

TARo: Token-level Adaptive Routing for LLM Test-time Alignment

Arushi Rai*^{1 2} Qiang Zhang¹ Hanqing Zeng¹ Yunkai Zhang³
Dipesh Tamboli¹ Xiangjun Fan¹ Zhuokai Zhao^{† 1} Lizhu Zhang^{† 1}

*Work done during internship at Meta.

† Joint last author

¹Meta ²University of Pittsburgh ³University of California, Berkeley
arr159@pitt.edu {qiangzhang, zhuokai, lizhu}@meta.com

Abstract

Large language models (LLMs) exhibit strong reasoning capabilities but typically require expensive post-training to reach high performance. Recent test-time alignment methods offer a lightweight alternative, but have been explored mainly for preference alignment rather than reasoning. To bridge this gap, we propose *Token-level Adaptive Routing* (TARo), which steers frozen LLMs toward structured reasoning entirely at inference time. Specifically, we first train reward models on step-wise mathematical traces to capture fine-grained logical consistency signals, then introduce a learnable token-level router that automatically controls the guidance of the reward model to the base model. Extensive experiments show that TARo significantly improves reasoning performance by up to +22.4% over base model and +8.4% over existing token-level test-time alignment methods, while also boosting out-of-distribution clinical reasoning (MedXpertQA) and instruction following (AlpacaEval). Furthermore, TARo also generalizes from small to large backbones without retraining, extending test-time alignment from preference optimization to robust, cross-domain reasoning.

1 Introduction

Large Language Models (LLMs) have achieved impressive performance across many natural language tasks (OpenAI, 2024; Guo et al., 2025; Team et al., 2025; Yang et al., 2025a). On complex domains, such as mathematics, science, and clinical reasoning, it remains challenging to reliably solve logically demanding problems (Mirzadeh et al., 2024; Wang et al., 2025a; Cui et al., 2025; Wang et al., 2025b; Xiong et al., 2026). Recent advances in LLMs post-training, especially reinforcement learning with verifiable reward (RLVR) approaches such as group relative policy optimization (GRPO) (Shao et al., 2024), dynamic sampling policy optimization (DAPO) (Yu et al., 2025a) and

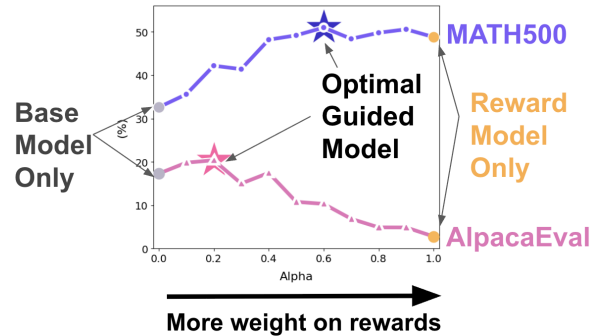


Figure 1: Performance on MATH500 (accuracy) and AlpacaEval (length-controlled win rate) for the state-of-the-art test-time alignment approach (GenARM) under different mixing coefficients $\alpha \in [0, 1]$. An $\alpha = 0$ corresponds to decoding solely from the base model, while $\alpha = 1$ uses only the reward model.

others (Zheng et al., 2025a; Yang et al., 2025b), have substantially improved reasoning. However, post-training approaches require costly model updates (Casper et al., 2023; Hou et al., 2024), tend to be domain-specific (Wu et al., 2025; Qi et al., 2024), and often degrade non-reasoning capabilities or disrupt previously learned user preferences (Chen et al., 2024; Xiao et al., 2025). Moreover, retraining becomes increasingly impractical for larger LLMs and especially prohibitive when robust reasoning is needed across multiple, frequently changing domains (e.g. clinical guidelines, legal standards, financial regulations, etc).

Test-time alignment offers a lighter, and versatile alternative by steering the *base* model (a frozen LLM) during decoding with a *reward* model (usually a smaller LLM) that provides domain expertise or user preference signals complementary to the base model (Pan et al., 2025; Zhang et al., 2025). While this paradigm avoids costly retraining, existing approaches typically rely on fixed interpolation weights between the base and reward models (Xu et al., 2024), requiring manual tuning and offering no mechanism to adapt guidance as decoding un-

folds or as domains change. In general-purpose deployments, where a model must handle diverse requests across tasks and domains, this rigidity becomes a significant limitation. Furthermore, as base models are scaled (Xu et al., 2024), the optimal balance between base and reward model guidance shifts, yet fixed interpolation weights offer no mechanism to accommodate this.

As shown in Fig. 1, the performance of existing token-level test-time alignment is sensitive to hyperparameter choices, and the optimal hyperparameter varies substantially across domains and model families. For instance, fixing the interpolation weight at $\alpha = 0.5$ as in GenARM (Xu et al., 2024) can even cause the guided model to underperform the base model on certain tasks. Although test-time alignment is attractive for its flexibility, this lack of adaptive control over reward guidance limits robustness when transferring across tasks, domains, and model scales.

In this work, we propose **Token-level Adaptive Routing (TARO)**, that enables robust reasoning improvement without retraining the base model. We first train the reward model on step-wise mathematical reasoning traces to capture fine-grained logical consistency signals. Then, to make this reward guidance effective across domains and model scales, we introduce a learnable token-level router that dynamically combines the base and reward model outputs at each decoding step, eliminating the need for manual hyperparameter tuning and improving stability when transferring across tasks, domains, and model families.

We evaluate TARO on both reasoning and non-reasoning benchmarks, including MATH500 (Lightman et al., 2023a) for mathematical reasoning, MedXpertQA (Zuo et al., 2025) for out-of-distribution clinical reasoning, and AlpacaEval (Li et al., 2023) for instruction following. Our method consistently outperforms state-of-the-art test-time alignment methods, achieving up to +22.4% accuracy over the base model and +8.4% over GenARM (Xu et al., 2024) on MATH500 (Lightman et al., 2023a), while also delivering robust gains on out-of-distribution tasks such as clinical reasoning and instruction following. Notably, the proposed router exhibits *weak-to-strong generalization*: when trained on smaller models, it transfers effectively to much larger backbones (base models) without retraining, indicating that the learned token-level modulation is both scale- and architecture-agnostic.

To summarize, our contributions are threefold:

- ① *Token-level reasoning rewards*: we show that step-wise mathematical traces can train effective reward models for test-time reasoning guidance.
- ② *Adaptive token-level router*: we propose a lightweight, learnable router that removes the need for manual interpolation tuning by dynamically blending base and reward logits.
- ③ *Robust, transferable reasoning*: TARO consistently improves reasoning across domains and model scales without additional training, extending test-time alignment from preference optimization to general, cross-domain reasoning.

2 Related Work

Test-time alignment. Expensive policy optimization methods have motivated a shift toward dynamic alignment approaches operating at inference. Some test-time alignment methods, such as Best-of-N sampling (Gao et al., 2022), rely on trajectory-level rewards and require multiple complete forward passes. Others apply trajectory-level reward models at each decoding step over full rollouts (Chakraborty et al., 2024; Huang et al., 2024) or partial rollouts (Khanov et al., 2024; Li et al., 2024a), making them prohibitively costly. In contrast, GenARM (Xu et al., 2024) learns to predict token-level rewards from preference data, eliminating the need for rollouts altogether. Concurrent work to ours, UniR (Kim et al., 2025) also explores test-time alignment for reasoning, training a reward model with GRPO (Shao et al., 2024) alongside a frozen base model. Our approach is more similar to GenARM: we learn a mathematical reasoning reward model from preference data, independently of the base model. Beyond both GenARM and UniR, we further study how to achieve robust and adaptive test-time reasoning without relying on fixed reward and base model interpolation.

Post-training methods for reasoning. Supervised finetuning (SFT) (Guha et al., 2025; Li et al., 2024b) has been used to enhance reasoning ability during post-training from datasets distilled from more advanced models (DeepSeek-AI et al., 2025; OpenAI, 2024) or carefully curated datasets like (Yue et al., 2023; Ye et al., 2025). Recently, reinforcement learning from verifiable rewards (Shao et al., 2024; Yu et al., 2025b; Liu et al., 2025; Zheng

et al., 2025b) have been used to significantly improve the reasoning ability of large language models. Our method seeks to improve the reasoning ability of LLMs as well, but does not require training the base or policy model.

Mixture of Experts. Mixture-of-Experts models have recently emerged as the state-of-the-art architecture for improving LLM capacity (Fedus et al., 2021; Dai et al., 2024; Yang et al., 2025a). In MoE, each expert specializes in a task domain (Li et al., 2022; Sukhbaatar et al., 2024) where a router selects the most suitable experts for different input data. Recently, MoE has also been applied as adapters (Li et al., 2024c; Tian et al., 2024; Zeng et al., 2025) on top of a frozen base model in parameter-efficient fine-tuning use cases. In most existing works, however, the router is conditioned on the model’s hidden embeddings, tightly coupling routing decisions to the current set of experts. Such designs make it less flexible if we were to replace experts at test time, as would be required to demonstrate weak-to-strong generalization (Xu et al., 2024). In this work, we adapt the idea of MoE for test-time alignment, where we treat the base and reward models as experts and train a lightweight router separately. We tailor the routing mechanism to be conditioned on output logits (i.e., token confidence scores) rather than hidden states, so that no retraining is needed when scaling to stronger base models. This yields a flexible and lightweight approach to test-time alignment.

