# Tool Verification for Test-Time Reinforcement Learning

**Ruotong Liao**[1,3,*] **Nikolai Röhrich**[1,*] **Xiaohan Wang**[2,†] **Yuhui Zhang**[2]
**Yasaman Samadzadeh**[1] **Volker Tresp**[1,3] **Serena Yeung-Levy**[2,†]
[1] Ludwig-Maximilians University of Munich [2] Stanford University [3] MCML

## Abstract

Test-time reinforcement learning (TTRL) has emerged as a promising paradigm for Recursive Self-Improving AI (RSI) by adapting Large Reasoning Models (LRMs) on unlabeled test inputs, using self-consensus rewards derived from majority voting over sampled rollouts. However, majority voting can mistake popularity for correctness: a spurious yet high-frequency *unverified consensus* may become a biased reward signal, causing test-time RL to reinforce frequent but wrong answers and collapse into an incorrect mode. We address this *false-popular* failure mode with **T$^3$RL** (**T**ool Verification for **T**est-**T**ime **R**einforcement **L**earning), a verification-aware test-time RL framework. **T$^3$RL** grounds pseudo-label construction in external tool evidence. Concretely, a verifier utilizes external tool evidence (e.g., from code execution) to upweight verified rollouts during verification-aware voting, producing more reliable pseudo-labels for training. Across math benchmarks of varying difficulty (MATH-500, AMC, and AIME 2024) and diverse backbone families, **T$^3$RL** significantly improves over TTRL, with larger gains on harder problems. **T$^3$RL** is a *verified online data synthesizer* that amortizes tool verification into the policy weights during online adaptation and forms a self-distillation loop. **T$^3$RL** offers both training and inference efficiency and extends to other verifiable domains.

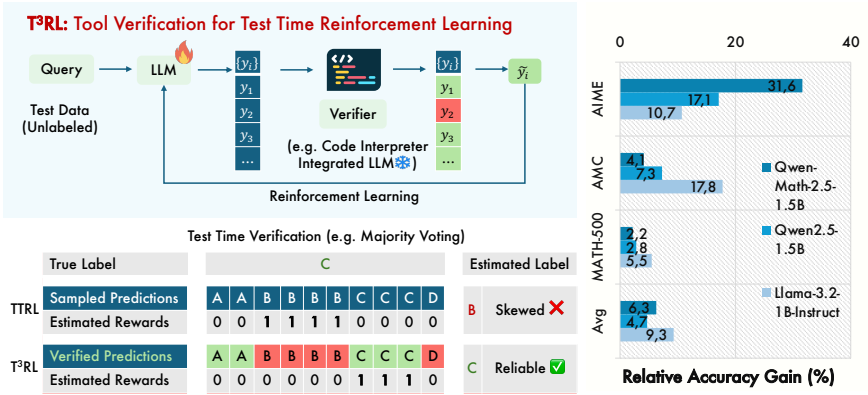



Figure 1: The concept of **T$^3$RL**. Top: **T$^3$RL** introduces test-time verification into self-evolvement via tool-executed evidence (e.g., code interpreter) to stabilize training with verified rollouts. Bottom: Majority-vote pseudo-labels can be spurious. **T$^3$RL** introduces verification to suppress false-popular pseudo-labels. Right: **T$^3$RL** achieves consistent gains with evidence-grounded test-time training.

*Equal contribution.
†Correspondence: liao@dbs.ifi.lmu.de; xhanwang@stanford.edu.
Code: https://github.com/mayhugotong/T3RL

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

# 1 Introduction

Modern language models often know more than they reliably say [Orgad et al., 2025, Wang et al., 2023]. Pre-training and post-training endow models with broad factual and reasoning capabilities, yet a single inference-time generation may fail to express them. This gap has motivated *test-time scaling* [Snell et al., 2024b, Muennighoff et al., 2025]: improving model performance by allocating additional computation at inference through sampling, search, verification, or aggregation [Snell et al., 2024a, Wu et al., 2025, Chang et al., 2025, Venktesh et al., 2025]. A key observation behind these methods is that, even when one rollout is unreliable, the distribution over many rollouts may contain a useful signal [Kang et al., 2026, Mekala et al., 2024, Lifshitz et al., 2025a].

*Test-Time Reinforcement Learning* (TTRL) turns this signal into online learning [Zuo et al., 2025] for Recursive Self-Improving AI (RSI). Given an unlabeled test input, a model samples multiple reasoning traces, uses majority voting to infer a pseudo-label, and reinforces rollouts that agree with this majority answer. When the correct answer is the dominant mode of the model distribution, this provides a simple label-free reward signal that sharpens the model at test time.

However, the majority answer is not always correct. If the model has a systematic bias, many rollouts may converge to the same wrong solution. In this case, voting selects popularity rather than truth: incorrect rollouts receive positive reward, correct minority rollouts are suppressed, and test-time RL amplifies the error. We call this failure mode **false-popular mode collapse**. As shown in Figure 1, a frequent wrong answer can dominate the vote over a less frequent correct one, causing self-training to reinforce a false mode. Worse, once mode collapse occurs, it becomes even harder to self-correct.

This failure arises because self-consistency forms a closed loop: the model generates candidates, votes on them, and learns from its own vote. To break this loop, test-time learning needs evidence beyond the model's internal distribution. Many reasoning tasks provide exactly such evidence. Math solutions can be checked by computation, code can be executed against tests, and structured reasoning can be validated by external tools. These tasks are hard to solve but easy to verify.

