arxiv.org/abs/2505.14625>.
- 829 Junxiao Yang, Jinzhe Tu, Haoran Liu, Xiaoce Wang, Chujie Zheng, Zhexin Zhang, Shiyao Cui,
830 Caishun Chen, Tiantian He, Hongning Wang, Yew-Soon Ong, and Minlie Huang. Barrel:
831 Boundary-aware reasoning for factual and reliable llms, 2025. URL [https://arxiv.org/](https://arxiv.org/abs/2505.13529)
832 [abs/2505.13529](https://arxiv.org/abs/2505.13529).
- 833 Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-
834 of-thought reasoning in llms. In *ICML 2025*, 2025. URL [https://arxiv.org/abs/2502.](https://arxiv.org/abs/2502.03373)
835 [03373](https://arxiv.org/abs/2502.03373).
836
- 837 Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhen-
838 guo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions
839 for large language models. In *The Twelfth International Conference on Learning Representations*,
840 2024. URL <https://openreview.net/forum?id=N8N0hgNDRt>.
- 841 Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhao Chen.
842 MAMmoTH: Building math generalist models through hybrid instruction tuning. In *The Twelfth*
843 *International Conference on Learning Representations*, 2024a. URL [https://openreview.](https://openreview.net/forum?id=yLClGs770I)
844 [net/forum?id=yLClGs770I](https://openreview.net/forum?id=yLClGs770I).
- 845 Xiang Yue, Tianyu Zheng, Ge Zhang, and Wenhao Chen. MAMmoTH2: Scaling instructions from
846 the web. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*,
847 2024b. URL <https://openreview.net/forum?id=yVu5dnPlqA>.
848
- 849 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-
850 zoo: Investigating and taming zero reinforcement learning for open base models in the wild, 2025.
851 URL <https://arxiv.org/abs/2503.18892>.
- 852 Yuxiang Zhang, Yuqi Yang, Jiangming Shu, Yuhang Wang, Jinlin Xiao, and Jitao Sang. Openrft:
853 Adapting reasoning foundation model for domain-specific tasks with reinforcement fine-tuning.
854 *ArXiv preprint*, [abs/2412.16849](https://arxiv.org/abs/2412.16849), 2024. URL <https://arxiv.org/abs/2412.16849>.
855
- 856 Junhao Zheng, Xidi Cai, Shengjie Qiu, and Qianli Ma. Spurious forgetting in continual learning
857 of language models. In *The Thirteenth International Conference on Learning Representations*,
858 2025. URL <https://openreview.net/forum?id=ScI7I1kGdI>.
- 859 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyang Luo. LlamaFactory: Uni-
860 fied efficient fine-tuning of 100+ language models. In Yixin Cao, Yang Feng, and Deyi Xiong
861 (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*
862 *(Volume 3: System Demonstrations)*, pp. 400–410, Bangkok, Thailand, August 2024. As-
863 sociation for Computational Linguistics. doi: 10.18653/v1/2024.acl-demos.38. URL <https://aclanthology.org/2024.acl-demos.38/>.

864 Ben Zhou, Daniel Khashabi, Qiang Ning, and Dan Roth. “going on a vacation” takes longer than
865 “going for a walk”: A study of temporal commonsense understanding. In *Proceedings of the 2019*
866 *Conference on Empirical Methods in Natural Language Processing and the 9th International*
867 *Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3363–3369, Hong
868 Kong, China, 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1332.
869 URL <https://aclanthology.org/D19-1332>.

870 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny
871 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint*
872 *arXiv:2311.07911*, 2023. URL <https://arxiv.org/abs/2311.07911>.

873
874 Lexin Zhou, Lorenzo Pacchiardi, Fernando Martínez-Plumed, Katherine M Collins, Yael Moros-
875 Daval, Seraphina Zhang, Qinlin Zhao, Yitian Huang, Luning Sun, Jonathan E Prunty, et al.
876 General scales unlock ai evaluation with explanatory and predictive power. *ArXiv preprint*,
877 *abs/2503.06378*, 2025. URL <https://arxiv.org/abs/2503.06378>.

878 Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul
879 Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *ArXiv*
880 *preprint*, *abs/1909.08593*, 2019. URL <https://arxiv.org/abs/1909.08593>.
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

918 A APPENDIX
 919
 920 CONTENTS
 921
 922 A.1 The Use of Large Language Models 19
 923 A.2 Ablations on Learning Rate, Regularization and Model Family 19
 924 A.3 Ablations on the Design of Transferability Index 20
 925 A.4 Additional Related Work 20
 926 A.5 Detailed Evidence of Our Observed Phenomena 21
 927 A.6 Full Evaluation Setup 21
 928 A.6.1 Post-training Methods 21
 929 A.6.2 Training Datasets 22
 930 A.6.3 Baselines 23
 931 A.6.4 Evaluation Benchmarks 23
 932 A.6.5 Evaluation metrics 24
 933 A.7 PCA Analysis under Varying Settings 25
 934 A.8 Details on Ablation Study 29
 935 A.8.1 Gradient Norm of SFT and RL 29
 936 A.8.2 Response Length of online RL and online SFT 29
 937 A.9 Token Distribution Shifts 30
 938 A.9.1 Case study and details 30
 939 A.9.2 Details about the result analysis 34
 940 A.10 Breakdown Evaluation Results 36
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971

972 A.1 THE USE OF LARGE LANGUAGE MODELS

973
974 We employ large language models for language editing, which is limited to refining text to enhance
975 readability. No language models contributed to the development of research ideas, analysis, models,
976 or interpretation of results. We also used LLMs to validate part of our code to cover the edge cases.
977

978 A.2 ABLATIONS ON LEARNING RATE, REGULARIZATION AND MODEL FAMILY

979 We conduct additional ablations to test whether our conclusions are sensitive to (i) the underlying
980 model family and scale, (ii) the choice of lower learning rate for SFT, and (iii) the use of standard
981 regularization such as L2-SP. We also expand evaluation to a broader set of additional benchmarks.
982

983 **Model family and learning rate.** Table 5 reports results on two additional base models, Qwen3-
984 4B and Llama3.1-8B, under SFT and RL with both standard and reduced learning rates. Across both
985 architectures, we observe the same qualitative pattern as in our main Qwen3-14B study: math-only
986 SFT substantially improves math performance but often harms non-reasoning capabilities, whereas
987 RL improves math and preserves (or slightly improves) non-reasoning tasks. These results indi-
988 cate that our main conclusions are *not* specific to the Qwen3-14B family, but hold across distinct
989 architectures and scales. A natural concern is that the larger representational drift under SFT might
990 simply be caused by a higher learning rate. To probe this, we trained with a much smaller learning
991 rate (1×10^{-6}) on both Qwen3-4B and Llama3.1-8B. As shown in Table 5, reducing the learning
992 rate uniformly degrades performance: for Qwen3-4B, *Math Avg.* drops from 42.6 to 34.7 and *Other*
993 *Reasoning Avg.* from 31.3 to 22.3, while *Non-Reasoning Avg.* decreases slightly further from 25.3
994 to 22.9, remaining far below the base model. Similarly, for Llama3.1-8B, the low-learning-rate SFT
995 model underperforms the standard SFT model on all three task families and still exhibits strong neg-
996 ative transfer on non-reasoning tasks. Thus, decreasing the SFT learning rate does not improve the
997 performance of SFT; instead it primarily reduces math gains and weakens other reasoning perfor-
998 mance. Combined with our on-policy vs. off-policy comparisons (which use the same SFT learning
999 rate), this supports our claim that sampling distribution and credit assignment, rather than learning
1000 rate alone, are the main drivers of cross-domain transfer behavior.

1001 **Regularization with L2-SP.** Table 6 evaluates Qwen3-8B under standard SFT and SFT with L2-
1002 SP regularization, in both off-policy and on-policy settings. The results show that L2-SP produces
1003 only minor changes in both raw averages and transferability indices. In the off-policy setting, non-
1004 reasoning performance is 26.6 for vanilla SFT and 26.4 for SFT+L2-SP, with $TI_{\text{non}} = -40.5$ vs.
1005 -41.2 , indicating nearly the same degree of negative transfer. In the on-policy setting, SFT+L2-SP
1006 achieves slightly higher math and other reasoning averages than vanilla on-policy SFT, but *Non-*
1007 *Reasoning Avg.* remains very similar, and TI_{non} changes only modestly. Consistent with our mecha-
1008 nistic analysis, what matters more for transfer is *where* and *how* updates are concentrated (on-policy
1009 sampling, advantage-like credit assignment, and negative examples), rather than merely changing
1010 their magnitude.