3 Method

3.1 Preliminaries

Our work builds on GenARM (Xu et al., 2024), which reformulates trajectory-level reward $r(x, y)$, a scalar score assigned to a full input-output pair, into token-level rewards produced by a reward model. Formally, let $x \in \mathcal{X}$ denote a prompt (input sequence), and let $y = (y_1, \dots, y_{|y|}) \in \mathcal{Y}$ denote the completion (LLM response) as a sequence of tokens. GenARM models the reward as the log-likelihood of the trajectory under a reward-parameterized language model π_{reward} , i.e.

$$r(x, y) = \sum_{t=1}^{|y|} \log \pi_{\text{reward}}(y_t \mid x, y_{<t})$$

where $y_{<t} = (y_1, \dots, y_{t-1})$ denotes the prefix up to step $t-1$. In GenARM, the reward model π_{reward} is trained on human preference data (Ouyang et al.,

2022) using a preference loss. At inference, GenARM combines the base model and the reward model at the token level. Specifically, the next-token distribution is given by a weighted sum of the base and reward model distributions:

$$\begin{aligned} \pi_{\text{guided}}(y_t \mid x, y_{<t}) &= \pi_{\text{base}}(y_t \mid x, y_{<t}) \\ &\quad + \alpha \pi_{\text{reward}}(y_t \mid x, y_{<t}) \end{aligned}$$

where α is a scalar controlling the influence of the reward model. Next, we introduce our proposed method, TARO. At a high-level, we first train a *reasoning reward model*, and then learn a lightweight *token-level router* that adaptively combines the logits from the base and reward models during decoding. We detail each component below.

3.2 Reasoning Reward LLM

Unlike GenARM (Xu et al., 2024), which learns rewards from preference pairs without explicitly modeling reasoning, we train a reasoning-aware reward model that directly targets stepwise logical correctness. In practice, we use the Math-StepDPO-10K (Lai et al., 2024) dataset, which captures fine-grained reasoning dynamics by constructing preference pairs between two step completions that share an identical correct reasoning *prefix* but diverge at the next step, yielding one logically valid continuation y_w and one incorrect continuation y_l . Our reasoning reward model is thus optimized to prefer logically coherent, step-by-step reasoning over erroneous continuations by minimizing the standard preference loss (Ouyang et al., 2022; Bradley and Terry, 1952):

$$l_{\text{pref}} = -\log \sigma \left(\beta_r r([x, \text{prefix}], y_w) - \beta_r r([x, \text{prefix}], y_l) \right),$$

where $[x, \text{prefix}]$ denotes the concatenation of the question x and the shared correct reasoning prefix, β_r is a temperature-like scaling factor, and $\sigma(\cdot)$ is the logistic sigmoid. This objective encourages the reward model π_{reward} to assign higher scores to steps that continue the reasoning correctly and lower scores to invalid ones.

To reconcile step-level supervision with the token-level reward parameterization, we treat each step $y = (y_1, \dots, y_{|y|})$ as a short trajectory and decomposes its scalar reward into per-token log-likelihood under the reward model.

$$r([x, \text{prefix}], y) = \sum_{t=1}^{|y|} \log \pi_{\text{reward}}(y_t \mid [x, \text{prefix}], y_{<t}),$$

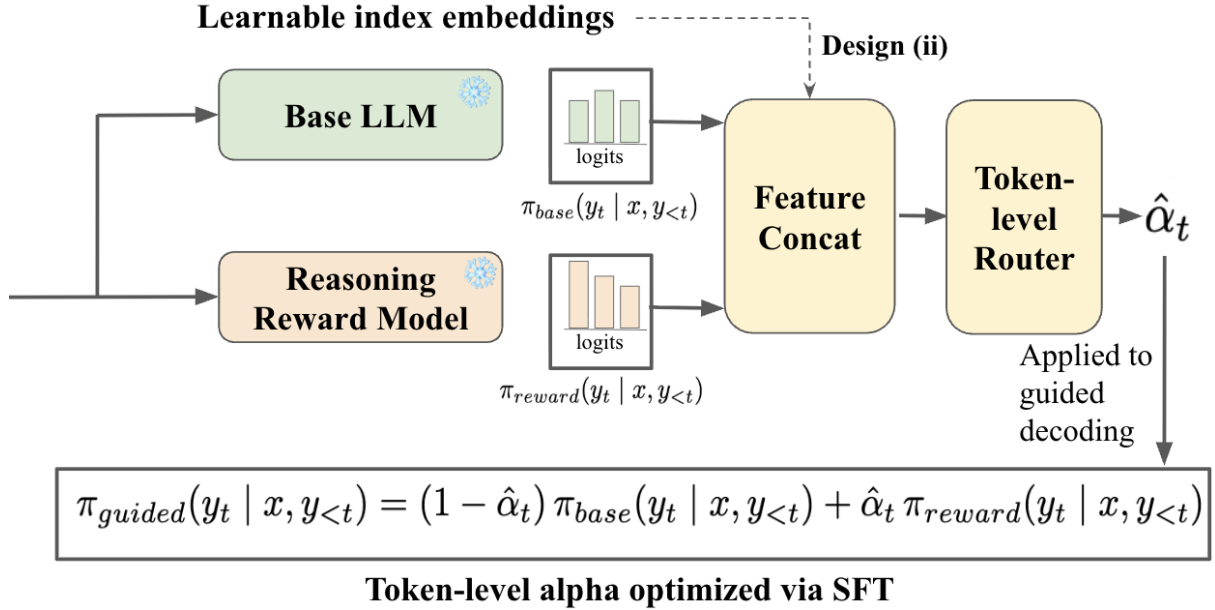


Figure 2: Learnable token-level router design. At each LLM decoding step t , the base and reward models produce logits z_t^{base} and z_t^{reward} . The logits are passed as input to *Feature Concat*, which either (i) concatenates logits, or (ii) concatenates logits plus learnable token-index embeddings (as discussed in §3.3). The router consumes the concatenated feature and outputs a routing weight $\alpha_t \in (0, 1)$. The guided distribution $(1 - \alpha_t) z_t^{\text{base}} + \alpha_t z_t^{\text{reward}}$ is then used for sampling next token. This design makes the router portable across base model scales and families.

where $y_{<t} = (y_1, \dots, y_{t-1})$. This decomposition enables step-level preferences to supervise token-wise reward signals, preserving the fine-grained token-level formulation while aligning it with fine-grained reasoning-driven correctness rather than flat, response-level preferences.

3.3 Learnable Token-level Router

A fixed interpolation between base and reward model logits, as shown in Fig. 1, is fragile: one global mixing coefficient (i.e., α) must simultaneously work across domains, model scales, and every decoding step. To make the alignment adaptive, we introduce a lightweight token-level router that dynamically chooses how much to trust and utilize each model at every generation step.

Concretely, at decoding step t , the frozen base and reward models produce logits z_t^{base} and z_t^{reward} . A lightweight feed-forward router g_θ takes input features derived from these two logits and outputs a scalar adaptive value α_t . We define this routing coefficient as $\alpha_t = \sigma(\alpha_t) \in (0, 1)$, which determines how much the guided decoding should follow the reward model versus the base model:

$$z_{\text{guided}}(\cdot | x, y_{<t}) = (1 - \alpha_t) z_t^{\text{base}} + \alpha_t z_t^{\text{reward}}.$$

Essentially, the proposed α_t adaptively adjusts the influence of the reward model token by token, in

contrast to a fixed α that is expected to work universally across domains and model scales.

To prepare the router’s input for predicting α_t , we investigate two feature constructions that fundamentally differ in whether token position (index) information from the base and reward models is explicitly encoded.

Full-logits concatenation. In this design, we concatenate both logits from the base and reward model, before passing them through a small multi-layer perceptron (MLP), i.e.

$$h_t^{\text{full}} = [z_t^{\text{base}}; z_t^{\text{reward}}] \in \mathbb{R}^{2V},$$

where V is the size of vocabulary. This design is straightforward and utilizes the existing logit distributions of both models.

Top- k logits with index embedding. Instead of operating purely in the raw logits space, we also consider explicitly encoding token index information by pairing each selected logit $z_{t,i}$ with a learnable index embedding $e_{t,i}$. For each chosen token $z_{t,i}$ from the base model we form a feature vector that combines its logit value with its embedding; the same is done for tokens j from the reward

model. Formally, we have:

$$\begin{aligned} u_{t,i}^{\text{base}} &= [z_{t,i}^{\text{base}}; e_i] \in \mathbb{R}^{d+1}, \\ u_{t,j}^{\text{reward}} &= [z_{t,j}^{\text{reward}}; e_j] \in \mathbb{R}^{d+1}. \end{aligned}$$

Here $e_i = \mathbb{E}(i)$ comes from a d -dimensional learnable encoder \mathbb{E} , allowing the router to represent each token’s identity rather than treating all tokens the same in the logit vector. Since the position information is explicitly encoded, in practice we can restrict the inputs to only the top- k tokens from each model. This keeps the feature representation compact while preserving the most informative candidates for routing.

We then concatenate all index-augmented features from both models into a single vector:

$$h_t^{\text{top-}k} = [u_{t,1}^{\text{base}}, \dots, u_{t,k}^{\text{base}}; u_{t,1}^{\text{reward}}, \dots, u_{t,k}^{\text{reward}}].$$

This vector $h_t^{\text{top-}k} \in \mathbb{R}^{2K(d+1)}$ is passed through the same MLP as in the full-logits design to produce α_t . Note that k is meant to be very small, i.e., $|h_t^{\text{top-}k}| \ll |h_t^{\text{full}}|$.

Router design. In both cases, the resulting representation h_t is passed through the same shallow MLP to predict the routing weight:

$$\hat{\alpha}_t = \sigma(W_2 \phi(W_1 h_t + b_1) + b_2), \quad (1)$$

where ϕ is the Tanh activation and σ is the sigmoid function, constraining $\hat{\alpha}_t \in (0, 1)$.