We therefore ask: **Can label-free test-time reinforcement learning be made robust to false-popular mode collapse?** To answer this question, we propose **T$^3$RL**, **Tool Verification for Test-Time Reinforcement Learning**. Rather than relying on majority vote alone, **T$^3$RL** uses tools to verify candidate rollouts and constructs pseudo-labels through *verification-weighted voting*. In short, the model should not only ask which answer is most frequent, but which answers can be verified.

Concretely, **T$^3$RL** adds three components to TTRL. (1) A *Verifier* extracts the final answer from each rollout and translates checkable reasoning steps into executable code. (2) A *Verification Tool*, such as a Python interpreter, executes the code and returns evidence about the rollout's validity. (3) A *Verification Weight* then gives verified rollouts greater vote mass when constructing the pseudo-label. The reward remains label-free, but is grounded in external tool feedback rather than raw self-consensus.

We evaluate **T$^3$RL** on MATH500, AMC, and AIME2024. **T$^3$RL** consistently improves over majority-vote TTRL, with larger gains on harder benchmarks and a maximum relative improvement of **31.6%** on AIME2024 (§5.2). We point out that **T$^3$RL** is implicitly *a verified synthetic data generator on the fly* (§6.1). Extensive ablations validate the importance of the *Verifier*, the *Verification Tool*, and the *Verification Weight* (§5.3). We also demonstrate that **T$^3$RL** is more robust (Figure 8b), test-time compute-efficient (Table 4) with half training cost and less than 1% inference cost, and can be improved by stronger verifiers (Figure 6).

More broadly, **T$^3$RL** offers a plug-and-play framework for reliable test-time learning on verifiable tasks. While we instantiate it primarily with code-based verification for mathematical reasoning, our non-math Knights-and-Knaves experiment demonstrates that the same principle extends to symbolic verification with logical puzzles using SMT-LIB2 and Z3 (§6.2). We envision further extensions to theorem checkers, simulators, retrieval systems, and other task-specific verification tools.

# 2 Related Works

**Verification for Test Time Scaling** Verification in test time scaling uses external verifiers to evaluate the quality of additional computation and select the best output from multiple candidates during inference. Verification mechanisms include reward models [Uesato et al., 2022, Lightman et al.,

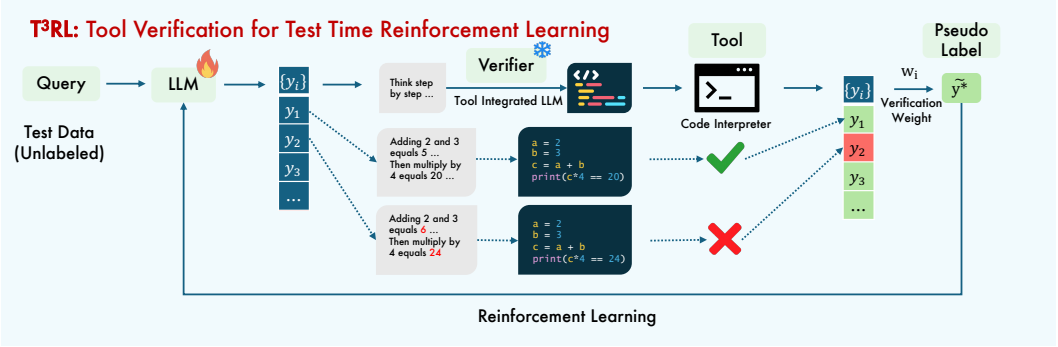



Figure 2: **T³RL: Tool Verification for Test-Time Reinforcement Learning.** *Verifier:* an LLM verifier parses each sampled rollout $y_i$ into a candidate answer $a_i$ and, with tool feedback, produces a binary validity flag $v_i \in \{0, 1\}$. *Tool verification:* the verifier compiles the rollout's claimed computations into lightweight Python and queries a code interpreter to obtain executable evidence. *Verification-weighted majority voting:* unverified rollouts receive unit vote mass, and verified rollouts receive $\omega$ vote mass. A verification-aware pseudo-label $\tilde{y}^*$ is formed by weighted voting.

2023, Cobbe et al., 2021], generative verifiers [Zhang et al., 2024], symbolic checks [Ling et al., 2023], or multi-agent systems [Jin et al., 2025, Lifshitz et al., 2025b]. Recently, tool-integrated reasoning (TIR) [Gou et al., 2024] has brought a new perspective to tool-integrated verification [Mekala et al., 2024, Kang et al., 2026], formalizing tool use as an additional robust evidence. However, none of the prior works explore verification in *test-time training*; we are the first to turn sampled rollouts into *online, evidence-labeled* training instances through verification and shape the training process as a verified data synthesizer on-the-fly.