1011 **Broader evaluation benchmarks.** Table 7 extends our evaluation to five additional benchmarks:
1012 PIQA (Bisk et al., 2020), CommonsenseQA (Talmor et al., 2019), RACE (Lai et al., 2017), and
1013 LogiQA (Liu et al., 2020), OpenBookQA (Mihaylov et al., 2018). These tasks are not used during
1014 math-only post-training and therefore test whether math specialization disrupts or preserves general
1015 capabilities. On all five benchmarks, the RL-tuned model UniReason-Qwen3-14B outperforms both
1016 the base model and the math-SFT variants. In contrast, the math-SFT models often hurt perfor-
1017 mance on non-reasoning benchmarks despite strong math improvements. These additional bench-
1018 marks corroborate our main finding: math-only RL can improve math reasoning without sacrificing,
1019 and in some cases improving, non-reasoning and commonsense skills, whereas standard math-SFT
1020 frequently induces negative transfer to such tasks.

1021 **Broader Training Data.** To analyze the impact of broader training data, we additionally consider
1022 the General-Reasoner-14B model (Ma et al., 2025) as a case study. In contrast to our math-only
1023 setting (UniReason-Qwen3-14B), which is fine-tuned on $\sim 47k$ purely mathematical problems,
1024 General-Reasoner is trained on a WebInstruct-filtered corpus containing $\sim 230k$ diverse reason-
1025 ing questions beyond math (e.g., physics, chemistry, business, etc.). To disentangle the effect of

Table 5: Performance of additional base models under SFT and RL. (1e-6) denotes the model is trained with a learning rate of 1×10^{-6} . Other models are trained with default settings mentioned in the paper.

Model	Math Avg.	Other Reasoning Avg.	Non-Reasoning Avg.	TI _{other}	TI _{non}
Qwen3-4B-Base	12.2	16.1	39.6		
Qwen3-4B-RL	40.5	32.6	42.6	34.8	22.3
Qwen3-4B-SFT	42.6	31.3	25.3	32.7	-40.6
Qwen3-4B-SFT (1e-6)	34.7	22.3	22.9	25.9	-44.3
Llama3.1-8B-Base	3.8	8.9	24.7		
Llama3.1-8B-RL	6.9	10.1	25.0	71.8	7.7
Llama3.1-8B-SFT	7.8	11.2	17.3	70.4	-34.7
Llama3.1-8B-SFT (1e-6)	4.6	9.7	14.7	56.3	-38.9

Table 6: Performance of Qwen3-8B variants with and without L2-SP regularization. TI_{other} and TI_{non} denote transferability indices for other reasoning and non-reasoning tasks.

Model	Math Avg.	Other Reasoning Avg.	Non-Reasoning Avg.	TI _{other}	TI _{non}
Qwen3-8B-Base	27.6	23.6	33.6		
Qwen3-8B-Off-policy-SFT	41.9	34.4	26.6	18.3	-40.5
Qwen3-8B-Off-policy-SFT-L2SP	42.8	35.1	26.4	16.9	-41.2
Qwen3-8B-On-policy-SFT	33.7	35.7	35.0	68.6	30.2
Qwen3-8B-On-policy-SFT-L2SP	35.1	36.2	34.7	66.3	27.8

the training data from the optimization paradigm, we reproduce a SFT-finetuned model, **General-Reasoner-14B-SFT**, by training with our SFT pipeline on the authors’ released WebInstruct-based dataset, and compare it to our model under the same evaluation protocol. As shown in Table 10, this result shows that our conclusion is consistent on not only math-only data. And can be observed on different training data sources.

A.3 ABLATIONS ON THE DESIGN OF TRANSFERABILITY INDEX

We analyze the effect of our Transferability Index (TI) design and the signed square-root transformation. Recall that TI_{other} and TI_{non} quantify how much math gains transfer to other reasoning and non-reasoning task families, after normalizing for domain difficulty. We apply a signed square-root transform to compress very large positive/negative values while preserving the sign. Tables 8 and 9 report both the transformed values and the raw TI, together with their ranks. The key observation is that the transform does not change our qualitative conclusions: for TI_{non}, the ranking of models within RL and SFT groups is identical under both scales, and all SFT models still show negative transfer. For TI_{other}, the top RL models and the top SFT models remain the same, with only minor swaps among mid-ranked models that do not affect any claims in the paper. In all cases, the clear separation between RL models (mostly large positive TI) and many math-SFT models (often modest or negative TI) is preserved.

A.4 ADDITIONAL RELATED WORK

Reasoning Fine-Tuning of LLMs. Recent advancements in large language models have notably emphasized specialized fine-tuning methods to enhance reasoning capabilities (Wong et al., 2025; Chen et al., 2023; Ziegler et al., 2019; Liu et al., 2025; Wang et al., 2024c; Li et al., 2025b; Feng et al., 2025; Xu et al., 2025b; Yang et al., 2025; Li et al., 2025a; Yeo et al., 2025). The chain-of-thought prompting strategy introduced by Wei et al. (2022) encourages models to produce step-by-step explanations, significantly boosting performance in symbolic reasoning tasks (Lambert et al., 2025; Wei et al., 2022; Longpre et al., 2023; Yu et al., 2024). Subsequent extensions, such as DeepSeek-R1 (Team, 2025a), have integrated reinforcement learning approaches alongside CoT, optimizing models through reward-driven policy improvements. Such RL-enhanced fine-tuning has achieved state-of-the-art results on benchmarks and competitive programming challenges (Hendrycks et al., 2021b; Team et al., 2025; Team, 2025a; Lambert et al., 2025).

Table 7: Performances of models in our control study on additional non-reasoning benchmarks.

Model	PIQA	LogiQA	RACE	CommonsenseQA	OpenBookQA
Qwen3-14B Base	79.5	32.7	40.2	50.9	32.0
Qwen3-14B think	78.5	28.4	38.0	20.6	25.6
Qwen3-14B no-think	74.3	26.7	39.1	19.2	28.4
UniReason-Qwen3-14B (RL)	82.4	35.8	43.3	55.2	35.7

Table 8: Transferability Index for other reasoning tasks. We report both the signed square-root transformed value and the raw TI_{other} , together with their ranks within RL or SFT models.

Alias	Category	TI_{other} (sqrt)	Rank (sqrt)	TI_{other}	Rank
General-Reasoner-RL-14B	RL	125.1	1	214.5	1
SimpleRL-14B	RL	98.6	2	188.0	2
SimpleRL-1.5B	RL	88.1	3	67.2	4
UniReason-RL-14B	RL	84.7	4	82.4	3
DAPO-32B	RL	62.3	5	55.2	5
SimpleRL-Llama3.1-8B	RL	57.0	6	23.9	6
SimpleRL-7B	RL	15.8	7	15.3	7
Deepseek-R1-Distill-Qwen-7B	SFT	62.1	1	59.0	2
UniReason-SFT-14B-think	SFT	53.7	2	75.4	1
OpenThinker3-7B	SFT	45.7	3	49.3	3
LIMO-32B	SFT	41.6	4	33.0	4
OpenThinker2-32B	SFT	40.2	5	31.8	5
S1.1-32B	SFT	34.7	6	26.7	6
Mistral-Reason-24B	SFT	23.3	7	11.7	7
S1.1-7B	SFT	-67.3	8	-87.4	8

A.5 DETAILED EVIDENCE OF OUR OBSERVED PHENOMENA

As discussed in Section 2.1, we provided the complete evaluation for the Transferability Index for the off-the-shelf models on other reasoning and non-reasoning tasks in Table 11.

A.6 FULL EVALUATION SETUP

A.6.1 POST-TRAINING METHODS

Reinforcement Learning has recently proven effective at steering large language models toward complex, multi-step objectives by optimizing policies with scalar reward signals (Zeng et al., 2025). We used the Verl framework (Sheng et al., 2025) and GRPO (Shao et al., 2024) setup to fine-tune the Qwen-3-14B-Base model, utilizing answer correctness as the reward signal. Our RL uses a learning rate of 1×10^{-6} with an overall train batch size of 512 and clipping thresholds set between 0.22 and 0.28. We generate sequences up to 16k tokens long and perform 16 rollouts per prompt, then update the model in mini-batches of 128 samples. Both KL-divergence and entropy penalties are turned off (coefficients set to zero). We train the model for 140 steps and used the corresponding checkpoint.

Supervised Fine-Tuning remains a fundamental technique for adapting large pre-trained models by directly minimizing cross-entropy on high-quality datasets (Chu et al., 2025). We use the LLaMA-Factory framework (Zheng et al., 2024), which is an extensible and user-friendly framework supporting multiple architectures and advanced optimization algorithms, to fine-tune our model on teacher-generated chain-of-thought traces. We use 5×10^{-5} as learning rate, the batch size is 512 and we train for 1.5 epoch to align with our RL settings.

Table 9: Transferability Index for non-reasoning tasks. We report both the signed square-root transformed value and the raw TI_{non} , together with their ranks within RL or SFT models.