To promote confident routing behavior, we optionally add an entropy regularizer on $\hat{\alpha}_t$. This encourages the router to avoid indecisive values (e.g., $\hat{\alpha}_t \approx 0.5$) when the base and reward models diverge, thereby helping it to commit to the source it considers more reliable.

The overall training objective combines standard negative log-likelihood with the entropy penalty:

$$\begin{aligned} \mathcal{L}_{\text{router}} &= - \sum_t \log \pi_{\text{guided}}(y_t^* | x, y_{<t}) \\ &\quad + \lambda_{\text{entropy}} \sum_t H(\hat{\alpha}_t), \end{aligned} \quad (2)$$

where y_t^* is the gold target token and

$$H(\hat{\alpha}_t) = -\hat{\alpha}_t \log \hat{\alpha}_t - (1 - \hat{\alpha}_t) \log(1 - \hat{\alpha}_t)$$

is the Bernoulli entropy of the router’s decision. The hyperparameter $\lambda_{\text{entropy}} \geq 0$ controls the strength of this confidence regularization.

Note that *no ground truth values of $\hat{\alpha}_t$ are required*. Instead, the router implicitly optimizes $\hat{\alpha}_t$

through $\mathcal{L}_{\text{router}}$: (1) the NLL term penalizes routing decisions that reduce the likelihood of gold tokens and (2) the entropy term encourages hard routing decisions over uniformly mixing between the two models.

Final guided decoding. With the learned router, the decoding distribution becomes:

$$\begin{aligned} \pi_{\text{guided}}(y_t | x, y_{<t}) &= (1 - \hat{\alpha}_t) \pi_{\text{base}}(y_t | x, y_{<t}) \\ &\quad + \hat{\alpha}_t \pi_{\text{reward}}(y_t | x, y_{<t}). \end{aligned}$$

This allows dynamic token-level modulation of reward guidance, improving reasoning ability while mitigating performance drop across domains.

4 Experiment

4.1 Experimental Setup

Benchmarks. We evaluate TARO on two reasoning domains: MATH500 (Lightman et al., 2023b), which is in-distribution with respect to the router training recipe, and MedXpertQA (Zuo et al., 2025), which is out-of-distribution. We also include AlpacaEval (Li et al., 2023), a general instruction-following benchmark for the multi-domain experiment. AlpacaEval mainly consists of knowledge-intensive question answering, but also includes simpler reasoning tasks in mathematics and coding.

Models. We experiment with two model families: Llama-3.1 (Llama Team, 2024) and Qwen-2.5 (Qwen et al., 2025), using their instruct variants unless otherwise stated. For the reward models, we use DeepSeek-R1-Distill-Llama 8B (DeepSeek-AI et al., 2025) and Qwen2.5-0.5B. Reward models are trained on the step-wise preference reasoning dataset as discussed in §3.2.

Implementation details. Following Xu et al. (2024) and Kim et al. (2025), we train a separate reward model for each base model family. In terms of router, we train the learnable router on examples from Math-StepDPO-10K (Lai et al., 2024) and UltraFeedback (Cui et al., 2023). More training details and hyperparameters are illustrated in Appendix A.

During decoding, following Xu et al. (2024), we use standard sampling with temperature $t = 0.5$ across all models. We generate up to 512 tokens for AlpacaEval and MedXpertQA, and up to 2,048 tokens for MATH500. Prompts used in our experiments are reported in Appendix G.

| Method | MATH500 | MedXpertQA | AlpacaEval | Average |
|--|-------------|-------------|-------------|-------------|
| <i>Llama-3.1-8B Instruct Base + DeepSeek Distill Llama-8B Reward</i> | | | | |
| UniR (Kim et al., 2025) | 51.2 | – | – | – |
| Llama-3.1-8B Instruct (Base) | 32.0 | 13.0 | 17.3 | 20.8 |
| DeepSeek Llama-8B (Reward) | 47.2 | 11.6 | 2.8 | 20.5 |
| GenARM (Xu et al., 2024)* | 49.2 | 11.2 | 10.8 | 23.7 |
| TARO (Ours) | 54.4 | 13.2 | 20.8 | 29.5 |
| <i>Qwen-2.5-3B Instruct Base + Qwen-2.5-0.5B Instruct Reward</i> | | | | |
| UniR (Kim et al., 2025) | 64.8 | – | – | – |
| Qwen-2.5 3B (Base) | 60.8 | 11.8 | 33.6 | 35.3 |
| Qwen-2.5 0.5B (Reward) | 24.2 | 8.0 | 4.5 | 12.2 |
| GenARM (Xu et al., 2024)* | 56.4 | 9.8 | 18.9 | 28.4 |
| TARO (Ours) | 64.8 | <u>9.8</u> | <u>21.6</u> | <u>32.0</u> |

Table 1: Performance across reasoning (MATH500, MedXpertQA) and instruction-following (AlpacaEval) benchmarks. Reward models are trained on Math-StepDPO-10K described in §3.2. Results for UniR are taken from its original paper, * indicates evaluated with our Math-StepDPO-10K trained reasoning reward model.

Baselines. We compare TARO against (i) the *base models*, (ii) the *reward models*, and (iii) two state-of-the-art test-time alignment methods: GenARM (Xu et al., 2024) and UniR (Kim et al., 2025). For GenARM, we use the same reward model trained on MATH-StepDPO-10K (Lai et al., 2024) as in our method ($\alpha = 0.5$, equivalent to equal base and reward weighting).

4.2 Results Across Diverse Domains

Table 1 reports results across MATH500, MedXpertQA, and AlpacaEval. For the Llama-3.1 family, our method improves over both the base and reward models individually, and also outperforms GenARM. The largest gains appear on MATH500, where the reasoning reward model is extremely helpful with mathematical reasoning, and we also observe improvements on the out-of-distribution MedXpertQA domain.

For the Qwen-2.5 family, the base model is already very strong, especially on AlpacaEval. While our method does not outperform the Qwen-2.5 base model on AlpacaEval and MedXpertQA, it consistently exceeds GenARM across domains, showing that token-level routing provides more effective reward guidance than static interpolation. Importantly, our reward models are trained using a preference loss on *step-wise mathematical preference data*, which is considerably simpler than the RL objective used in GRPO for UniR (Kim et al., 2025). Despite the relatively low standalone performance of the reward model, this does not imply a lack of utility: the reward model may be overfitting to signals of step-wise mathematical reasoning, which could remain highly beneficial when used to steer

the base model. However, we are only able to demonstrate utility when this weak reward model is combined with token-level routing. Overall, we are able to show that effective *token-level reward models for math* can be constructed directly from math step-wise reasoning traces preferences.

We also compare TARO against majority voting (N=8) and show that TARO achieves higher accuracy on MATH500 with approximately 4x less compute. The results are detailed in Appendix D.

Interestingly, GenARM notably improved the performance of Llama-3.1-8B Instruct (Base) but failed to yield gains on Qwen-2.5 3B (Base) when evaluated on MATH500. We hypothesize that GenARM and other similar test-time alignment methods may be ineffective when the reward model performs substantially worse than the base model. In contrast, our proposed method adaptively controls the reward weight, which can significantly mitigate this limitation.

4.3 Weak-to-Strong Generalization

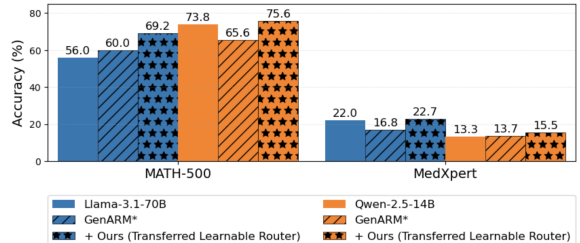


Figure 3: Weak-to-strong generalization of learned router on reasoning. Learned router and reasoning reward model are not retrained for this scale.

Using the router trained in §4.2, we evaluate

its transferability by scaling to larger base models *without* any re-training. Specifically, we pair the learned router with Llama-3.1-70B and Qwen-2.5-14B backbones. As shown in Fig. 3, our approach consistently improves over both the base and GenARM. On MATH500, the transferred router achieves substantial gains, while on MedXpertQA, improvements are smaller but still positive.

This setting highlights the weak-to-strong generalization property of both our router and reward model: routers and reward models trained on relatively small backbones can effectively steer much larger frozen LLMs. Importantly, despite being trained on limited step-wise mathematical preference data, the router provides transferable benefits even in out-of-domain reasoning tasks such as MedXpertQA.

4.4 Discussion

From our experiments in Sec. 4.2 and Sec. 4.3, we find that two factors jointly determine the effectiveness of test-time alignment: the degree of domain overlap between the reward model and the evaluation task, and the relative capability of the base model on that task. When the training domain (Math-StepDPO-10K) is the same as the evaluation benchmark (MATH500), even a weak reward model (Qwen-2.5-0.5B, 24.2% standalone) substantially improves performance when used with TARO (64.8% on MATH500). On MedXpertQA, the same pairing fails to improve over the base model, yet the performance recovers when the base model is stronger, both in the main experiments (Llama-3.1-8B in Table 1) and more strongly in the weak-to-strong generalization setting. This suggests that mathematical reasoning training has partially *transferable* benefits in the form of logical consistency and step coherence signals, but only when the base model is sufficiently capable for the router to distinguish reliable from noisy reward guidance. Gains are more limited on AlpacaEval, where the domain gap between mathematical reasoning and instruction-following is more extreme, despite AlpacaEval including some coding and mathematical questions.