**Test-Time Training** Test Time Training (TTT) adapts model parameters during inference to handle distribution shifts on new tasks [Sun et al., 2020, 2024b,a, Behrouz et al., 2024, Liu et al., 2021], in domains like video generation and understanding [Wang et al., 2025a, Dalal et al., 2025], large language models [Hardt and Sun, 2023], or scientific discovery [Yuksekgonul et al., 2026]. Recently, TTT advanced to test-time reinforcement learning (TTRL) [Zuo et al., 2025] that combines unsupervised reinforcement learning [Prasad et al., 2024, Zhang et al., 2025] and reinforcement learning with verifiable rewards (RLVR) [Zeng et al., 2025, Wang et al., 2025b], which has been widely discussed in the era of self-evolving artificial intelligence. However, in the face of the challenge that spurious reward estimation presents to test-time training, none of the existing work has discussed verification. To the best of our knowledge, we are the first to propose test-time verification for test-time training, especially *tool-assisted verification* for evidence-grounded test-time reinforcement learning. This provides a broader vision of agentic systems that increasingly rely on tool interaction, yet require reliable reward signals to for stable online learning.

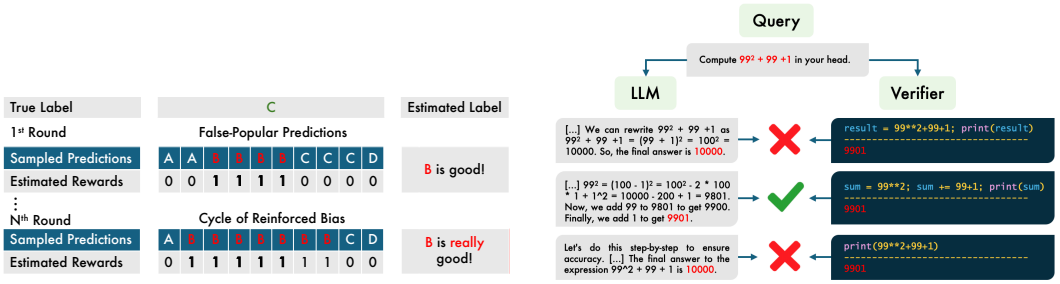
## 3 The Failure Mode: How Unverified Consensus Induces Reward Bias

### 3.1 Test Time Reinforcement Learning

Unlike traditional RL, where models learn from known reward signals, TTRL operates on unlabeled test data without access to explicit supervision. TTRL is defined as follows:

Given a state represented by the prompt $x$, the model acts by producing an output $y$ sampled from a policy $\pi_\theta(y \mid x)$, parameterized by $\theta$. To construct a reward signal without ground-truth labels, TTRL generates multiple candidate outputs $\{y_1, y_2, \ldots, y_N\}$ through repeated sampling. A consensus output $y^*$ is derived by *majority voting*, serving as a proxy for the optimal estimation. The environment then provides a reward $r(y, y^*)$ based on the alignment between the sampled estimation $y$ and the consensus estimation $y^*$. The RL objective is thus to maximize the expected reward:

$$\max_\theta \mathbb{E}_{y \sim \pi_\theta(\cdot|x)}[r(y, y^*)], \tag{1}$$




(a) Failure case of TTRL.

(b) Success case of **T³RL**.

Figure 3: **Unverified consensus versus tool-verified consensus.** (a) Spurious reward in the reinforced cycle of TTRL: the false-popular answer is rewarded and its vote share grows over rounds. (b) Tool verification can adjust the false-popular answer.

and parameters $\theta$ are updated through gradient ascent:

$$\theta \leftarrow \theta + \eta \nabla_\theta \mathbb{E}_{y \sim \pi_\theta(\cdot|x)}[r(y, y^*)], \tag{2}$$

where $\eta$ denotes the learning rate.

### 3.2 Spurious Majority as a Biased Pseudo-Label

**Self-consensus can estimate wrong labels.** Let the generator induce a distribution over final answers, and consider two competing modes: the correct answer $y^\star$ and an incorrect but high-frequency answer $\tilde{y}$, e.g. `B` and `C` in Figure 3a respectively. If the win probability of `B` over `C` is non-zero, the majority vote can select $\tilde{y}$ as the pseudo-label. We measure the fraction of questions for which the unverified majority answer is wrong (Figure 4) and the fraction for which verification-weighted voting changes that answer (Figure 3b). The false-majority share rises sharply with difficulty.

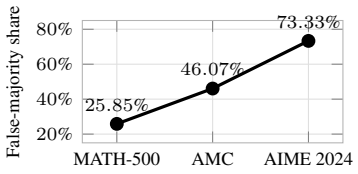



Figure 4: **False-majority share** across benchmarks.

**Self-reinforcing feedback loop and incorrect mode collapse.** Once the pseudo-label is set to the false-popular $\tilde{y}$, the majority-based reward assigns positive reinforcement to rollouts that agree with the false signal $\tilde{y}$ and zero out rewards for truthful rollouts. The RL update increases the likelihood of sampling $\tilde{y}$ in subsequent rollouts, which further increases its vote share, making the pseudo-label even more confidently wrong. This self-reinforcing dynamic can drive TTRL toward *incorrect mode collapse*. Furthermore, once mode collapse occurs, the model becomes even harder to self-correct in an echo chamber.

## 4 Method: Tool Verification for Test Time Reinforcement Learning

To prevent this failure mode, we present **T³RL**, an RL framework which integrates tool verification into the aggregation mechanism of test-time RL, thus achieving grounded and more robust reward estimation. Specifically, **T³RL** has three core components: (1) An external *Verifier*, an LLM that is tasked with verifying a given reasoning trace by compiling the trace into executable Python code, and judging its validity based on the execution output (§4.1); (2) The *Verification Tool*, a code interpreter that executes the generated Python program and returns signals to the verifier (§4.2); (3) The *Verification Weight*, which is used to replace the majority vote with a verification-aware weighted vote that boosts the voting power of verified rollouts. (§4.3).