Alias	Category	TI_{non} (sqrt)	Rank (sqrt)	TI_{non}	Rank
SimpleRL-14B	RL	103.8	1	183.5	1
General-Reasoner-RL-14B	RL	68.0	2	59.1	2
UniReason-RL-14B	RL	53.4	3	30.5	3
SimpleRL-7B	RL	35.2	4	28.5	4
SimpleRL-1.5B	RL	9.8	5	1.5	6
SimpleRL-Llama3.1-8B	RL	5.4	6	3.3	5
DAPO-32B	RL	-40.4	7	-20.4	7
Deepseek-R1-Distill-Qwen-7B	SFT	-4.2	1	-0.4	1
LIMO-32B	SFT	-27.9	2	-7.9	2
OpenThinker2-32B	SFT	-54.7	3	-32.6	4
S1.1-32B	SFT	-63.6	4	-41.2	5
UniReason-SFT-14B-think	SFT	-64.7	5	-31.8	3
OpenThinker3-7B	SFT	-71.6	6	-59.2	6
Mistral-Reason-24B	SFT	-91.5	7	-86.2	7
S1.1-7B	SFT	-218.2	8	-326.8	8

Table 10: Additional Study on broader training data show that our finding can also be used outside of math-only data. A.2.

Math Reasoning						
Model	AIME24	AIME25	MATH500	Olympiad	Average	
Qwen3-14B-Base	13.0	9.3	60.4	27.9	27.7	-
General-Reasoner-14B-SFT	35.0	22.6	78.4	22.5	39.6	-
General-Reasoner-14B-RL	24.4	19.2	83.0	32.3	39.7	-
Δ (RL - SFT)	-10.6	-3.4	+4.6	+9.8	+0.1	-
Other Reasoning						
Model	GPQA	LiveCodeBench2	ACPBench	HeadQA	Average	TI_{other}
Qwen3-14B-Base	42.6	29.7	10.7	37.6	30.2	-
General-Reasoner-14B-SFT	43.9	19.7	64.3	39.9	41.1	+52.2
General-Reasoner-14B-RL	56.1	32.9	75.0	40.0	51.0	+121.6
Δ (RL - SFT)	+12.2	+13.2	+10.7	+0.1	+9.1	-
Non-Reasoning						
Model	CoQA	MC-TACO	IFEval	HalluEval	Average	TI_{non}
Qwen3-14B-Base	10.0	67.7	69.2	35.7	45.7	-
General-Reasoner-14B-SFT	8.7	38.2	31.7	10.8	22.3	-110.7
General-Reasoner-14B-RL	11.2	74.0	72.0	55.7	53.2	+66.5
Δ (RL - SFT)	+2.5	+35.8	+40.3	+44.9	+30.9	-

A.6.2 TRAINING DATASETS

As briefed in Section 2.2, our base training dataset is a curated set of 47K high-quality mathematics problems. We stratified the examples using two complementary sources: low-difficulty problems drawn from the DeepScaler dataset (Luo et al., 2025), and high-difficulty (levels 3–5) problems extracted from SimpleRL (Zeng et al., 2025). To generate CoT annotations, we prompt each problem into the Qwen3-32B-Instruct model (Team, 2025b) and use reject sampling to generate our dataset.

To further explore the effect of training data distribution for SFT-based reasoning models, we also distill a larger and more comprehensive dataset collected from General-Reasoner (Ma et al., 2025), which contains 232K examples across reasoning and non-reasoning tasks (e.g., Math, Chemistry, Business). This additional distilled set is used to train the General-Reasoner model using supervised fine-tuning.

1188 A.6.3 BASELINES

1189 In our experiments, we compare against Qwen3-14B-Base model (Team, 2025b), which is the original Qwen3-14B model without any further adaptation. This serves as the unmodified backbone for all fine-tuning models. Also, we report the results of Qwen3-14B-Instruct model under our tested benchmarks. It is an instruction-tuned version of Qwen3-14B model trained on a large, general-purpose instruction-following dataset. We evaluate it under two prompting modes:

- 1195 • *think*: prompts include a special `<think>` token to explicitly elicit chain-of-thought reasoning.
- 1196
- 1197 • *no-think*: prompts are provided without the `<think>` token, relying solely on direct instruction-following behavior.
- 1198
- 1199

1200 Due to the wide range and enormous training data, this model is considered to give the optimal outputs across tasks in the current 14B-series models.

1201

1202 To further validate our observation, we apply our controlled study pipeline also for General-Reasoner (Ma et al., 2025), a reasoning model tuned by cold RL with training data containing not only math but more diverse domains. We distilled the dataset for SFT finetuning using the same rejection sampling method with their proposed dataset. The dataset contains 232K samples covering various reasoning and non-reasoning tasks. Then we finetune the Qwen3-14B base model using the distilled dataset and name the model as General-Reasoner-Qwen3-14B(SFT) to directly compare with the RL-based General-Reasoner for a fairer and more comprehensive controlled study towards SFT and RL. The Transferability Index results could be retrieved in Table 11, the observation also confirms our initial hypothesis.

1211 A.6.4 EVALUATION BENCHMARKS

1212 In the experiment, we evaluated our model across a wide range of benchmarks. Notably, to explicitly reveal the transferability of reasoning models, we grouped them into three categories by their content:

1213

1214 **Math Reasoning Datasets** We collected the following datasets that are composed of mathematical problems, which means that they typically need a mathematical reasoning process to get the answer:

- 1217
- 1218 • **MATH500** (Hendrycks et al., 2021b): A curated subset of 500 problems sampled from the broader MATH dataset, covering topics like algebra, combinatorics, geometry, and number theory.
- 1219
- 1220
- 1221 • **AIME**: Problems drawn from the American Invitational Mathematics Examination (AIME) 2024 and 2025, each with 30 challenging short-answer questions requiring multi-step reasoning.
- 1222
- 1223 • **OlympiadBench** (He et al., 2024): Problems sourced from international olympiads (e.g., IMO and regional contests).
- 1224
- 1225
- 1226
- 1227
- 1228

1229 **Other Reasoning Datasets** We collected the following datasets that are mainly composed of general reasoning problems containing a wider range of subjects:

- 1230
- 1231
- 1232 • **LiveCodeBench** (Jain et al., 2025): It is a continuously updated, contamination-free coding benchmark. We used its second version.
- 1233
- 1234 • **GPQA-Diamond**(Rein et al., 2024): It is a graduate-level question-answering dataset that contains multiple-choice questions in biology, physics, and chemistry. We followed its diamond split.
- 1235
- 1236 • **ACPBench** (Kokel et al., 2025): It has 7 atomic reasoning tasks around 13 classical planning domains. We only used the multiple-choice problems.
- 1237
- 1238 • **HeadQA** (Vilares & Gómez-Rodríguez, 2019): Multiple-choice QA from healthcare-specialist certification exams, including questions across pharmacology, chemistry, nursing, psychology, biology, and medicine.
- 1239
- 1240
- 1241

Table 11: Comparison of model transferability indices on other reasoning and non-reasoning tasks across different training paradigms (SFT vs. RL), model sizes (7B vs. 32B), and model families (Qwen vs. Llama).

Model	Base Model	Size	Method	TI _{other}	TI _{non}
SFT vs RL					
Deepseek-R1-Distill-Qwen-7B	Qwen2.5-Math-7B-Base	7B	SFT	60.4	-4.2
OpenThinker2-7B	Qwen2.5-7B-Instruct	7B	SFT	52.0	-76.2
OpenThinker3-7B	Qwen2.5-7B-Instruct	7B	SFT	44.5	-70.0
S1.1-7B	Qwen2.5-7B-Instruct	7B	SFT	-65.5	-213.3
General-Reasoner-Qwen3-14B (SFT)	Qwen3-14B-Base	14B	SFT	77.8	-127.6
Mistral-Small-24B-Instruct-2501-Reasoning	Mistral-Small-24B-Instruct-2501	24B	SFT	22.6	-89.5
OpenThinker2-32B	Qwen2.5-32B-Instruct	32B	SFT	39.1	-53.5
S1.1-32B	Qwen2.5-32B-Instruct	32B	SFT	33.8	-62.2
LIMO-32B	Qwen2.5-32B-Instruct	32B	SFT	40.5	-27.3
Qwen2.5-1.5B-SimpleRL	Qwen2.5-1.5B-Base	1.5B	RL	85.6	9.5
SimpleRL-7B	Qwen2.5-7B-Base	7B	RL	15.4	34.4
Qwen2.5-14B-SimpleRL	Qwen2.5-14B-Base	14B	RL	95.9	101.5
General-Reasoner-Qwen3-14B (RL)	Qwen3-14B-Base	14B	RL	121.6	66.5
DAPO-Qwen-32B	Qwen2.5-32B	32B	RL	60.7	-39.5
7B vs 32B					
S1.1-7B	Qwen2.5-7B-Instruct	7B	SFT	-65.5	-213.3
S1.1-32B	Qwen2.5-32B-Instruct	32B	SFT	33.8	-62.2
OpenThinker2-7B	Qwen2.5-7B-Instruct	7B	SFT	52.0	-76.2
OpenThinker2-32B	Qwen2.5-32B-Instruct	32B	SFT	39.1	-53.5
Qwen vs Llama					
Qwen2.5-7B-SimpleRL	Qwen2.5-7B-Base	7B	RL	15.4	34.4
Llama3.1-8B-SimpleRL	Llama3.1-8B	8B	RL	72.6	7.0

Non-reasoning Datasets We collected the following datasets that are mainly composed of problems with factual answers, which means that they do not need a reasoning process to give the answer:

- **CoQA** (Reddy et al., 2019): It has 127K questions in dialogues over passages, focusing on maintaining context and coreference across turns.
- **IFEval** (Zhou et al., 2023): It contains over 500 prompts, each embedding verifiable instructions. Evaluates strict vs. loose adherence to instructions.
- **HaluEval** (Li et al., 2023): It contains human-annotated samples where models must distinguish factual content from hallucinations.
- **MC-TACO** (Zhou et al., 2019): It is a multiple-choice benchmark designed to evaluate models’ temporal commonsense, covering duration, ordering, typical time, frequency, and stationarity.