Two distinct failure modes emerge from these experiments. When the base model is insufficiently strong on the evaluation task and there is domain gap, the router cannot reliably suppress noisy reward guidance. When the domain gap is extreme, as on AlpacaEval, TARO partially mitigates the negative influence of the reward model and im-

| High α Tokens | Low α Tokens |
|----------------------|---------------------|
| rightarrow | period |
| Exp | -coordinate |
| Step | profit |
| cases | factors |
| critical | shaded |
| Part | squares |
| Find | students |
| understand | time |
| begin | digits |
| Setting | bag |
| equiv | houses |
| Evaluate | digit |
| analyzing | What |
| opens | mean |
| back | pets |
| determine | sums |
| compute | bles |
| end | region |
| follow | already |

Table 2: Tokens in top 0.1% and bottom 0.1% of α from generated responses to MATH500 responses; filtered out tokens with less than 50 occurrences and shorter than 2 characters (difficult to interpret). Tokens on the left show strong reward model influence, reflecting mathematical operators, formatting, and reasoning scaffolds (e.g. “cases”, “Step”). Tokens on the right are dominated by the base model, largely consisting of vocabulary from the problem context. Base Model = Qwen2.5-3B; Reward Model = Qwen2.5-0.5B.

proves over GenARM, though gains remain limited. A notable exception is Llama-3.1-8B on AlpacaEval, where TARO improves over the base model despite the reward model having minimal standalone performance in this domain; the conditions under which this occurs remain an open question. These failure modes suggest that the limiting factor could be the breadth of the reward model’s domain coverage, and pairing TARO with multiple reasoning-specific and general-purpose reward models is a promising direction for extending robust test-time alignment across a broader range of domains.

5 Analysis and Ablation Studies

5.1 Understanding Router Behavior

We further investigate which tokens are more influenced by the reward model by analyzing the learned $\hat{\alpha}_t$ values. Table 2 shows tokens from the top 0.1% and bottom 0.1% of average $\hat{\alpha}_t$ on MATH500. High- α tokens include mathematical operators, formatting, and scaffolding symbols (e.g., “cases”, “Step”), where the reward model contributes most strongly. In contrast, low- α tokens consist mainly of problem context tokens from the base model, such as common words (“period”, “students”,

| Router Input Feature | MATH500 Accuracy (%) |
|-----------------------------|----------------------|
| Reward hidden state only | 51.2 |
| Base + reward hidden states | 49.6 |
| Reward logits only | 52.6 |
| Base + reward logits | 54.4 |

Table 3: Ablation on router input features for Llama-3.1-8B Instruct with DeepSeek Distill Llama-8B reward. Results are reported on MATH500. See Appendix C for results on all three benchmarks.

| Model | Variant | MATH | MedX |
|-------|--|-------------|-------------|
| Qwen | $K = 100$ | 60.2 | 7.7 |
| Qwen | $K = 1,000$ | 60.2 | 7.3 |
| Qwen | $K = 10,000$ | 61.8 | 7.4 |
| Qwen | $K = 100 + \lambda_{\text{entropy}} = 0.5$ | 59.8 | 7.3 |
| Qwen | $K = 1,000 + \lambda_{\text{entropy}} = 0.5$ | 64.8 | 9.8 |
| Llama | $K = 1,000 + \lambda_{\text{entropy}} = 0.5$ | 50.6 | 14.5 |
| Llama | Full logits | 54.4 | <u>13.2</u> |

Table 4: Ablation on Top- K , entropy regularization, and full-logit router designs. Accuracy (%) is reported on MATH500 (MATH) and MedXpertQA (MedX).

“already”). This analysis shows that the router allocates reward-model guidance to tokens that shape the underlying reasoning process, while relying on the base model to preserve fluent language and accurately convey problem-specific details.

5.2 Qualitative Analysis

To better understand model behavior, Table 5 shows responses from the base model, reward model, and our method to a problem from MATH500. The base model over-simplifies and arrives at an incorrect numeric answer, while the reward model identifies the correct reasoning path but becomes stuck mid-derivation. Our method succeeds by explicitly decomposing the reasoning into interpretable steps, ultimately producing the correct final answer. This illustrates the benefit of dynamically leveraging both base and reward signals at the token level. We present more examples in Appendix H.

5.3 Ablations on Token-level Router

We first validate the necessity of token-level routing granularity. As shown in Appendix E, a prompt-level router that predicts a single α for the entire sequence significantly under-performs token-level routing on MATH500 (33.2% vs. 49.6%), confirming that fine-grained, per-token control is essential for structured reasoning. Having established token-level routing, we next ablate the choice of input features for the router and effect of restricting

router inputs to the Top- k logits.

Router input feature choice. On Llama-3.1-8B Instruct and DeepSeek Distill Llama-8B reward, we experimented with several input choices for the router: (1) reward hidden state only, (2) base and reward hidden states, and (3) reward logits. Since hidden states encode token-level context, it may be natural to consider them as router inputs. As shown in Table 3, we found that using the reward hidden state alone outperforms adding the base hidden state on MATH500 (51.2% vs 49.6%), while on out-of-domain tasks adding the base hidden state improves performance (18.7% vs 15.7% on AlpacaEval and 13.7% vs 12.6% on MedXpertQA). However, base and reward hidden-state-based routers cannot reliably generalize to stronger base models, as the base hidden-state distribution shifts with model scale. This motivated our use of logits, which are scale- and domain-agnostic since they reflect model confidence over the predicted token distribution. We also found using base and reward logits yields the strongest results on MATH500 (54.4%).

Top- k ablation. We next study the effect of restricting the router inputs to the Top- k logits. The motivation for this design is to reduce noise from the full vocabulary distribution and focus the router on the most confident token candidates. As shown in Table 4, on Qwen-2.5, we find that using $K = 1000$ performs well, and adding entropy regularization further improves MATH500 accuracy to 64.8%. The entropy penalty encourages the router to make more decisive choices between the base and reward models, which is particularly effective when the reward model provides complementary signal.

On Llama-3.1, however, we observe that the entropy regularization is suboptimal on MATH500. Note in Table 1, the reasoning reward model performs relatively closer to the base model on MATH500 than on AlpacaEval. This could suggest that on most of the soft router training dataset, the base model has a balanced performance, and this could drive the router to predict consistently smaller alpha values, since the base model is more consistent in its performance. Possibly, the reasoning reward model is then underused in this configuration. This effect could suggest that the benefit of entropy regularization is sensitive to the relative strengths of the base and reward models.

Question: A worker receives an annual wage of \$20,000, which he always deposits into a savings account at the end of the year. By the end of the third year (when he makes the third deposit), he wants to have at least \$66,200 in the account to finance the purchase of a house. What is the minimal compound interest rate that the savings account must provide? Express your answer as a percentage, but do not include the percent sign.

Ground Truth: 10

| | Base Model | Reward Model | Ours |
|--------------------------|--|--|--|
| Output (abridged) | <p>Let r denote the interest rate. The worker will have a total of $20,000(1+r)^3$ in the account by the end of the third year. This is at least 66,200, so we have the inequality:</p> $20,000(1+r)^3 \geq 66,200$ $(1+r)^3 \geq 3.31$ $1+r \geq \sqrt[3]{3.31}$ $r \geq \sqrt[3]{3.31} - 1$ $r \geq 0.322$ <p>Thus, the smallest possible value of r is 0.322, or 32.2. Final Answer: 32.2. I hope it is correct.</p> | <p>Let the annual interest rate be r. The worker makes three deposits of \$20,000 at the end of each year. After the first year: $20,000(1+r)$. After the second year: $20,000(1+r)^2 + 20,000(1+r)$. After the third year: $20,000(1+r)^3 + 20,000(1+r)^2 + 20,000(1+r)$. Wait, this seems incorrect... Wait, no... The correct amount... [... stuck in confusion]</p> | <p>Step 1: The worker deposits \$20,000 each year, wants at least \$66,200 after three years.</p> <p>Step 2: Total after 3 years is $20,000(1+r)^3 + 20,000(1+r)^2 + 20,000(1+r)$.</p> <p>Step 3: Divide by 20,000: $(1+r)^3 + (1+r)^2 + (1+r) \geq 3.31$.</p> <p>Step 4: Let $x = 1+r$, then $x^3 + x^2 + x \geq 3.31$.</p> <p>Step 5: Solving gives $x \geq 1.1$, so $r \geq 0.1$.</p> <p>Final Answer: 10</p> |

Table 5: Qualitative example of base (Llama-3.1-8B Instruct), reward and our model outputs (abridged to fit) for a MATH500 question. The base model oversimplifies, while the reward model understands the right plan but gets confused. Ours solves correctly by engaging in reasoning steps.

5.4 Inference Cost

We compare the inference-time efficiency of TARO against GenARM in Table 6, measured on a single node with 8xH100 GPUs. TARO introduces a lightweight router module to dynamically adjust α , and we find that throughput is comparable to GenARM when the Top-K logits (Design ii in §3.3) is used. Expanding to full logits (Design i in §3.3) reduces throughput, particularly for large-vocabulary models such as Llama3.1-8B; however, vocabulary sizes appear to have stabilized in recent model releases, suggesting this overhead will not grow with newer models. We report both tokens-per-second (TPS) and queries-per-second (QPS). The latter does not penalize concise generations unlike TPS, making it a more complete measure of end-to-end throughput. In terms of QPS, the overhead is minimal in both settings and even lower than GenARM in the Top-K setting due to more concise outputs. More details are discussed in Appendix B.