### 4.1 Verifier

A *verifier* $\mathcal{V}$ is designed to evaluate each rollout before aggregation derives the estimated consensus label $y^*$:

**Verifier.** Given an input prompt $x$, we sample $N$ rollouts $\{y_i\}_{i=1}^{N} \sim \pi_\theta(\cdot \mid x)$ and extract the rollout's own final answer from each rollout,

$$\hat{a}_i = \mathrm{Extract}(y_i), \qquad \hat{a}_i \in \mathcal{A} \tag{3}$$

$\mathcal{V}$ then evaluates each rollout and returns only a validity flag

$$(v_i) \;=\; \mathcal{V}(x, y_i, t_i), \tag{4}$$

where $v_i \in \{0, 1\}$ indicates whether the rollout passes executable checks. $t_i$ represents the tool output from the verification tool $\mathcal{T}$ used by the verifier.

In **T$^3$RL**, $\mathcal{V}$ is implemented as an LLM-based verifier that performs the following tasks: 1) generating the tool-calling query, by transforming the rollout rationale into a lightweight python program; 2) the execution of the *verification tool* (§4.2), and 3) returning the verification validity results. Note that the verifier provides a reliability signal for $\hat{a}_i$.

### 4.2 Verification Tool

A *verification tool* $\mathcal{T}$ provides external, deterministic, and executable evidence to assist the verifier $\mathcal{V}$ by offloading specific verification tasks from $\mathcal{V}$ to an external tool.

**Tool execution as external evidence.** For math problems, many rollouts hinge on intermediate correctness (e.g., arithmetic computations, calculation hallucinations, etc). In **T$^3$RL**, *tool verification* is initiated as a code interpreter $\mathcal{T}$, that is an executable checker that evaluates a verifier-generated Python program and returns the output of the executed code

$$t_i \;=\; \mathcal{T}(\mathrm{Code}(x, y_i)). \tag{5}$$

The verifier then *contrasts* the tool result with the rollout's extracted answer $\hat{a}_i$ and produces a tool-verified validity indicator

$$v_i \;=\; \mathbb{1}[t_i = \hat{a}_i]\,. \tag{6}$$

### 4.3 Verification Weight

**Heuristic Weighting.** In standard TTRL, the consensus label is derived via simple majority voting, where every reasoning path contributes equally regardless of its logical soundness. However, in our framework, rollouts that pass the tool verification, i.e. for which $v_i = 1$, are heuristically assumed to be more reliable than those that do not. Since the ground truth is inaccessible at test time, we cannot calculate the exact probability of correctness for verified rollouts. Instead, we introduce a *Verification Weight* hyperparameter, $\omega$, to quantify the voting power of tool-verified traces.

**Verification-Aware Consensus.** We modify the aggregation mechanism to use a verification-weighted majority vote. To maintain the diversity of rollouts, we ensure that each rollout maintains its contribution by setting the default weight to 1. For a set of rollouts, we assign a weight $w_i$ to each derived answer $\hat{a}_i$ based on the verification indicator $v_i$:

$$w_i \;=\; (1 - v_i) \cdot 1 \;+\; v_i \cdot \omega, \tag{7}$$

where $\omega \geq 1$ is a fixed scalar hyperparameter. This ensures that unverified rollouts retain a unit vote, while each verified rollout contributes an upweight of $\omega$ votes. The *verification-aware* consensus label $\tilde{y}^*$ is then obtained by maximizing the total weighted vote mass:

$$\tilde{y}^* \;=\; \arg\max_{a \in \mathcal{A}} \sum_{i=1}^{N} w_i \cdot \mathbb{1}[\hat{a}_i = a]. \tag{8}$$

This mechanism allows **T$^3$RL** to shift the consensus derivation from the distribution of the most *frequent* answer in the whole observed candidate set $\mathcal{A}$ towards the *verified* consensus (see Figure 3b), provided that the verified group possesses sufficient cumulative weight.

**Reward Calculation.** Consistent with standard TTRL, the final reward signal remains binary but is now anchored to the robust, verification-aware consensus $\tilde{y}^*$. The reward for the $i$-th rollout is computed as:

$$r_i^{\mathrm{v}} \;=\; \mathbb{1}[\hat{a}_i = \tilde{y}^*]. \tag{9}$$

Table 1: **Main performance of T$^3$RL.** (a) Improvement across benchmark and backbone models. (b) Difficulty breakdown performance on MATH-500 (L5 is the hardest). The relative improvement of **T$^3$RL** over TTRL increases with difficulty level (Figure 5).