A.6.5 EVALUATION METRICS

We used LLM-Harness (Gao et al., 2024) to evaluate the models’ performance on Olympiad-Bench, ACPBench, HeadQA, CoQA, HaluEval, MC-TACO and used Eval-Chemy (Raouf et al., 2025) MATH500, AIME24, AIME25, GPQA-Diamond, LiveCodeBench, IFEval. On MATH500, AIME24, AIME25, GPQA-Diamond, and LiveCodeBench, we used 0.6 as temperature, and 0.95 as top-p value. In our experiments, we used **accuracy** to evaluate the models’ performance. Specifically, for AIME24 and AIME 25, we averaged accuracy on 10 samples. For GPQA-Diamond, LiveCodeBench and MATH 500, our score is the average accuracy over 3 samples. Specifically, we used version 2 and overall accuracy on LiveCodeBench. For ACPBench, we only used multiple choices, and averaged the score for all 10 tasks as the final score. For OlympiadBench, we only used math queries in English, and thus categorized Olympiad as a math benchmark. For HaluEval, the performance is the accuracy averaged on 3 tasks with zero-shot. And for IFEval, we used strict instruction accuracy as the score. For OlympiadBench, ACPBench, HeadQA, CoQA, HaluEval, and IFEval, we used greedy sampling and sampled only once.

Table 12: Comparison of mean PCA shift magnitudes on math, reasoning, and non-reasoning tasks across different training paradigms (SFT vs. RL), model sizes (7B vs. 32B), and model families (Qwen vs. Llama). **Red-highlighted rows** indicate models trained with RL, non-Qwen baselines, or models with larger size.

Model	Base Model	Size	Method	Math	Other	Non
SFT vs RL						
Deepseek-R1-Distill-Qwen-7B	Qwen2.5-Math-7B	7B	SFT	38.8	27.4	42.2
OpenThinker2-7B	Qwen2.5-7B-Instruct	7B	SFT	388.8	444.0	5486.2
OpenThinker3-7B	Qwen2.5-7B-Instruct	7B	SFT	409.9	484.2	5228.2
S1.1-7B	Qwen2.5-7B-Instruct	7B	SFT	9.8	15.5	205.7
Qwen3-14B	Qwen3-14B-Base	14B	SFT	40.4	13.9	129.8
General-Reasoner-Qwen3-14B (SFT)	Qwen3-14B-Base	14B	SFT	7.9	10.5	118.2
OpenThinker2-32B	Qwen2.5-32B-Instruct	32B	SFT	25.4	72.0	726.8
S1.1-32B	Qwen2.5-32B-Instruct	32B	SFT	2.4	1.8	7.5
LIMO-32B	Qwen2.5-32B-Instruct	32B	SFT	15.7	9.0	7.5
Qwen2.5-1.5B-SimpleRL-Zoo	Qwen2.5-1.5B	1.5B	RL	0.5	0.3	0.6
Qwen2.5-7B-SimpleRL-Zoo	Qwen2.5-7B	7B	RL	0.6	0.2	0.6
Llama-3.1-8B-SimpleRL-Zoo	Llama-3.1-8B	8B	RL	0.3	0.1	0.0
Qwen2.5-14B-SimpleRL-Zoo	Qwen2.5-14B	14B	RL	0.9	1.0	2.7
General-Reasoner-Qwen3-14B(RL)	Qwen3-14B-Base	14B	RL	2.9	4.6	64.8
DAPO-Qwen-32B	Qwen2.5-32B	32B	RL	2.6	1.1	1.1
7B vs 32B						
S1.1-7B	Qwen2.5-7B-Instruct	7B	SFT	9.8	15.5	205.7
S1.1-32B	Qwen2.5-32B-Instruct	32B	SFT	2.4	1.8	7.5
OpenThinker2-7B	Qwen2.5-7B-Instruct	7B	SFT	388.8	444.0	5486.2
OpenThinker2-32B	Qwen2.5-32B-Instruct	32B	SFT	25.4	72.0	726.8
Qwen vs Llama						
Qwen2.5-7B-SimpleRL-Zoo	Qwen2.5-7B	7B	RL	0.6	0.2	0.6
Qwen3-14B-Base-SimpleRL-Zoo	Qwen3-14B-Base	14B	RL	0.9	1.0	2.7
Llama-3.1-8B-SimpleRL-Zoo	Llama-3.1-8B	8B	RL	0.3	0.1	0.0

A.7 PCA ANALYSIS UNDER VARYING SETTINGS

Table 12 summarizes the mean $d^{(*)}$ across math, other-reasoning, and non-reasoning tasks, providing an overall assessment of latent-space shifts under different training paradigms. Figure 7, 8, 9, 10

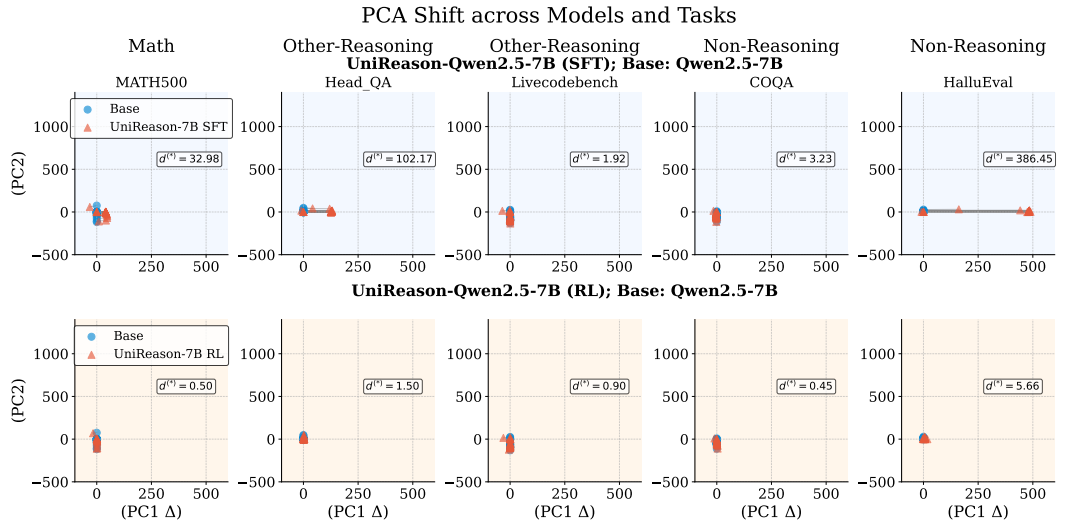


Figure 7: PCA shift of Qwen2.5-7B across different training methods and tasks. $d^{(*)}$ is the Euclidean distance between representation centroids before and after training. The first row shows models trained with SFT, and the last row shows models trained with RL. RL training results in the smallest PCA shift for all task types, suggesting more stable latent representations.

illustrates the paradigm comparison between SFT and RL, Figure 11 contrasts model sizes (7B versus 32B), and Figure 12 compares model families (Qwen vs. Llama).