6 Conclusion

In this paper, we introduced TARO, a test-time alignment framework that improves LLM reasoning by adaptively routing between a frozen base model and a reward model at the token level. Across mathematical reasoning, clinical reasoning, and instruction-following benchmarks, TARO consistently outperforms fixed-weight decoding base-

| Model | Method | TPS | GPU Mem. (%) | QPS |
|-------------|--------------|------|--------------|--------|
| Qwen2.5-3B | GenARM | 85.6 | 4.29 | 0.0651 |
| | TARo (Top-K) | 85.9 | 5.22 | 0.0448 |
| Llama3.1-8B | GenARM | 80.7 | 6.93 | 0.0141 |
| | TARo (Full) | 32.5 | 10.45 | 0.0161 |

Table 6: Throughput analysis on MATH500. TPS = tokens per second; QPS = queries per second. Note that these results reflect unoptimized implementations; techniques such as speculative decoding and fully sharded data-parallel inference have not been applied.

lines while preserving the flexibility and low training cost of inference-time alignment. Our results show that fine-grained reward signals, even when trained from step-wise mathematical preference data, can generalize beyond their source domain when applied through adaptive routing rather than static interpolation. We further find that the learned routing policy transfers to larger backbones without retraining, suggesting that token-level logit routing provides a scalable and portable interface for test-time reasoning control. Overall, TARO highlights that lightweight adaptive routing can be a practical path toward stronger, more robust reasoning in frozen LLMs without expensive post-training.

Limitations

While our proposed token-level adaptive routing method demonstrates strong improvements in rea-

soning performance and generalization across domains and model scales, several limitations remain. First, the effectiveness of our approach is contingent on the quality of the reward model; if the reward model is poorly calibrated or misaligned with the target reasoning task, the guidance provided may be suboptimal or even detrimental. Second, our experiments primarily focus on mathematical and clinical reasoning benchmarks as well as AlpacaEval, and the generalizability of our method to other domains or more open-ended tasks remains to be fully explored. Third, although our router is lightweight and does not require retraining of the base model, it does introduce additional inference-time computation, which may impact latency in real-world applications. Finally, our ablation studies indicate that the optimal router configuration can vary across model families, suggesting that some manual tuning or domain expertise may still be required to achieve the best results. We encourage future work to address these limitations and further investigate the robustness and applicability of token-level adaptive routing in broader settings.

References

- Ralph Allan Bradley and Milton E. Terry. 1952. [Rank analysis of incomplete block designs: I. the method of paired comparisons](#). *Biometrika*, 39:324.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, and 1 others. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Souradip Chakraborty, Soumya Suvra Ghosal, Ming Yin, Dinesh Manocha, Mengdi Wang, A. S. Bedi, and Furong Huang. 2024. Transfer q star: Principled decoding for llm alignment. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Angelica Chen, Sadhika Malladi, Lily H Zhang, Xinyi Chen, Qiuyi Zhang, Rajesh Ranganath, and Kyunghyun Cho. 2024. Preference learning algorithms do not learn preference rankings. *Advances in Neural Information Processing Systems*, 37:101928–101968.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. 2023. [Ultrafeedback: Boosting language models with high-quality feedback](#). *ArXiv*, abs/2310.01377.
- Hao Cui, Zahra Shamsi, Gowoon Cheon, Xuejian Ma, Shutong Li, Maria Tikhonovskaya, Peter Norgaard, Nayantara Mudur, Martyna Plomecka, Paul Raccuglia, and 1 others. 2025. Curie: Evaluating llms on multitask scientific long context understanding and reasoning. *arXiv preprint arXiv:2503.13517*.
- Damai Dai, Chengqi Deng, Chenggang Zhao, R. X. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y. K. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. 2024. [Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models](#). *Preprint*, arXiv:2401.06066.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Jun-Mei Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiaoling Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 179 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *ArXiv*, abs/2501.12948.
- William Fedus, Barret Zoph, and Noam Shazeer. 2021. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *CoRR*, abs/2101.03961.
- Leo Gao, John Schulman, and Jacob Hilton. 2022. [Scaling laws for reward model overoptimization](#). In *International Conference on Machine Learning*.
- Etash Kumar Guha, Ryan Marten, Sedrick Scott Keh, Negin Raoof, Georgios Smyrnis, Hritik Bansal, Marianna Nezhurina, Jean-Pierre Mercat, Trung Vu, Zayne Sprague, Ashima Suvarna, Ben Feuer, Liangyu Chen, Zaid Khan, Eric Frankel, Sachin Grover, Caroline Choi, Niklas Muennighoff, Shiye Su, and 31 others. 2025. [Openthoughts: Data recipes for reasoning models](#). *ArXiv*, abs/2506.04178.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*.
- Zhenyu Hou, Pengfan Du, Yilin Niu, Zhengxiao Du, Aohan Zeng, Xiao Liu, Minlie Huang, Hongning Wang, Jie Tang, and Yuxiao Dong. 2024. Does rlhf scale? exploring the impacts from data, model, and method. *arXiv preprint arXiv:2412.06000*.
- James Y. Huang, Sailik Sengupta, Daniele Bonadiman, Yi an Lai, Arshit Gupta, Nikolaos Pappas, Saab Mansour, Katrin Kirchoff, and Dan Roth. 2024. [Deal: Decoding-time alignment for large language models](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Maxim Khanov, Jirayu Burapachee, and Yixuan Li. 2024. [Args: Alignment as reward-guided search](#). *ArXiv*, abs/2402.01694.
- Jaemin Kim, Hangeol Chang, Hyunmin Hwang, Choonghan Kim, and Jong Chul Ye. 2025. [Universal reasoner: A single, composable plug-and-play reasoner for frozen llms](#). *ArXiv*, abs/2505.19075.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. 2024. [Step-dpo: Step-wise preference optimization for long-chain reasoning of llms](#). *ArXiv*, abs/2406.18629.
- Bolian Li, Yifan Wang, Ananth Y. Grama, and Ruqi Zhang. 2024a. [Cascade reward sampling for efficient decoding-time alignment](#). *ArXiv*, abs/2406.16306.
- Chen Li, Weiqi Wang, Jingcheng Hu, Yixuan Wei, Nan-ning Zheng, Han Hu, Zheng Zhang, and Houwen Peng. 2024b. [Common 7b language models already possess strong math capabilities](#). *ArXiv*, abs/2403.04706.
- Dengchun Li, Yingzi Ma, Naizheng Wang, Zhengmao Ye, Zhiyuan Cheng, Yinghao Tang, Yan Zhang, Lei Duan, Jie Zuo, Cal Yang, and Mingjie Tang. 2024c. [Mixlora: Enhancing large language models fine-tuning with lora-based mixture of experts](#). *Preprint*, arXiv:2404.15159.
- Margaret Li, Suchin Gururangan, Tim Dettmers, Mike Lewis, Tim Althoff, Noah A. Smith, and Luke Zettlemoyer. 2022. [Branch-train-merge: Embarrassingly parallel training of expert language models](#). *Preprint*, arXiv:2208.03306.

- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023a. Let's verify step by step. *arXiv preprint arXiv:2305.20050*.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023b. [Let's verify step by step](#). *ArXiv*, abs/2305.20050.
- Zi-Yan Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025. [Understanding rl-zero-like training: A critical perspective](#). *ArXiv*, abs/2503.20783.
- Llama Team. 2024. Introducing llama 3.1: Our most capable models to date. <https://ai.meta.com/blog/meta-llama-3-1/>. Meta AI Blog.
- Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. 2024. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *arXiv preprint arXiv:2410.05229*.
- OpenAI. 2024. Gpt-4 technical report. <https://cdn.openai.com/papers/gpt-4.pdf>. Accessed: 2025-10-05.
- OpenAI. 2024. [Learning to reason with llms](https://openai.com/index/learning-to-reason-with-llms/). <https://openai.com/index/learning-to-reason-with-llms/>. Accessed: 2025-05-01.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis Christiano, Jan Leike, and Ryan J. Lowe. 2022. [Training language models to follow instructions with human feedback](#). *ArXiv*, abs/2203.02155.
- Birong Pan, Yongqi Li, Weiyu Zhang, Wenpeng Lu, Mayi Xu, Shen Zhou, Yuanyuan Zhu, Ming Zhong, and Tiejun Qian. 2025. A survey on training-free alignment of large language models. *arXiv preprint arXiv:2508.09016*.
- Zhenting Qi, Hongyin Luo, Xuliang Huang, Zhuokai Zhao, Yibo Jiang, Xiangjun Fan, Himabindu Lakkaraju, and James Glass. 2024. Quantifying generalization complexity for large language models. *arXiv preprint arXiv:2410.01769*.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, and 1 others. 2024. Deepseek-math: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Sainbayar Sukhbaatar, Olga Golovneva, Vasu Sharma, Hu Xu, Xi Victoria Lin, Baptiste Rozière, Jacob Kahn, Daniel Li, Wen tau Yih, Jason Weston, and Xian Li. 2024. [Branch-train-mix: Mixing expert llms into a mixture-of-experts llm](#). *Preprint*, arXiv:2403.07816.
- Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, and 1 others. 2025. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*.
- Chunlin Tian, Zhan Shi, Zhijiang Guo, Li Li, and Chengzhong Xu. 2024. [Hydralora: An asymmetric lora architecture for efficient fine-tuning](#). *Preprint*, arXiv:2404.19245.
- Peng-Yuan Wang, Tian-Shuo Liu, Chenyang Wang, Yi-Di Wang, Shu Yan, Cheng-Xing Jia, Xu-Hui Liu, Xin-Wei Chen, Jia-Cheng Xu, Ziniu Li, and 1 others. 2025a. A survey on large language models for mathematical reasoning. *arXiv preprint arXiv:2506.08446*.
- Wenxuan Wang, Zizhan Ma, Meidan Ding, Shiyi Zheng, Shengyuan Liu, Jie Liu, Jiaming Ji, Wenting Chen, Xiang Li, Linlin Shen, and 1 others. 2025b. Medical reasoning in the era of llms: A systematic review of enhancement techniques and applications. *arXiv preprint arXiv:2508.00669*.
- Juncheng Wu, Sheng Liu, Haoqin Tu, Hang Yu, Xiaoke Huang, James Zou, Cihang Xie, and Yuyin Zhou. 2025. Knowledge or reasoning? a close look at how llms think across domains. *arXiv preprint arXiv:2506.02126*.
- Jiancong Xiao, Ziniu Li, Xingyu Xie, Emily Getzen, Cong Fang, Qi Long, and Weijie J Su. 2025. On the algorithmic bias of aligning large language models with rlhf: Preference collapse and matching regularization. *Journal of the American Statistical Association*, (just-accepted):1–21.
- Nuoya Xiong, Yuhang Zhou, Hanqing Zeng, Zhaorun Chen, Furong Huang, Shuchao Bi, Lizhu Zhang, and Zhuokai Zhao. 2026. Token-level llm collaboration via fusionroute. *arXiv preprint arXiv:2601.05106*.