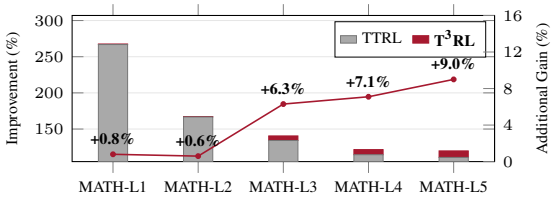
(a) Main results

| Model / Method | AIME 2024 | AMC | MATH-500 | Avg |
|---|---|---|---|---|
| **Math Model** | | | | |
| Qwen-2.5-Math-1.5B | 7.7 ± 2.0 | 28.6 ± 4.5 | 32.7 ± 2.7 | 23.0 ± 1.9 |
| w/ TTRL | 15.8 ± 2.1 | 48.9 ± 1.6 | 73.0 ± 3.8 | 45.9 ± 1.5 |
| **w/ T$^3$RL** | **20.8 ± 1.6** | **50.9 ± 2.7** | **74.6 ± 3.2** | **48.8 ± 1.5** |
| **Vanilla Model** | | | | |
| Qwen-2.5-1.5B | 0.2 ± 1.2 | 0.6 ± 2.2 | 7.7 ± 2.0 | 2.8 ± 1.1 |
| w/ TTRL | 3.5 ± 1.2 | 28.6 ± 4.5 | 63.2 ± 3.5 | 31.8 ± 1.9 |
| **w/ T$^3$RL** | **4.1 ± 1.2** | **30.7 ± 3.7** | **65.0 ± 2.1** | **33.3 ± 1.5** |
| Qwen-3-4B | 0.0 ± 0.0 | 12.5 ± 4.4 | 50.4 ± 3.5 | 21.0 ± 1.9 |
| w/ TTRL | 36.4 ± 1.9 | 71.7 ± 3.0 | 88.4 ± 1.4 | 65.5 ± 1.3 |
| **w/ T$^3$RL** | **40.0 ± 1.6** | **74.2 ± 2.8** | **89.5 ± 1.4** | **67.9 ± 1.2** |
| **Instruct Model** | | | | |
| Llama-3.2-1B-Instruct | 0.8 ± 0.0 | 4.2 ± 2.6 | 4.4 ± 1.5 | 3.1 ± 1.0 |
| w/ TTRL | 7.5 ± 2.5 | 16.9 ± 1.6 | 40.0 ± 1.6 | 21.5 ± 1.1 |
| **w/ T$^3$RL** | **8.3 ± 2.4** | **19.9 ± 0.9** | **42.2 ± 2.6** | **23.5 ± 1.2** |
| Llama-3.2-3B-Instruct | 6.0 ± 1.4 | 19.4 ± 1.8 | 43.9 ± 2.1 | 23.1 ± 1.0 |
| w/ TTRL | 13.3 ± 3.8 | 31.3 ± 4.0 | 61.6 ± 5.1 | 35.4 ± 2.5 |
| **w/ T$^3$RL** | **17.1 ± 1.2** | **34.2 ± 2.7** | **63.3 ± 3.3** | **38.2 ± 1.5** |

(b) MATH-500 difficulty breakdown.

| | L1 | L2 | L3 | L4 | L5 |
|---|---|---|---|---|---|
| Qwen-2.5-Math-1.5B | 25.9 | 33.0 | 36.3 | 32.5 | 22.3 |
| w/ TTRL | 95.0 | 88.0 | 85.0 | 69.7 | 47.0 |
| **w/ T$^3$RL** | **95.2** | **88.2** | **87.3** | **72.0** | **49.0** |
| Rel. (% over Baseline) | ↑267.6% | ↑167.3% | ↑140.5% | ↑121.5% | ↑119.7% |
| Rel. (% over TTRL) | **↑0.2%** | **↑0.2%** | **↑2.7%** | **↑3.3%** | **↑4.3%** |

Figure 5: Relative gains across MATH-500 levels.




**Training objective.** We keep the TTRL update of Eq. (2) unchanged in form, replacing the majority-vote pseudo-label $y^*$ with the verification-aware consensus $\tilde{y}^*$ of Eq. (8) and its induced reward, Eq. (9):

$$\max_{\theta} \mathbb{E}_{y\sim\pi_\theta(\cdot|x)}[r^{\mathrm{v}}(y,\tilde{y}^*)]\,, \qquad \theta \leftarrow \theta + \eta\nabla_\theta \mathbb{E}_{y\sim\pi_\theta(\cdot|x)}[r^{\mathrm{v}}(y,\tilde{y}^*)]\,, \tag{10}$$

where $\tilde{y}^*$ is held fixed within each update, and we optimize this objective with GRPO in practice.

# 5 Experiments

## 5.1 Experimental Setup

**Benchmarks** We evaluate **T$^3$RL** on 3 mathematical reasoning benchmarks: AIME 2024 [Li et al., 2024], AMC [Li et al., 2024], and MATH-500 [Hendrycks et al., 2021].

**Models.** We evaluate **T$^3$RL** under the *test-time reinforcement learning* setting across diverse backbone configurations spanning model families, alignment regimes, and scales (1B–4B), which appeals to the call to validate methods on diverse models [Shao et al., 2025]. Our experiments cover both *base* and *instruction-tuned* models, as well as *math-specialized* backbones. Concretely, we study LRMs from the Qwen and LLaMA families: (1) **Vanilla:** `Qwen2.5-1.5B` [Qwen and el al, 2024] and `Qwen-3-4B` [Yang et al., 2025]; (2) **Math:** `Qwen2.5-Math-1.5B` [Qwen and el al, 2024]; (3) **Instruct:** `Llama-3.2-1B-Instruct` and `Llama-3.2-3B-Instruct` [Grattafiori and el al, 2024].