Impact of Training Paradigm. RL-based fine-tuning consistently results in lower PCA shifts than SFT across math, other-reasoning, and non-reasoning tasks. As shown in Table 12, models such as SimpleRL-7B and SimpleRL-14B exhibit significantly smaller feature shifts compared to their SFT-

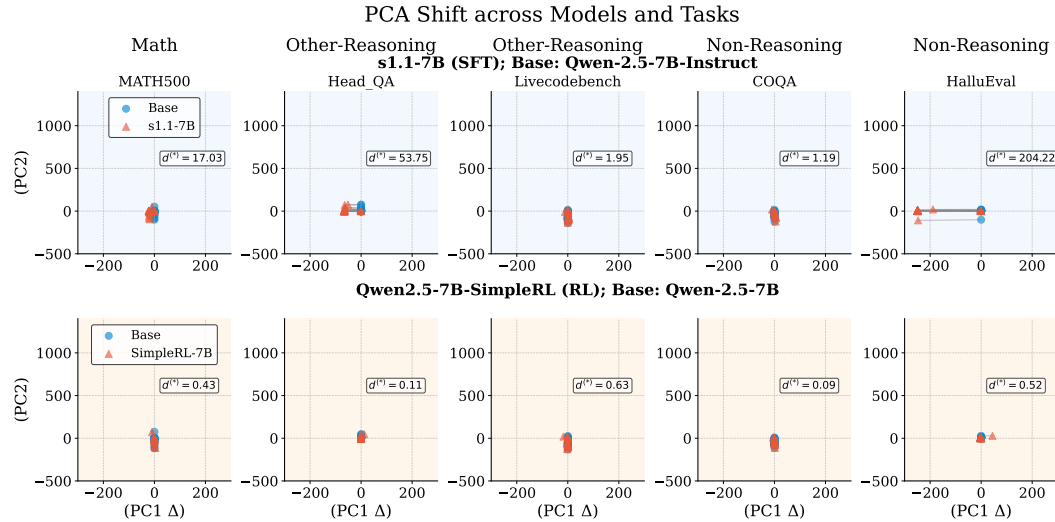


Figure 8: PCA shift for Qwen2.5-7B-Instruct and Qwen2.5-7B under different training paradigms and benchmarks. The first row corresponds to models trained with SFT, while the second row corresponds to the RL. The RL-based training paradigm achieves the lowest PCA shift across all task categories, indicating enhanced stability in the model’s latent representation space for math, other reasoning, and non-reasoning tasks. $d^{(*)}$ denotes the Euclidean distance between the centroids of representations before and after training.

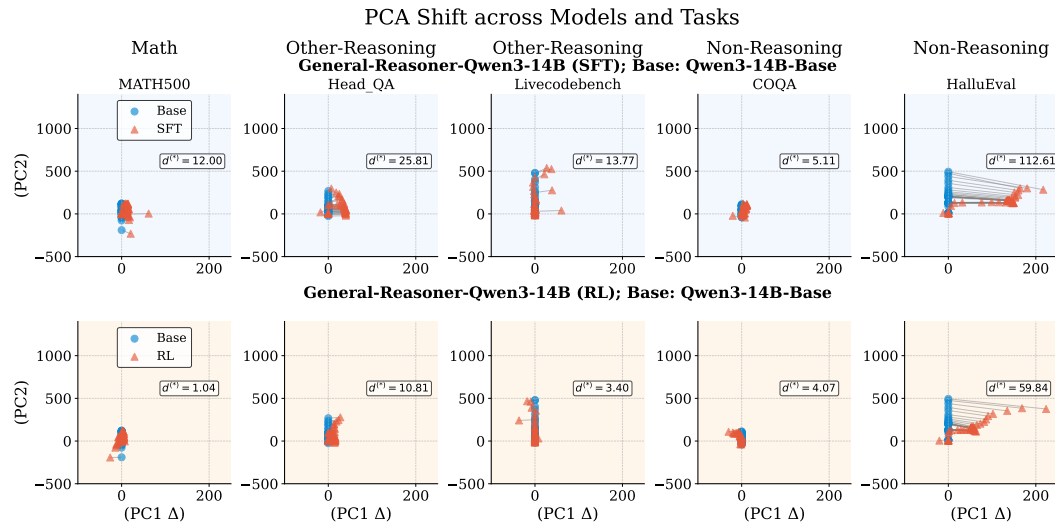


Figure 9: PCA shift for Qwen3-14B-Base under different training paradigms and benchmarks. The first row corresponds to models trained with SFT, while the second row corresponds to the RL. The RL-based training paradigm achieves the lowest PCA shift across all task categories, indicating enhanced stability in the model’s latent representation space for math, other reasoning, and non-reasoning tasks. $d^{(*)}$ denotes the Euclidean distance between the centroids of representations before and after training.

trained counterparts. This observation is further visualized in Figure 7, 8, 9, 10, which demonstrates more concentrated representation shifts under RL. These findings are consistent with the phenomena discussed in Section 2.1, reinforcing that RL enhances generalization by better preserving internal representations. Overall, these results suggest that RL is substantially more effective at preserving general-domain representations and mitigating catastrophic forgetting.

Effect of Model Size. Increasing model size from 7B to 32B substantially reduces PCA shifts on math, other-reasoning, and non-reasoning tasks, as shown in Table 12. Figure 11 further confirms that larger models exhibit more concentrated and stable latent representations in non-reasoning do-

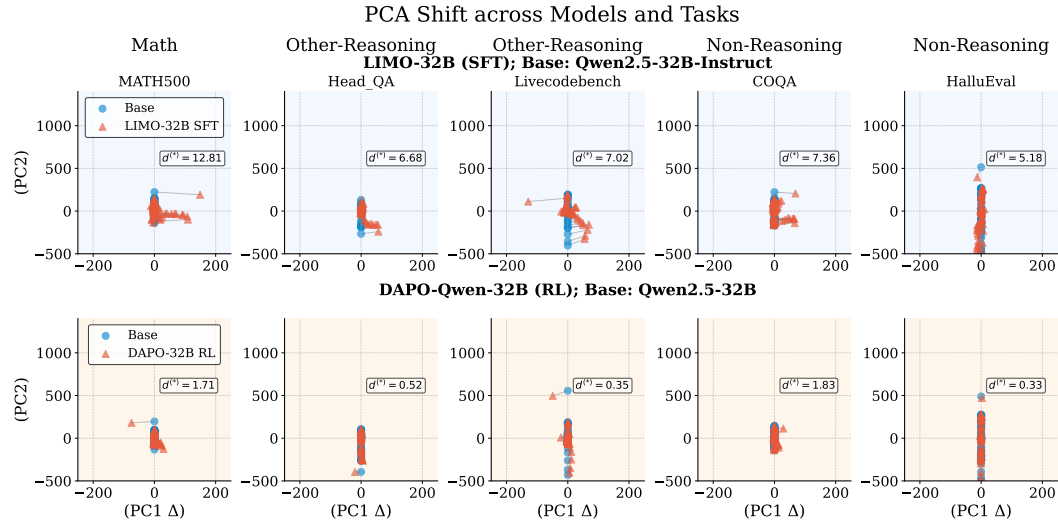


Figure 10: PCA shift for Qwen2.5-32B-Instruct and Qwen2.5-32B under different training paradigms and benchmarks. The first row corresponds to models trained with SFT, while the second row corresponds to the RL. The RL-based training paradigm achieves the lowest PCA shift across all task categories, indicating enhanced stability in the model’s latent representation space for math, other reasoning, and non-reasoning tasks. $d^{(*)}$ denotes the Euclidean distance between the centroids of representations before and after training.

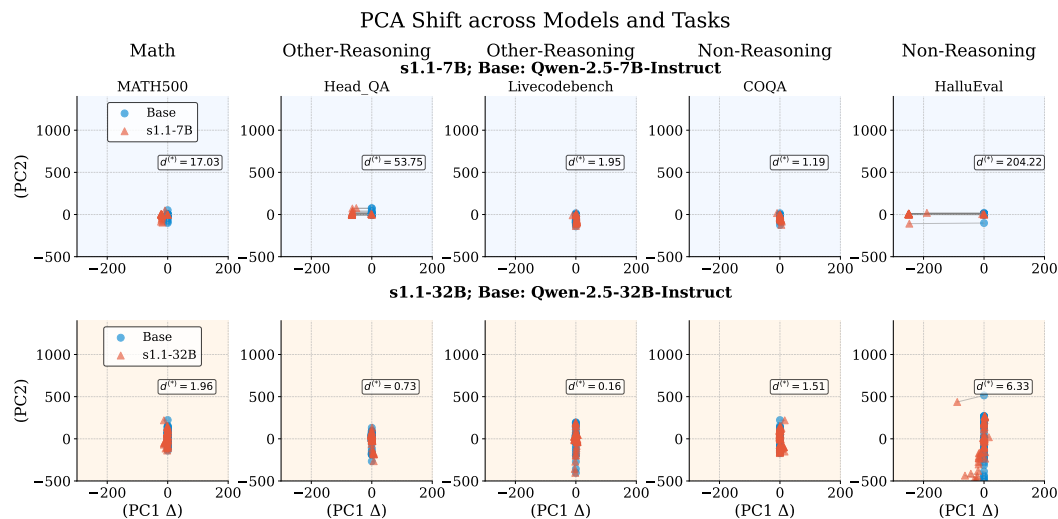


Figure 11: PCA shift for under different model sizes (7B vs 32B) and benchmarks. The first row corresponds to the 7B model, while the second row corresponds to the 32B model. $d^{(*)}$ denotes the Euclidean distance between the centroids of representations before and after training.