- Yuancheng Xu, Udari Madhushani Sehwag, Alec Koppel, Sicheng Zhu, Bang An, Furong Huang, and Sumittra Ganesh. 2024. Genarm: Reward guided generation with autoregressive reward model for test-time alignment. *arXiv preprint arXiv:2410.08193*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Chenghao Yang, Lin Gui, Chenxiao Yang, Victor Veitch, Lizhu Zhang, and Zhuokai Zhao. 2025b. Let it calm: Exploratory annealed decoding for verifiable reinforcement learning. *arXiv preprint arXiv:2510.05251*.
- Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. 2025. Limo: Less is more for reasoning. *ArXiv*, abs/2502.03387.
- Qiyong Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, and 1 others. 2025a. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Qiyong Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, and 16 others. 2025b. Dapo: An open-source llm reinforcement learning system at scale. *ArXiv*, abs/2503.14476.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. 2023. Mammoth: Building math generalist models through hybrid instruction tuning. *ArXiv*, abs/2309.05653.
- Hanqing Zeng, Yinglong Xia, Zhuokai Zhao, Gilbert Jiang, Qiang Zhang, Jiayi Liu, Lizhu Zhang, Xiangjun Fan, and Benyu Zhang. 2025. S'more: Structural mixture of residual experts for llm fine-tuning. *Preprint*, arXiv:2504.06426.
- Qiyuan Zhang, Fuyuan Lyu, Zexu Sun, Lei Wang, Weixu Zhang, Wenyue Hua, Haolun Wu, Zhihan Guo, Yufei Wang, Niklas Muennighoff, and 1 others. 2025. A survey on test-time scaling in large language models: What, how, where, and how well? *arXiv preprint arXiv:2503.24235*.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, and 1 others. 2025a. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xionghui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. 2025b. Group sequence policy optimization. *ArXiv*, abs/2507.18071.
- Yuxin Zuo, Shang Qu, Yifei Li, Zhangren Chen, Xuekai Zhu, Ermo Hua, Kaiyan Zhang, Ning Ding, and Bowen Zhou. 2025. Medxpertqa: Benchmarking expert-level medical reasoning and understanding. *ArXiv*, abs/2501.18362.

Appendix

The appendix provides supporting details across eight sections. Appendix A describes implementation and hyperparameter details for training the reward model and token-level router. Appendix B analyzes the parameter overhead introduced by the router for each design variant. Appendix C extends the router input ablation to all three benchmarks. Appendix D compares TARo against majority voting on Qwen2.5-3B. Appendix E validates the choice of token-level over prompt-level routing. We provide generation and evaluation (LLM-as-a-judge) prompts in Appendix F and Appendix G, respectively. Finally, Appendix H presents additional qualitative examples illustrating cases where TARo succeeds and where it fails.

A Experiments Implementation and Hyperparameter Details

Reasoning reward model. Reward model training uses AdamW with a learning rate of 2×10^{-5} , cosine learning rate scheduler, batch size 32, and 3 epochs. We set $\beta_r = 0.1$ in the preference loss. LoRA adapters are applied with rank 8 and scaling factor $\alpha = 16$ ¹.

Token-level router. Router training uses supervised fine-tuning (SFT) to train the MLP layers (hidden dimension size $H = 128$) on 1,000 samples from each dataset for three epochs. We use a learning rate of 5×10^{-6} , batch size 32, and 10 warmup steps. Training is conducted on 1-4 NVIDIA H100 GPUs with bfloat16 precision.

B Router Complexity

Compared with GenARM, the router introduces $(2V + 1) \times (H + 1)$ learnable parameters for Design (i) and $(2Kd + 1) \times (H + 1) + d \times V$ for Design (ii). Given the hidden dimension of router prediction head $H = 128$, for Llama ($V = 128256$) and Qwen 2.5 ($V = 151936$), the router would have $33M$ and $12M$ learnable parameters respectively.

C Router Input Ablation: Full Results

| Router Inputs | Scalable to Larger Base | MATH500 (Acc %) | AlpacaEval (LC WR %) | MedXpertQA (Acc %) | Avg. |
|-----------------------------|-------------------------|-----------------|----------------------|--------------------|-------------|
| Base + reward hidden states | X | 49.6 | 18.7 | 13.7 | 27.3 |
| Reward hidden state only | ✓ | 51.2 | 15.7 | 12.6 | 26.5 |
| Reward logits only | ✓ | 52.6 | 19.7 | 14.4 | 28.9 |
| Base + reward logits | ✓ | 54.4 | 20.8 | 13.2 | 29.5 |

Table 7: Ablation of router input configurations on Llama3.1-8B across all three benchmarks. *Scalable to Larger Base* indicates whether the router can be applied to a larger base model without retraining. The greyed row is excluded from the scalable comparison. Among scalable designs, base and reward logits on average outperforms using hidden states or logits from the reward model alone. Bold indicates best result per column across all rows.

Table 7 extends the router input ablation from the main paper (§5.3) to all three benchmarks. We first compare two router input configurations: using reward hidden states alone vs. using both base and reward hidden states. The latter consistently improves performance on tasks outside the reward model’s training domain (MedXpertQA: 12.6% \rightarrow 13.7%, AlpacaEval: 15.7% \rightarrow 18.7%), suggesting that the reward model’s hidden states carry a domain bias from mathematical training that the base model’s representation can counteract. However, this configuration is not base-model-scalable, as a larger base model’s hidden states may differ in both representation and dimensionality from those seen during router training.

Logits provide a scale-agnostic alternative, where the vocabulary space remains the same across scales. We observe that reward logits alone already generalize better on OOD tasks than their hidden state counterpart (AlpacaEval: 19.7% vs. 15.7%, MedXpertQA: 14.4% vs. 12.6%). We posit this is because logits encode model confidence over a shared, domain-agnostic vocabulary space, making them a more universal signal for routing regardless of the target domain. Adding base logits further improves in-domain

¹This α differs from the interpolation coefficient in GenARM (Xu et al., 2024) and our router.

performance (MATH500: 52.6%→54.4%) at a slight OOD cost, and achieves the best average across benchmarks, motivating our choice of base and reward logits as the default configuration.

D Comparison with Majority Voting

| Method | MATH500 (Acc %) | AlpacaEval (LC WR %) | MedXpertQA (Acc %) | Est. TFLOPs |
|---------------------------|--------------------|-------------------------|-----------------------|-------------|
| Majority Voting ($N=8$) | 56.4 | — | 10.6 | 113.6 |
| GenARM | 56.4 | 18.9 | 9.8 | 28.4 |
| TARo | 64.8 | 21.6 | 9.8 | 28.4 |

Table 8: Comparison with majority voting on Qwen2.5-3B. TFLOPs are estimated for a single forward pass with a 2048-token context. Majority voting is not applicable to AlpacaEval (—) due to its open-ended generation format. Bold indicates best result per column.

Table 8 compares TARo against majority voting and GenARM on Qwen2.5-3B. Note that majority voting cannot be applied to AlpacaEval, as it depends on a known, extractable answer format. This can make majority voting poorly suited for agentic systems, as a potential application setting, that need flexible logical reasoning and instruction-following, not fixed-format outputs. TARo outperforms majority voting on MATH500 while requiring only $4\times$ less compute, as majority voting requires $N=8$ full response samples rather than token-wise generation. On MedXpertQA, TARo underperforms majority voting.

E Prompt-level vs. Token-level Routing

| Routing Granularity | MATH500 (Acc %) | AlpacaEval (LC WR %) | MedXpertQA (Acc %) | Avg. |
|---------------------|--------------------|-------------------------|-----------------------|------|
| Prompt-level | 33.2 | 22.1 | 13.3 | |
| Token-level | 49.6 | 18.7 | 13.7 | |

Table 9: Comparison of prompt-level versus token-level routing with hidden state inputs on Llama3.1-8B. Token-level routing significantly outperforms prompt-level routing on mathematical reasoning, validating the design choice of adaptive token-level α prediction.

Table 9 compares router performance when trained to predict a single prompt-level α (uniform across all tokens) versus a token-level α at each decoding step. Prompt-level routing performs significantly worse on mathematical reasoning (MATH500: 33.2% vs. 49.6%), demonstrating that fine-grained token-level control is essential for structured reasoning tasks where the reward model’s guidance should be concentrated on specific tokens such as operators, variables, and reasoning scaffolds (see Tab. 2 from main text). Performance differences on OOD tasks (AlpacaEval, MedXpertQA) are smaller, suggesting that token-level granularity matters most in the in-domain reasoning setting.