**Evaluation Setup** We apply **T$^3$RL** to each benchmark utilizing a maximum token limit of 2,560. For our *main experiments*, we follow the DeepSeek-R1 [Guo et al., 2025] protocol: we evaluate each experiment 4 times (temperature 0.6, top-$p$ 0.95) and report pass@1 as the average correctness, formalized as: For AIME 2024, whose test set is small, we average 16 independent evaluations; for AMC and MATH-500 we average 4. The reported ± values are standard deviations aross runs.

$$\text{pass@1} = \frac{1}{k}\sum_{i=1}^{k} p_i,$$

where $p_i$ indicates if response $i$ is correct.

**Baselines** We compare our method against base models, as well as against TTRL [Zuo et al., 2025], where we report the original results where possible. For experiments where TTRL results are not available, we reproduce training runs using the implementation details in Zuo et al. [2025].

**Implementation Details** We implement **T$^3$RL** using GRPO [Shao et al., 2024] on each benchmark, training the policy model via AdamW with a cosine schedule (peak learning rate: $5 \times 10^{-7}$). To balance performance with computational efficiency, we follow TTRL [Zuo et al., 2025] in employing a vote-then-sample strategy: 64 responses are initially generated for label estimation (temperature 0.6) and downsampled to 32 for training. The maximum token length is 2,560. Regarding the verifier model, we find that we can restrict the generation length to $1{,}024$ tokens, since code snippets are relatively short. We also use a temperature of $0.6$ for the verifier across all benchmarks. Training runs are conducted on 8 NVIDIA A100 GPUs, with varying episode counts based on dataset size: 10 (MATH-500), 30 (AMC), and 80 (AIME 2024). The pseudo code is listed in Appendix A and the system prompt is in Appendix B.

### 5.2 Main Results

Table 1 shows that **T$^3$RL** repeatedly outperforms TTRL across all evaluated models and benchmarks, supporting our central hypothesis that verification-shaped majority voting mitigates unverified-consensus bias during test-time RL. Across the 15 paired **T$^3$RL**–TTRL model–benchmark comparisons, **T$^3$RL** wins all 15.

**Consistent gains across benchmarks and model types.** **T$^3$RL** improves over TTRL on all three evaluated benchmarks, spanning an easy-to-hard spectrum from MATH-500 (easiest), to AMC (medium), to AIME 2024 (hardest). **T$^3$RL** consistently improves over TTRL across five backbones spanning three regimes (math, vanilla, instruction-tuned) and two scales (1B–4B): the largest gains are on AIME 2024 for `Qwen-Math-1.5B` ($+31.6\%$) and `Qwen2.5-1.5B` ($+17.1\%$), and on AMC for `Llama-3.2-1B-Instruct` ($+17.8\%$). The benefit is tied with three standing out trends:

**Math-specialised backbones, harder benchmarks, and larger models benefit most.**

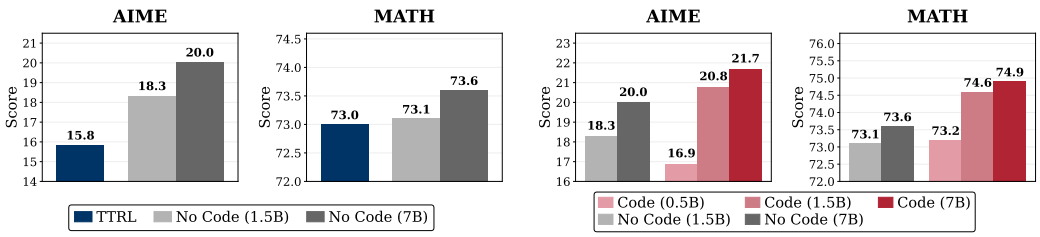
- *(i) Math-specialised models benefit more.* `Qwen-Math-1.5B` achieves a larger average gain over TTRL than the vanilla `Qwen2.5-1.5B` ($+6.3\%$ vs. $+4.7\%$): math backbones generate longer calculation chains whose small arithmetic slips are precisely the kind of error that executable checks can catch.
- *(ii) Harder benchmarks benefit more.* Within each backbone, gains peak on AIME 2024 (e.g., $+31.6\%$ for `Qwen-Math-1.5B`) and on the hardest MATH-500 level L5: longer computation chains accumulate errors that deterministic tool execution can still verify, preserving signal where self-consensus degrades.
- *(iii) Gains scale to larger backbones.* On `Llama-3.2-3B-Instruct` AIME rises from 13.3 to 17.1 ($+28.6\%$); on the much stronger `Qwen-3-4B` where TTRL already achieves 36.4, it reaches 40.0 ($+9.9\%$). The gain is largely preserved under self-verification ($+9.3\%$ / $+27.1\%$, Table 2), so the contribution scales with backbone capability rather than coming from external-coder distillation.

### 5.3 Ablation Studies

This section presents an analysis of the three key factors that support **T$^3$RL**: (1) test-time verification, (2) tool-assisted verification, and (3) verification-weighted majority voting. Unless otherwise specified, we run ablations on `Qwen-Math-1.5B` with rollout size $N{=}64$ and report Pass@1 performance.

**The Contribution of Test-Time Verification.** We isolate the effect of *test-time verification* by comparing (i) vanilla TTRL against (ii) **T$^3$RL** *without* tool access, i.e., using LLM alone for verification. In Figure 6a, improvements are seen on both AIME and MATH with 1.5B or 7B verifier on 1.5B policy model. These results showcase that introducing verification on rollouts yields better online updates with improved reward estimation, even before introducing tool-assisted verification.