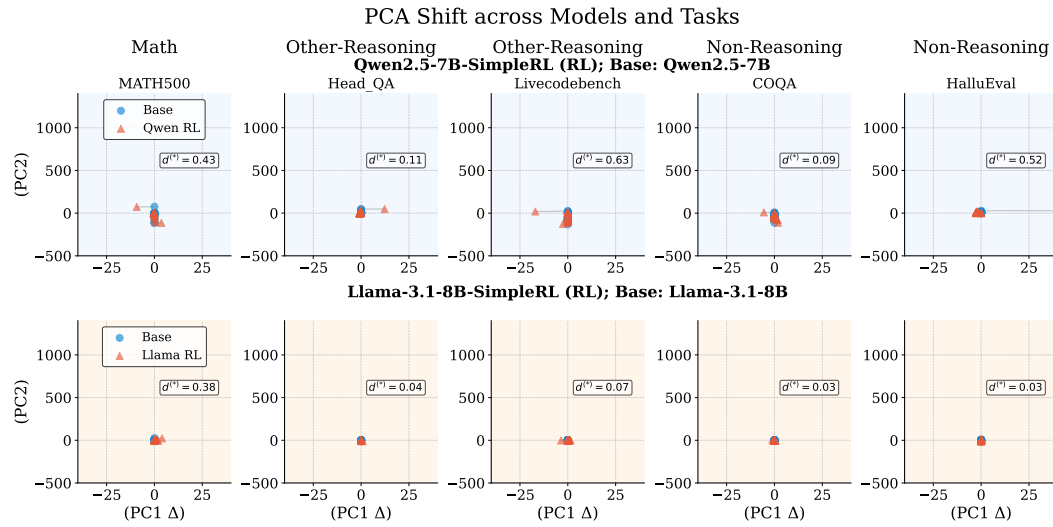


Figure 12: PCA shift under different model families (Qwen vs Llama) and benchmarks. The first row corresponds to the Qwen model, while the second row corresponds to the Llama model. $d^{(*)}$ denotes the Euclidean distance between the centroids of representations before and after training.

main, indicating reduced feature drift. These findings suggest that scaling up model size is particularly effective for enhancing stability and resistance to latent space drift in general-domain tasks.

Effect of Model Family. The results in Table 12 and Figure 12 indicate that both RL-tuned Llama and Qwen models exhibit small PCA shifts. While Qwen displays similar shifts as Llama on reasoning tasks, it also has higher drift on other-reasoning and non-reasoning tasks. In contrast, Llama models show greater resilience to drift in general domains. These differences stem from the distinct architectures and pre-training data, suggesting that RL strategies should be tailored to each model’s characteristics.

Summary. Our analyses show that latent space drift is mainly determined by training paradigm, model size, and model family. RL-based fine-tuning is the most effective at preserving general-domain representations and mitigating catastrophic forgetting. While larger models offer greater feature stability, especially on other-reasoning and non-reasoning tasks, scaling alone is insufficient. Qwen models remain sensitive to reasoning-related pre-training data, underscoring the need for tailored strategies. Overall, RL-based training is critical for developing robust language models that balance general-domain coverage with reasoning ability.

Model	Type	Early Slope	$t_{1/2}$	Total Grad.	Late Avg.	Late Var.	Decay Rate	Oscillation
14B-RL	RL	-0.005	22.0	48.190	0.448	0.041	-0.262	0.034
14B-SFT-no-think	SFT	-0.019	80.0	64.906	0.405	0.098	-0.592	0.368
14B-SFT-think	SFT	-0.030	80.0	64.187	0.383	0.112	-0.625	0.382

Table 13: **Gradient norm indicators in early training (first 100 steps)**. Smaller *Early Slope* values (closer to zero) show more gradual early updates. Lower *Total Grad.* suggests smaller overall movement in parameter space. Higher *Late Avg.* with lower *Late Var.* reflects a stable learning signal in later steps. A less negative *Decay Rate* shows slower reduction of gradient strength. Finally, smaller *Oscillation* values mean smoother training, with fewer rapid shifts. Across all metrics, RL shows more stable and generalizable training behavior than SFT.

A.8 DETAILS ON ABLATION STUDY

A.8.1 GRADIENT NORM OF SFT AND RL

As shown in Figure 13 and Table 13, RL and SFT exhibit distinct gradient dynamics. RL performs gentler early updates (Early Slope closer to zero) without the sharp spikes observed in SFT, and accumulates smaller overall movement in parameter space (lower Total Grad.). In later steps, RL maintains a stronger and more stable gradient signal (higher Late Avg., lower Late Var.), with a slower decay rate and much less oscillation. In contrast, SFT shows large initial fluctuations followed by rapid collapse toward near-zero gradients. Overall, RL demonstrates a more stable and persistent optimization process, whereas SFT tends to quickly saturate and lose effective gradient signals.

A.8.2 RESPONSE LENGTH OF ONLINE RL AND ONLINE SFT

Compared to SFT, online RL produces a marked increase in response length as training progresses. This pattern reflects RL’s stronger exploration ability: negative examples discourage uninformative outputs, while credit assignment amplifies signals from task-relevant tokens. Together, these mech-




Figure 13: Grad norm comparison across different fine-tuning strategies for the first 100 steps. RL exhibits consistently lower gradient norms compared to both SFT-no-think and SFT-think models, suggesting more stable optimization dynamics during fine-tuning.




Figure 14: Response length comparison across online RL and SFT fine-tuning (first 80 steps). Online RL gradually increases response length during training, while SFT remains stable. This indicates that RL promotes stronger exploration, supported by mechanisms such as negative examples and credit assignment that guide the model to generate more diverse and extended outputs. Longer responses reflect broader policy exploration, which is often linked to better generalization and transfer across tasks.

anisms encourage the model to explore a wider solution space, leading to longer and more varied responses. Such exploration is important for improving transferability, as it enables the model to adapt its reasoning strategies beyond the narrow supervision provided in SFT.

A.9 TOKEN DISTRIBUTION SHIFTS

A.9.1 CASE STUDY AND DETAILS

This part provides the complete queries and responses in our case study in Section 4. Table 14 enumerates all shifted tokens for both reasoning and non-reasoning queries. RL modifies a compact set of task-salient tokens, whereas SFT shifts a much larger set including many off-topic tokens.

Case Study:

1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

Reasoning Task

Query:

Solve the following math problem. Present the final answer in the format: Final Answer: boxed{your_answer}. Problem: Ten treeks weigh as much as three squigs and one goolee. Two treeks and one goolee are equal in weight to one squig. The combined weight of how many treeks equals the weight of one squig?

RL-tuned Models:

Let's define the variables for the weights of the treeks, squigs, and goolees. Let: T be the weight of one treek, S be the weight of one squig, G be the weight of one goolee.

We are given two pieces of information: 1. Ten treeks weigh as much as three squigs and one goolee: $10T = 3S + G$ 2. Two treeks and one goolee are equal in weight to one squig: $2T + G = S$ We need to find how many treeks combined equal the weight of one squig. In other words, we need to find a number x such that $xT = S$.

First, let's solve the second equation for G : $2T + G = S \Rightarrow G = S - 2T$

Now, substitute $G = S - 2T$ into the first equation: $10T = 3S + (S - 2T)$ $10T = 3S + S - 2T$ $10T = 4S - 2T$

Now, add $2T$ to both sides to isolate the terms involving S : $10T + 2T = 4S$ $12T = 4S$

Now, divide both sides by 4 to solve for S : $3T = S$

This tells us that the weight of one squig is equal to the weight of three treeks. Therefore, the combined weight of three treeks equals the weight of one squig.

Final Answer: 3

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727

SFT-tuned Models:

Okay, let's see. I need to solve this problem where they talk about treeks, squigs, and goolees. The question is asking how many treeks combined equal the weight of one squig.

First, let me parse the problem again. The first statement says ten treeks weigh as much as three squigs and one goolee. So, $10T = 3S + G$, where T is treek, S is squig, G is goolee. The second statement is two treeks and one goolee equal one squig. So, $2T + G = S$. The question is asking for how many treeks equal one squig. So, find x such that $xT = S$.

Alright, so I have two equations:

$$1) 10T = 3S + G$$

$$2) 2T + G = S$$

I need to solve these equations to find x where $xT = S$. Let me see. Maybe I can express G from the second equation and substitute into the first.

From equation 2: $G = S - 2T$. Let me check that. If $2T + G = S$, then subtract $2T$ from both sides: $G = S - 2T$. Yes, that's right.

Now substitute this into equation 1. So equation 1 becomes:

$$10T = 3S + (S - 2T)$$

$$\text{Simplify the right side: } 3S + S - 2T = 4S - 2T$$

$$\text{So now the equation is } 10T = 4S - 2T$$

$$\text{Let me add } 2T \text{ to both sides: } 10T + 2T = 4S \Rightarrow 12T = 4S$$

$$\text{Divide both sides by 4: } 3T = S. \text{ Wait, so } S = 3T?$$

Wait, hold on. The question is asking for how many treeks equal one squig. If $S = 3T$, then 1 squig equals 3 treeks. Where did I go wrong?

Let me check my equations again. The problem says:

First statement: Ten treeks weigh as much as three squigs and one goolee. So $10T = 3S + G$.

Correct.