F Generation Prompts

We share the MATH500 and MedXpertQA generation prompts in this section.

MATH500 Generation Prompt

Solve the following math problem step by step. Please reason step by step, and put your final answer within `\texttt{\textbackslash boxed\{\}}`.

Here are some examples:

Problem:

Find the domain of the expression $\frac{\sqrt{x-2}}{\sqrt{5-x}}$.

Solution:

The expressions inside each square root must be non-negative. Therefore, $x-2 \geq 0$, so $x \geq 2$, and $5-x \geq 0$, so $x \leq 5$.

Also, the denominator cannot be equal to zero, so $5 - x > 0$, which gives $x < 5$. Therefore, the domain of the expression is $[2, 5)$.

Final Answer: The final answer is $[2, 5)$. I hope it is correct.

Problem:

If $\det \mathbf{A} = 2$ and $\det \mathbf{B} = 12$, then find $\det(\mathbf{AB})$.

Solution:

We have that $\det(\mathbf{AB}) = (\det \mathbf{A})(\det \mathbf{B}) = (2)(12) = 24$.

Final Answer: The final answer is 24. I hope it is correct.

Problem:

Terrell usually lifts two 20-pound weights 12 times. If he uses two 15-pound weights instead, how many times must Terrell lift them in order to lift the same total weight?

Solution:

If Terrell lifts two 20-pound weights 12 times, he lifts a total of $2 \cdot 12 \cdot 20 = 480$ pounds of weight. If he lifts two 15-pound weights instead for n times, he will lift a total of $2 \cdot 15 \cdot n = 30n$ pounds of weight. Equating this to 480 pounds, we can solve for n :

$$30n = 480 \\ \Rightarrow n = 480/30 = 16$$

Final Answer: The final answer is 16. I hope it is correct.

Problem:

If the system of equations

$$\begin{cases} 6x - 4y = a \\ 6y - 9x = b \end{cases}$$

has a solution (x, y) where x and y are both nonzero, find $\frac{a}{b}$, assuming b is nonzero.

Solution:

If we multiply the first equation by $-\frac{3}{2}$, we obtain

$$6y - 9x = -\frac{3}{2}a. \text{ Since we also know that } 6y - 9x = b, \text{ we have}$$

$$-\frac{3}{2}a = b \Rightarrow \frac{a}{b} = -\frac{2}{3}.$$

Final Answer: The final answer is $-\frac{2}{3}$. I hope it is correct.

MedXpertQA Generation Prompt

Instructions

- * Select the multiple choice answer that best answers the question.
- * Please think step by step.
- * Write your final answer as "Selected Answer Choice: {answer}"

Examples

User: A 29-year-old woman presents to the emergency department with a 1-day history of fever. She says that it started spontaneously and that it hasn't gotten any better with either rest or acetaminophen. Her past medical history is significant for homelessness, unprotected sexual intercourse with multiple partners, and substance use disorder with multiple substances. Her temperature is 103.5F (39.5C), blood pressure is 100/70 mmHg, pulse is 112/min, respirations are 18/min, and oxygen saturation is 93% on room air. On physical exam, a murmur is heard along the left mid-sternal border. The pulmonary exam reveals minor bibasilar crackles. The patient's upper limbs demonstrate many bruises and scars in the antecubital fossa. The patient is started on vancomycin and gentamicin and sent for echocardiography. Based on the results of the echocardiogram, the patient undergoes surgery to remove multiple vegetations from the tricuspid valve. Vancomycin and gentamicin are continued over the next 5 days. On postoperative day 5, the patient presents with bleeding from her gums, oozing from her surgical sites, and recurrent epistaxis. Lab values are obtained as seen below:

Serum:Na+: 135 mEq/L

Cl-: 90 mEq/L

K+: 4.4 mEq/L

HCO3-: 23 mEq/L

BUN: 20 mg/dL

Glucose: 110 mg/dL

Creatinine: 1.0 mg/dL

Ca2+: 10.1 mg/dL

AST: 9 U/L

ALT: 9 U/L

Leukocyte count and differential:

Leukocyte count: 6,000 cells/mm³

Lymphocytes: 20%

Monocytes: 1%

Neutrophils: 78%

Eosinophils: 1%

Basophils: 0%

PT: 27 seconds

aPTT: 84 seconds

D-dimer: < 50 µg/L

Hemoglobin: 14 g/dL

Hematocrit: 40%

Platelet count: 150,000/mm³

Mean corpuscular volume: 110 fL

Mean corpuscular hemoglobin concentration: 34 g/dL

RDW: 14%

Which of the following is the most likely cause of this patient's current symptoms?

Answer Choices: (A) Bacterial infection of the bloodstream (B) Defect in von Willebrand factor (C) Acquired hemophilia A (D) Factor VIII deficiency (E) Coagulation cascade activation (F) Disseminated intravascular coagulation (G) Vitamin K deficiency due to antibiotic use (H) Immune-mediated thrombocytopenia (I) Antibiotic therapy (J) Hypersensitivity reaction to antibiotics

<leot_id>

Assistant: To determine the most likely cause of this patient's current symptoms, let's analyze the situation step by step. A 29-year-old woman presents to the emergency department with a 1-day history of fever. She has a history of homelessness, unprotected sexual intercourse with multiple partners, and substance use disorder. Her vital signs show a high fever, tachycardia, and slightly low oxygen saturation. A murmur is heard along the left mid-sternal border, and there are minor bibasilar crackles on pulmonary exam. The presence of bruises and scars in the antecubital fossa suggests intravenous drug use. She is started on vancomycin and gentamicin and undergoes surgery to remove vegetations from the tricuspid valve, indicating infective endocarditis. On postoperative day 5, she presents with bleeding from her gums, oozing from surgical sites, and recurrent epistaxis. Lab values show prolonged PT and aPTT, normal platelet count, and normal D-dimer, which suggests a coagulopathy rather than thrombocytopenia or DIC. The prolonged PT and aPTT with normal platelet count and D-dimer are consistent with a coagulation factor deficiency or inhibition. Given the recent antibiotic use, vitamin K deficiency due to antibiotic therapy is a likely cause, as antibiotics can disrupt gut flora that synthesize vitamin K, leading to deficiencies in vitamin K-dependent clotting factors. Therefore, the most likely cause of this patient's current symptoms is: Antibiotic therapy.

Selected Answer Choice: I

<leot_id>

User: A 55-year-old man is brought to the emergency department after a witnessed tonic-clonic seizure. There was no trauma during the seizure. The patient's medical history includes hypertension, coronary artery disease, and type 2 diabetes. His current medications are hydrochlorothiazide, atenolol, metformin, and aspirin. According to his family, he smokes a pack of cigarettes and consumes 6 to 8 beers daily but does not use recreational drugs. His last alcohol consumption occurred approximately 36 hours before presentation.

His blood pressure is 160/110 mm Hg, his heart rate is 98 beats per minute and regular, and his respiratory rate is 22 breaths per minute. He is afebrile but confused and diaphoretic.

A complete blood count shows mild macrocytic anemia and thrombocytopenia. A biochemical profile shows a normal glucose level, mild hypokalemia, hypophosphatemia, normal kidney function, and an elevated aspartate aminotransferase level.

Chest radiography is unremarkable. Noncontrast CT of the brain shows mild volume loss. Blood and urine drug test results are negative. An electrocardiogram is notable for sinus tachycardia and motion artifact.

Thiamine, multivitamins, folate, and electrolyte repletion are administered. Three hours later, the patient becomes increasingly agitated and confused, with incoherent speech, diaphoresis, and combative behavior that persist even after the cumulative administration of 40 mg of intravenous lorazepam. Physical restraints to prevent him from disrupting his intravenous access only contribute to his agitation.

What is the most appropriate initial management option in this patient?

Answer Choices: (A) Obtain repeat CT of the brain (B) Transfer to the intensive care unit, and initiate sedation (C) Start a continuous infusion of midazolam (D) Obtain electroencephalography (E) Start a continuous infusion of phenytoin (F) Administer phenobarbital for refractory agitation (G) Perform lumbar puncture to rule out meningitis (H) Administer haloperidol for agitation (I) Initiate propofol infusion for sedation (J) Administer dexmedetomidine for sedation

<leot_idl>

Assistant: To determine the most appropriate initial management option for this patient, let's analyze the situation step by step. The patient is a 29-year-old woman with a history of homelessness, unprotected sexual intercourse with multiple partners, and substance use disorder. She presents with a high fever, tachycardia, and slightly low oxygen saturation. Physical examination reveals a murmur along the left mid-sternal border and minor bibasilar crackles. The presence of bruises and scars in the antecubital fossa suggests intravenous drug use. She is started on vancomycin and gentamicin due to suspected infective endocarditis and undergoes surgery to remove vegetations from the tricuspid valve. On postoperative day 5, she presents with bleeding from her gums, oozing from surgical sites, and recurrent epistaxis. Lab values show prolonged PT and aPTT, normal platelet count, and normal D-dimer, indicating a coagulopathy rather than thrombocytopenia or disseminated intravascular coagulation (DIC). The prolonged PT and aPTT with normal platelet count and D-dimer suggest a coagulation factor deficiency or inhibition. Given the recent antibiotic use, vitamin K deficiency due to antibiotic therapy is a likely cause, as antibiotics can disrupt gut flora that synthesize vitamin K, leading to deficiencies in vitamin K-dependent clotting factors. The most likely cause of this patient's current symptoms is antibiotic therapy.

Selected Answer Choice: B

<leot_idl>

G Evaluation prompts.