**The Contribution of Tool-Assisted Verification.** We next control for the *tool execution* effect by comparing verification **T$^3$RL** *with* and *without* tool execution. The results in Figure 6b show that the verification tool contributes a clear additional lift beyond verifier-only checking, improving AIME from 18.3 to 20.8 for the 1.5B verifier and from 20.0 to 21.7 for the 7B, confirming that executable evidence reduces verifier uncertainty and yields a more reliable verification signal.




(a) Ablating verification.

(b) Ablating tool execution and verifier size.

Figure 6: **Ablation on *verifier* and *verification tool*.** Left: Adding an LLM verifier improves TTRL even without tool execution. Right: Code execution significantly strengthens verification.

**Self-Verifier vs. External Verifier.** We further ablate the distillation effect by comparing **T$^3$RL** with the policy model as self-verifier on AIME. Note that, the external verifier `Qwen-2.5-Coder-1.5B` is an even weaker model (1.7 on AIME). Table 2 shows that self-verification beats TTRL on all three backbones. So the gain stems from *verification itself* instead of distilling math knowledge from other model. Verification leads to flip rates on 5.40% of AIME 2024, 10.75% of AMC, and 12.87% of MATH-500.

Table 2: Self vs. Ext. verifier.

| | Math 1.5B | Llama 3B | Qwen3 4B |
|---|---|---|---|
| TTRL | 15.8 | 13.3 | 36.4 |
| Self | 16.7 | 16.9 | 39.8 |
| Ext. | **20.8** | **17.1** | **40.0** |

**The Contribution of Verification-Weighted Voting.** We ablate the *verification weight* $\omega$ in **T$^3$RL** by sweeping its value, which controls the voting power of rollouts that pass the tool-verifier check. A weight of $\omega$=1 degenerates to majority-vote TTRL, and $\omega \to \infty$ approximates *binary hard filtering* of unverified rollouts. As shown in Figure 7, finite weights $\omega > 1$ improve AIME accuracy across all three backbones. The largest gain occurs for `Qwen-2.5-Math-1.5B` at $\omega$=5 (AIME +31.6%, MATH-500 +2.2%), with similar performance at $\omega$=10; hard filtering is less reliable. This behavior arises because the verifier is itself noisy due to the stochastic nature of LLMs in rollout-to-code translation, although tool execution provides a deterministic check: Table 3 shows that the verifier accuracy with `Coder-1.5B` reaches only 73.2%, so hard filtering can propagate false positives and discard true positives, while $\omega$=1 ignores the verifier entirely. Therefore, tool-assisted verification should act as a ***soft*** preference signal.

Table 3: **Verifier quality on policy rollouts.**

| Policy | Verifier | Acc. | Prec. | Rec. | Δ AIME |
|---|---|---|---|---|---|
| `Math-1.5B` | `Coder-1.5B` | 73.2 | 25.0 | 51.5 | +5.0 |
| `Llama-3.2-3B` | `Coder-1.5B` | 73.3 | 39.4 | 59.1 | +3.8 |
| `Qwen-3-4B` | `Coder-1.5B` | 75.0 | 38.2 | 40.7 | +3.6 |

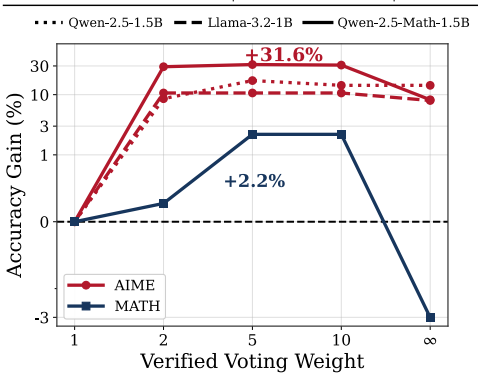



Figure 7: The effect of the vote weight $\omega$.

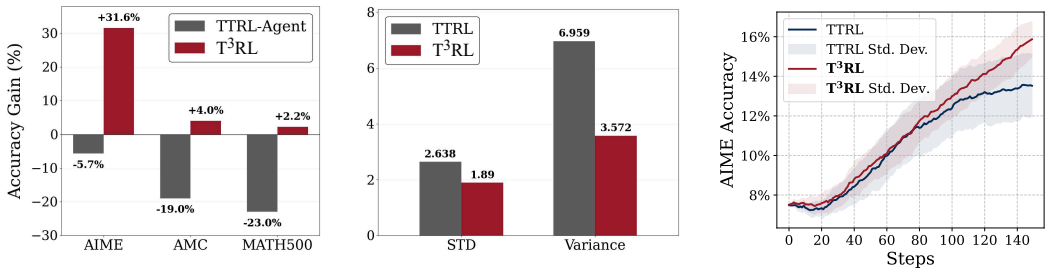
*Sensitivity of the verification weight:* Every $\omega$>1 improves over TTRL generally, and moderate values plateau ($\omega$=2 is within 0.4 points of $\omega$=5); only hard filtering degrades.