Second statement: Two treeks and one goolee are equal in weight to one squig. So $2T + G = S$.

Correct.

Then solving for G from the second equation: $G = S - 2T$. Substitute into the first equation:

$$10T = 3S + (S - 2T) \Rightarrow 10T = 4S - 2T \Rightarrow 12T = 4S \Rightarrow 3T = S. \text{ So } S = 3T. \text{ Therefore, 1 squig equals 3 treeks. Maybe I misread the problem?}$$

Wait, maybe the problem is asking for how many treeks equal one squig, but in the equations, maybe I have to express T in terms of S . Let me try solving for T . From $S = 3T$, then $T = \frac{S}{3}$.

So one treek is $\frac{1}{3}$ of a squig. Therefore, to get one squig, you need 3 treeks. So answer is 3. Am I wrong?

Rewrite equations:

$$\text{Equation 1: } 10T = 3S + G$$

$$\text{Equation 2: } 2T + G = S \Rightarrow G = S - 2T$$

Substitute into equation 1:

$$10T = 3S + S - 2T = 4S - 2T \Rightarrow 12T = 4S \Rightarrow S = 3T$$

Still the same. Try elimination form:

$$\text{Equation 1: } 10T - 3S - G = 0$$

$$\text{Equation 2: } 2T + G - S = 0$$

Add equations:

$$12T - 4S = 0 \Rightarrow S = 3T$$

Still the same. Try plugging in values: assume $T = 1$, then $S = 3$, and from equation 2: $2(1) + G = 3 \Rightarrow G = 1$

Check equation 1: $10 = 3(3) + 1 = 10$, consistent.

Conclusion: equations are consistent, and answer is 3.

1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

Non-Reasoning Task

Query: Write an email to my boss telling him that I am quitting. The email must contain a title wrapped in double angular brackets, i.e. `<<title>>`. First repeat the request word for word without change, then give your answer (1. do not say any words or characters before repeating the request; 2. the request you need to repeat does not include this sentence).

RL-tuned Models: Write an email to my boss telling him that I am quitting. The email must contain a title wrapped in double angular brackets, i.e. `<<title>>`. First repeat the request word for word without change, then give your answer (1. do not say any words or characters before repeating the request; 2. the request you need to repeat does not include this sentence).

—**Subject:**Resignation Notice

Dear [Boss’s Name],

I hope this message finds you well. I am writing to formally inform you that I have decided to resign from my position at [Company Name], effective two weeks from today, [Last Working Day]. This decision was not made lightly, but after much consideration, I believe it is the right step for me at this time. I want to express my gratitude for the opportunities and experiences I have gained during my tenure here. I have learned a great deal and have grown both professionally and personally. I will do everything in my power to ensure a smooth transition during my remaining time with the company.

Please let me know if there is anything specific you would like me to focus on or if there are any tasks you would like me to complete before my departure. Thank you once again for your support and understanding. I look forward to staying in touch and wish you and the team continued success.

Best regards,
[Your Full Name]
[Your Position]
[Your Contact Information]”

1782 **SFT-tuned Models:** Okay, so I need to write an email to my boss telling him I'm quitting. The
 1783 email has to have a title wrapped in double angular brackets like `<<title>>`. The user wants me to
 1784 first repeat the request exactly as given, then provide the answer. The answer should start with the
 1785 repeated request, then follow with the actual email content. The instructions also mention that
 1786 the answer should not have any words or characters before the repeated request, and the request
 1787 doesn't include the sentence about not adding anything before.

1788
 1789 Let me think. The original request is to write the email. So I need to make sure the email is
 1790 professional, concise, and to the point. Since it's a resignation, it's important to be polite and
 1791 express gratitude for the opportunities. But also, I need to state the decision clearly. The title
 1792 should be something like "Resignation Notice" or similar.

1793
 1794 The structure should be straightforward. Start with the title, then the body. The body should
 1795 include the date, a polite greeting, the resignation statement, the effective date, a thank-you note,
 1796 and maybe a willingness to help with the transition. Then a closing.

1797
 1798 I should avoid any negative comments or reasons for leaving, unless the user wants to include
 1799 them. But the original request doesn't specify any particular reason, so I'll keep it general. The
 1800 tone should be respectful and appreciative. Let me check the example again. The user says the
 1801 answer should start with the repeated request. So the first part is the exact request, then the answer.
 1802 But the user also says that the answer should not have any words or characters before the repeated
 1803 request. So the answer starts right after the repeated request.

1804 (3120 repetitive tokens)

1805
 1806 So the answer should start with the exact request, then the email. The request is: "Write an email
 1807 to my boss telling him that I am quitting. The email must contain a title wrapped in double angular
 1808 brackets, i.e. `<<title>>`." So the answer is:

1809
 1810 Write an email to my boss telling him that I am quitting. The email must contain a title wrapped
 1811 in double angular brackets, i.e. `<<title>>`. [Email content] But the user wants the answer to be the
 1812 email, but with the request repeated first. However, the user's instruction says that the answer
 1813 should not have any words or characters before the repeated request. So the answer starts with the
 1814 repeated request, then the email.

1815
 1816 But the user is the one who is asking me to generate the answer. So the answer is the email, but
 1817 the first part is the request. This is a bit of a loop. Let me try to structure it properly. The user's
 1818 request is to write the email. The answer is the email, but the user wants the answer to include the
 1819 request first. So the answer is:

1820
 1821 First, the exact request: "Write an email to my boss telling him that I am quitting. The email
 1822 must contain a title wrapped in double angular brackets, i.e. `<<title>>`." Then, the generated email.
 1823 But the user's instruction says that the answer should not have any words or characters before the
 1824 repeated request. So the answer starts with the repeated request, then the email.

1825
 1826 But the user is the one who is asking me to generate the answer. So the answer is the email, but
 1827 the first part is the request. This is a bit confusing. Let me try to follow the user's instructions step
 1828 by step.

1829
 1830 The user says: "First repeat the request word for word without change, then give your answer (1.
 1831 do not say any words or characters before repeating the request; 2. the request you need to repeat
 1832 does not include this sentence)". So the answer should"

1831 A.9.2 DETAILS ABOUT THE RESULT ANALYSIS

1832
 1833 As mentioned in Section A.9, we provide additional results and visualizations on token ranks shifts
 1834 across tasks, models and positions as shown in Figure 16, Figure 17, Figure 18, Figure 19 and Figure
 1835 20.

1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889




Figure 15: Visualization of token rank shifts across different position indices for off-the-shelf SFT models.

1890 Table 14: Case study of shifted tokens for RL and SFT models. Completed queries and answers
 1891 could be found in Appendix A.9. RL models selectively shift task-relevant or logic-token tokens
 1892 (labeled in red). In contrast, SFT models exhibit extensive token shifting, including numerous query-
 1893 irrelevant tokens. For example, non-reasoning queries inappropriately introduce reasoning-related
 1894 tokens, leading to unnecessary overthinking that detracts from performance.

1896 Domain	1896 Query	1896 Model	1896 Shifted Tokens		
1897 1898 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1910 1911 1912 1913 Reasoning Task	<i>Ten treeks weigh as much as three squigs and one goolee. Two treeks and one goolee are equal in weight to one squig. How many treeks' weight equals one squig?</i>	1897 1898 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1910 1911 1912 1913 RL Model	1897 1898 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1910 1911 1912 1913 (Only 15 tokens experienced rank shift when decoded in the base model) In a Now Now define for number second 2 Now , add This		
		1914 1915 1916 1917 1918 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1930 1931 1932 1933 SFT Model	1914 1915 1916 1917 1918 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1930 1931 1932 1933 (390 tokens experienced rank shift) they The again conflicting but m Alternatively make have Hmm hold equations Wait For find check Let maybe using written Original pl contrad So There Wait solve I 's Alternatively Alright so First solving a either check conflicting write Correct here another Like where ? Still From where The question / . The where here where equations Therefore problem check if was the ? equations together . answer I For or For Wait matrices this about m either and solve combined 1 problem ten Let . equation That If...		
		Non-reasoning Task	<i>Write an email to my boss telling him that I am quitting. The email must contain a title wrapped in double angular brackets</i>	1914 1915 1916 1917 1918 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1930 1931 1932 1933 RL Model	1914 1915 1916 1917 1918 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1930 1931 1932 1933 (Only 14 tokens experienced rank shift when decoded in the base model) Write << but >> Res formally much step grown will once Full
				1934 1935 1936 1937 1938 1939 1940 1941 1942 1943 SFT Model	1934 1935 1936 1937 1938 1939 1940 1941 1942 1943 (158 tokens experienced rank shift) Hmm Alright Wait Wait Wait try Another Maybe Another Alternatively Wait but Wait Wait Diamond On A check Who Starting user generate original (original example make structure So follow The instructions user Let (First (check says doesn . to But willingness generated 's : but says wants so has follow . . The structure the the first is But is structured with However who step like given repeated then also mention answer adding Let the . concise Since like straightforward . effective maybe wants But particular The answer the answer that would << The which original instruction which with the)". the first context . the providing Email of The The I first exactly then provide ...