In this section, we share the evaluation prompt used by Llama3.1-70B to evaluate the extracted answer for MedXpertQA and MATH500. First, we extract the model's answer from the generated output by (1) taking the content inside “\boxed{ }” for Qwen, (2) taking the span after “Final answer:” for Llama, or (3) using the full generated output as a fallback for either if extraction based on formatting fails. The evaluation prompt considers equivalence between the extracted answer and the ground truth answer for MedXpertQA. For MATH500, the evaluation prompt provides examples to evaluate based on mathematical equivalence.

MATH500 Evaluation Prompt

Look at the following two expressions (answers to a math problem) and judge whether they are equivalent.
Only perform trivial simplifications

Examples :

Expression 1: $2x+3$
Expression 2: $3+2x$

Yes

Expression 1: $\frac{3}{2}$
Expression 2: 1.5

Yes

Expression 1: x^2+2x+1
Expression 2: y^2+2y+1

No

Expression 1: x^2+2x+1
Expression 2: $(x+1)^2$

Yes

Expression 1: $\frac{3245}{5}$
Expression 2: 649

No

(these are actually equal, don't mark them equivalent if you need to do nontrivial simplifications)

Expression 1: $\frac{2}{(-3)}$
Expression 2: $-\frac{2}{3}$

Yes

(trivial simplifications are allowed)

Expression 1: 72 degrees
Expression 2: 72

Yes

(give benefit of the doubt to units)

Expression 1: 64
Expression 2: 64 square feet

Yes

(give benefit of the doubt to units)

YOUR TASK

Respond with only "Yes" or "No" (without quotes). Do not include a rationale.

Expression 1: %(expression1)s
Expression 2: %(expression2)s

MedXpertQA Evaluation Prompt

Look at the following two answers and judge whether they are equivalent.

YOUR TASK

Respond with only "Yes" or "No" (without quotes). Do not include a rationale.

Answer 1: %(expression1)s

Answer 2: %(expression2)s

H Additional Qualitative Examples

We present additional qualitative examples comparing the base, reasoning reward, and TARO model outputs in this section. Table 10 illustrates a case where the base model bypasses the reasoning process, producing an incorrect result (200), while both the reward and guided models reason systematically and arrive at the correct answer (2220). In contrast, Table 11 presents a case where the base model quickly identifies the correct cube (27) through trial and error, but the reward and TARO follow the right strategy, but make an error during reasoning, leading to an incorrect answer. This suggests that the reward and guided model lacks mechanisms to revise their problem-solving strategy and check their work. Future work can enhance the reward model’s ability to reward backtracking and verification.

| Question: What is the least positive integer multiple of 30 that can be written with only the digits 0 and 2? | | | |
|--|---|--|---|
| Ground Truth: 2220 | | | |
| | Base Model | Reward Model | TARO (Ours) |
| Output | The least positive integer multiple of 30 that can be written with only the digits 0 and 2 is 200 . Final Answer: 200. I hope it is correct. | We need to find the smallest positive integer multiple of 30 that can be written with only the digits 0 and 2. Since 30 is a multiple of 2, 3, and 5, the number must end with a 0 (to be divisible by 10) and contain at least one 3 or 9 (to be divisible by 3). However, we are only allowed to use the digits 0 and 2, so the number must be written with only the digits 0 and 2. The smallest such number is 200, but 200 is not divisible by 3. The next possible number is 2220, which is divisible by 6 (and therefore by 3) and ends with a 0 (so it is divisible by 10). Therefore, 2220 is a multiple of 30. The smallest such number is 60, but 60 is not written with only the digits 0 and 2. The next possible number is 22020, but 22020 is not written with only the digits 0 and 2. The smallest such number is 2220, which is written with only the digits 0 and 2 and is a multiple of 30. Therefore, the answer is 2220. Final Answer: 2220. I hope it is correct. | ## Step 1: Understand the problem We need to find the smallest multiple of 30 that can be expressed using only the digits 0 and 2. ## Step 2: Recall divisibility rules Multiples of 30 are divisible by both 2 and 3, and must end with 0. ## Step 3: Divisibility by 3 The sum of digits must be a multiple of 3. With digits 0 and 2, the sum = 2×(number of 2s), so the number of 2s must be a multiple of 3. ## Step 4: Try smallest combination Three 2s and one 0 gives 2220. ## Step 5: Verification 2220 is divisible by 2, 3, and 5. ⇒ It is divisible by 30. Final Answer: 2220 |

Table 10: Qualitative example comparing **Llama-3.1-8B Instruct** (base), the **reasoning reward model**, and TARO on a divisibility problem. The base model bypasses reasoning, while both the reward and guided models reason systematically to reach the correct answer.

Question: What is the smallest positive perfect cube that can be written as the sum of three consecutive integers?
Ground Truth: 27

| | Base Model | Reward Model | TARO (Ours) |
|---------------|--|--|--|
| Output | <p>We can check small sets of three consecutive integers: $\{1, 2, 3\}$, $\{2, 3, 4\}$, $\{3, 4, 5\}$, ... The sums are 6, 9, 12, 15, 18, 21, 24, 27. The first perfect cube among these is 27, which corresponds to $\{8, 9, 10\}$. Final Answer: 27.</p> | <p>We need to find the smallest positive perfect cube that can be written as the sum of three consecutive integers. Let's denote the three consecutive integers as $n - 1$, n, and $n + 1$. Then, the sum of these three integers is $(n - 1) + n + (n + 1) = 3n$. We need to find the smallest positive perfect cube that can be written as $3n$. Let's consider the cubes of the positive integers and see if any of them can be written as the sum of three consecutive integers. For example, the cube of 1 is 1, the cube of 2 is 8, the cube of 3 is 27, the cube of 4 is 64, the cube of 5 is 125, and the cube of 6 is 216. The cube of 6 is 216. Let's check if 216 can be written as the sum of three consecutive integers. We can write 216 as $3n$, so $n = 72$. Then, the three consecutive integers are 71, 72, and 73. The sum of these three integers is $71 + 72 + 73 = 216$. Therefore, the smallest positive perfect cube that can be written as the sum of three consecutive integers is 216. Final Answer: The final answer is 216. I hope it is correct.</p> | <p>Let's take a look at the problem. We're looking for the smallest positive perfect cube that can be written as the sum of three consecutive integers. Let's start by finding the smallest perfect cube. The smallest positive perfect cube is $1^3 = 1$. Now, we need to find three consecutive integers that add up to 1. The only set of three consecutive integers that add up to 1 is 0, 1, and 2, but their sum is 3, not 1. Let's try the next perfect cube. The next smallest positive perfect cube is $2^3 = 8$. We need to find three consecutive integers that add up to 8. The set of three consecutive integers that add up to 8 is -3, -2, and -1. Their sum is -6, not 8. The next perfect cube is $3^3 = 27$. We need to find three consecutive integers that add up to 27. The set of three consecutive integers that add up to 27 is -10, -9, and -8. Their sum is -27, not 27. The next perfect cube is $4^3 = 64$. We need to find three consecutive integers that add up to 64. The set of three consecutive integers that add up to 64 is -21, -20, and -19. Their sum is -60, not 64. The next perfect cube is $5^3 = 125$. We need to find three consecutive integers that add up to 125. The set of three consecutive integers that add up to 125 is -44, -43, and -42. Their sum is -129, not 125. The next perfect cube is $6^3 = 216$. We need to find three consecutive integers that add up to 216. The set of three consecutive integers that add up to 216 is -71, -70, and -69. Their sum is -210, not 216. The next perfect cube is $7^3 = 343$. We need to find three consecutive integers that add up to 343. The set of three consecutive integers that add up to 343 is -112, -111, and -110. Their sum is -333, not 343. The next perfect cube is $8^3 = 512$. We need to find three consecutive integers that add up to 512. The set of three consecutive integers that add up to 512 is -169, -168, and -167. Their sum is -504, not 512. The next perfect cube is $9^3 = 729$. We need to find three consecutive integers that add up to 729. The set of three consecutive integers that add up to 729 is -242, -241, and -240. Their sum is -723, not 729. The next perfect cube is $10^3 = 1000$. We need to find three consecutive integers that add up to 1000. The set of three consecutive integers that add up to 1000 is -333, -332, and -331. Their sum is -996, not 1000. The next perfect cube is $11^3 = 1331$. We need to find three consecutive integers that add up to 1331. The set of three consecutive integers that add up to 1331 is -444, -443, and -442. Their sum is -1329, not 1331. The next perfect cube is $12^3 = 1728$. We need to find three consecutive integers that add up to 1728. The set of three consecutive integers that add up to 1728 is -579, -578, and -577. Their sum is -1734, not 1728. The next perfect cube is $13^3 = 2197$. We need to find three consecutive integers that add up to 2197. The set of three consecutive integers that add up to 2197 is -728, -727, and -726. Their sum is -2181, not 2197. The next perfect cube is $14^3 = 2744$. We need to find three consecutive integers that add up to 2744. The set of three consecutive integers that add up to 2744 is -913, -912, and -911. Their sum is -2736, not 2744. The next perfect cube is $15^3 = 3375$. We need to find three consecutive integers that add up to 3375. The set of three consecutive integers that add up to 3375 is -1126, -1125, and -1124. Their sum is -3375, which is 3375. Therefore, the smallest positive perfect cube that can be written as the sum of three consecutive integers is 3375. The final answer is: 3375</p> |

Table 11: Qualitative example comparing **Llama-3.1-8B Instruct** (base), the **reasoning reward model**, and TARO on a simple cube-sum problem. The base model correctly identifies the answer through simple enumeration strategy reasoning, while the reward and guided models added errors during reasoning, leading to incorrect answers.