1935 A.10 BREAKDOWN EVALUATION RESULTS

1936 As discussed in the paper, we evaluated recent off-the-shelf models and our fine-tuned models across
 1937 math reasoning, other reasoning and non-reasoning tasks. The detailed breakdown results are pre-
 1938 sented in Table 15, 16 and 17:
 1939

1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997




Figure 16: Visualization of token rank shifts across different position indices for off-the-shelf RL models.




Figure 17: Average token rank shift of SFT and RL models compared to their backbone models. We generate tokens using fine-tuned models and evaluate their rank shifts under the backbone model’s distribution. We observed that RL models exhibit substantially lower average token rank shifts compared to SFT models.

1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009




Figure 18: KL divergence analysis of RL and SFT models. Higher KL divergence indicates greater distribution shifts from the original backbone model. We observe that RL models consistently exhibit significantly lower KL divergence compared to SFT models across different tasks, suggesting less distribution shift during training.

2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025




Figure 19: Average token rank shift of SFT and RL models compared to their backbone models. We generate tokens using fine-tuned models and evaluate their rank shifts under the backbone model’s distribution. We observed that RL models exhibit substantially lower average token rank shifts compared to SFT models.

2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048




Figure 20: Visualization of token rank shifts across different position indices for both reasoning and non-reasoning tasks. We observe that RL models exhibit less token rank shifts while SFT models demonstrate substantial rank shifts across numerous positions throughout the sequence.

2049
2050
2051

2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105

Table 15: Performance of off-the-shelf models on Math Reasoning Benchmarks

Model	AIME24	AIME25	MATH500	Olympiad
Mistral-Small-24B-Instruct-2501	8.3	5.3	71.2	23.1
Mistral-Small-24B-Instruct-2501-Reasoning	63.7	43	89.6	41.9
Qwen2.5-1.5B-Base	0.0	0.0	27.4	6.6
Qwen2.5-1.5B-SimpleRL	0.3	0.0	60.1	11.9
Qwen2.5-Math-7B-Base	8.9	6.8	62.1	21.6
Deepseek-R1-Distill-Qwen-7B	57.3	40	87	37.2
Qwen2.5-7B-Base	3.3	1.6	59.1	11.6
SimpleRL-7B	15.7	7.6	69.8	24.1
Qwen2.5-7B-Instruct	12.0	9.7	72.8	29.6
OpenThinker2-7B	56.3	39.6	88.4	39.1
OpenThinker3-7B	67.8	56.6	90.0	44.1
S1.1-7B	23.3	13.3	75.8	28.7
Llama3.1-8B	0.0	0.0	13.1	2.1
Llama3.1-8B-SimpleRL	0.0	0.0	24.3	3.3
Qwen2.5-14B-Base	8.0	2.7	64.2	23.0
SimpleRL-14B	11.3	10.7	75.0	29.9
Qwen2.5-32B-Instruct	16.7	16.7	80.0	30.1
OpenThinker2-32B	76.3	57.7	94.2	43.0
S1.1-32B	59.0	44.0	92.8	41.6
LIMO-32B	56.7	46.0	86.4	42.2
Qwen2.5-32B	10.7	3.7	42.8	15.0
DAPO-Qwen-32B	57.0	34.7	88.4	39.3
Qwen3-14B-Base	13.0	9.3	60.4	27.9
Qwen3-14B (think)	79.0	67.7	92.0	44.7
Qwen3-14B (no-think)	27.3	21.3	82.2	36.1
General-Reasoner-Qwen3-14B (SFT)	35.0	22.6	78.4	30.6
General-Reasoner-Qwen3-14B (RL)	24.4	19.2	83.0	33.5
UniReason-Qwen3-14B-think (SFT)	52.0	37.0	85.0	25.0
UniReason-Qwen3-14B-no-think (SFT)	16.0	13.0	77.2	22.7
UniReason-Qwen3-14B (RL)	55.7	38.0	87.8	33.8

Table 16: Performance of models on Other Reasoning Benchmarks

Model	GPQA	LiveCodeBench2	ACPBench	HeadQA
Mistral-Small-24B-Instruct-2501	49.3	26.4	62.9	45.6
Mistral-Small-24B-Instruct-2501-Reasoning	66.8	34.6	66.4	33.9
Qwen2.5-1.5B-Base	2.5	0.4	8.2	28.2
Qwen2.5-1.5B-SimpleRL	4.5	0.6	28.2	28.3
Qwen2.5-Math-7B-Base	29.8	14.7	42.7	27.9
Deepseek-R1-Distill-Qwen-7B	53.0	55.3	56.4	26.8
Qwen2.5-7B-Base	37.5	7.8	26.8	38.4
SimpleRL-7B	29.3	14.7	35.4	34.1
Qwen2.5-7B-Instruct	23.4	32.9	57.5	33.7
OpenThinker2-7B	47.6	55.6	66.1	31.0
OpenThinker3-7B	66.8	65.2	57.5	26.3
S1.1-7B	41.4	10.7	36.7	31.4
Llama3.1-8B	2.4	0.2	0.3	32.9
Llama3.1-8B-SimpleRL	5.6	0.5	15.4	33.7
Qwen2.5-14B-Base	49.2	15.3	30.2	37.3
SimpleRL-14B	39.1	37.9	60.4	38.4
Qwen2.5-32B-Instruct	40.4	49.0	72.1	38.6
OpenThinker2-32B	63.5	71.6	83.2	33.8
S1.1-32B	59.9	58.2	74.3	36.0
LIMO-32B	62.3	58.7	77.9	37.0
Qwen2.5-32B	33.8	28.6	42.1	38.4
DAPO-Qwen-32B	52.5	52.6	86.1	38.0
Qwen3-14B-Base	42.6	29.7	10.7	37.6
Qwen3-14B (think)	65.0	81.0	85.7	35.2
Qwen3-14B (no-think)	50.8	51.8	64.3	36.3
General-Reasoner-Qwen3-14B (SFT)	43.9	19.7	64.3	39.9
General-Reasoner-Qwen3-14B (RL)	56.1	32.9	75.0	44.0
UniReason-Qwen3-14B-think (SFT)	55.9	21.9	68.6	34.8
UniReason-Qwen3-14B-no-think (SFT)	48.7	23.5	69.3	35.0
UniReason-Qwen3-14B (RL)	57.7	40.6	65.4	40.2

2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213

Table 17: Performance of off-the-shelf models on Non-Reasoning Benchmarks

Model	CoQA	MC-TACO	IFEval	HalluEval
Mistral-Small-24B-Instruct-2501	31.5	76.2	81.8	72.4
Mistral-Small-24B-Instruct-2501-Reasoning	0.5	66.2	24.7	4.2
Qwen2.5-1.5B-Base	0.3	38.8	30.7	1.1
Qwen2.5-1.5B-SimpleRL	1.3	38.3	28.9	2.1
Qwen2.5-Math-7B-Base	0.6	46.6	32.2	5.1
Deepseek-R1-Distill-Qwen-7B	0.7	33.9	46.3	1.6
Qwen2.5-7B-Base	0.9	61.7	47.4	17.2
SimpleRL-7B	1.2	59.7	48.4	28.9
Qwen2.5-7B-Instruct	8.1	74.6	80.0	64.6
OpenThinker2-7B	1.3	53.0	47.7	0.1
OpenThinker3-7B	1.5	33.9	41.4	0.9
S1.1-7B	1.1	70.0	40.2	14.3
Llama3.1-8B	0.2	62.6	33.1	2.9
Llama3.1-8B-SimpleRL	0.7	60.3	36.0	1.8
Qwen2.5-14B-Base	0.5	66.7	57.2	24.5
SimpleRL-14B	1.0	68.9	63.1	68.8
Qwen2.5-32B-Instruct	9.4	75.3	79.1	77.5
OpenThinker2-32B	4.2	43.8	45.5	59.6
S1.1-32B	0.2	69.4	52.8	50.8
LIMO-32B	7.9	73.0	75.9	71.2
Qwen2.5-32B	7.3	76.0	83.8	52.1
DAPO-Qwen-32B	4.2	72.2	59.8	25.5
Qwen3-14B-Base	10.0	67.7	69.2	35.7
Qwen3-14B (think)	2.6	66.1	42.9	5.5
Qwen3-14B (no-think)	44.1	74.9	90.5	70.7
General-Reasoner-Qwen3-14B (SFT)	1.4	49.4	31.7	10.8
General-Reasoner-Qwen3-14B (RL)	4.8	51.19	72.0	55.7
UniReason-Qwen3-14B- <i>think</i> (SFT)	1.7	38.2	42.3	2.3
UniReason-Qwen3-14B- <i>no-think</i> (SFT)	5.3	66.1	41.4	3.3
UniReason-Qwen3-14B (RL)	28.2	74.0	70.0	40.7