## 6 Discussions and Analysis

### 6.1 Q1: Why Does T$^3$RL Work?

> ***T$^3$RL** functions as a verified self-distillation loop that generates synthetic training data online.*

Tool-use of reasoning models is learned from execution feedback and improved through interaction with the environment [Silver and Sutton, 2025]. In the test-time training loop, there's an open design choice: *tools as policy actions* or *tools as verification evidence.* The impressive contribution of tool verification in **T$^3$RL** raises the question of whether tool access alone suffices, or whether the position in the loop also matters.




(a) Tool in policy vs. in verification. (b) Training robustness. (c) Training dynamics comparison.

Figure 8: **Ablations on tool position and training robustness.** Left: allowing the *policy* to call tools during rollouts (TTRL-AGENT) can hurt performance, while **T³RL** improves by restricting tool use to the *verifier* for reward shaping. Right: verifier-shaped rewards reduce run-to-run variability, indicating more stable optimization under unlabeled test-time RL.

*Setting.* To further examine why **T³RL** works, we compare (i) TTRL; (ii) TTRL-AGENT, which extends TTRL and grants the trained policy access to tool calling directly with majority voting over execution results, and (iii) **T³RL**, which trains a non-agentic policy with a tool-assisted *verifier*.

*Observation & analysis.* Figure 8a shows TTRL-AGENT degrades, while **T³RL** yields growing positive gains with benchmark difficulty. We attribute TTRL-AGENT's failure to *error-signal mixture*: tool calls inside rollouts conflate reasoning errors with tool-usage errors (e.g., malformed code generation or brittle execution artifacts), and self-consensus rewards amplify this noise. In contrast, **T³RL** *decouples* reasoning from tool execution: the verifier provides evidence *after* generation and converts it into a verification-informed reward, turning rollouts into verified labeled training data on the fly: an implicit *verified online data synthesiser* that stabilizes test-time learning.

**T³RL is more robust under tool-assisted verification voting.** Tool verification further stabilizes training, as shown in Figure 8c. Across multiple runs, TTRL exhibits noticeable run-to-run variability, consistent with the sensitivity of self-induced rewards to sampling noise and pseudo-label estimation. In contrast, **T³RL** anchors reward construction with tool verification, yielding substantially lower dispersion in peak performance in Figure 8b: both the standard deviation and variance of the best accuracy after 100 steps. Overall, **T³RL** produces more consistent learning dynamics and more reliable accuracy on AIME.

Table 4: **Training computation.** AIME runs; arrows are relative to TTRL@64. **T³RL** achieves higher accuracy with less than half of the training cost, time and computations.

| Method (rollout) | AIME ↑ | Tokens ↓ | Time (s) ↓ | TFLOPs ↓ |
|---|---|---|---|---|
| TTRL@64 | 15.8 | 365K | 324 | 7.9K |
| **T³RL**@16 | **20.0** (↑27%) | **147K** (↓60%) | **126** (↓61%) | **3.0K** (↓63%) |
| **T³RL**@32 | 20.8 (↑32%) | 293K (↓20%) | 218 (↓33%) | 5.9K (↓25%) |
| **T³RL**@64 | 20.8 (↑32%) | 586K (↑61%) | 406 (↑26%) | 11.8K (↑50%) |

Table 5: **Inference computation.** Per-query costs; arrows are relative to **T³RL**.

| Method | Inf. Tokens ↓ | AIME | AMC | MATH |
|---|---|---|---|---|
| Base | 1,630 | 7.7 (↓63%) | 28.6 (↓44%) | 32.7 (↓56%) |
| TTRL | 1,630 | 15.8 (↓24%) | 48.9 (↓4%) | 73.0 (↓2%) |
| **T³RL (ours)** | **1,630** | **20.8** | **50.9** | **74.6** |
| Test-time majority ($N$=16) | 61K (↑37×) | 19.2 (↓8%) | 49.4 (↓3%) | 65.0 (↓13%) |
| Tool-verified majority ($N$=16) | 96K (↑59×) | 20.8 (→0%) | 51.5 (↑1%) | 65.5 (↓12%) |
| T1 ($N$=16, $V$=4) | 200K (↑123×) | 20.0 (↓4%) | 52.1 (↑2%) | 66.3 (↓11%) |
| T1 + PRM ($N$=16, $V$=4) | 200K (↑123×) | 20.4 (↓2%) | 52.7 (↑4%) | 67.9 (↓9%) |
| Buffer-of-Thoughts | 2,120 (↑30%) | 9.6 (↓54%) | 34.9 (↓31%) | 38.6 (↓48%) |
| BoN-RR ($N$=64) | 243K (↑148×) | 21.0 (↑1%) | 51.8 (↑2%) | 51.7 (↓31%) |

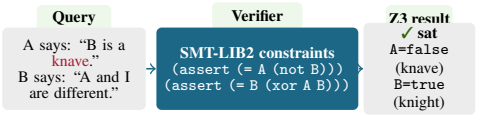
**Allocating test-time computation wisely: verification beats brute-force scaling.** How should the test-time budget be split between (i) *sampling more rollouts* for self-consensus and (ii) *verifying rollouts* to make each sample more informative? We find two things. *First, **T³RL** is cheaper during training (60% less) to achieve better accuracy than* TTRL. Holding all else fixed and accounting for both policy and verifier costs, **T³RL**@$r$=16 already surpasses TTRL@64 using 59.8% fewer tokens, 61.3% less wall-clock time, and 62.5% fewer TFLOPs (Table 4); sandbox execution adds only ∼49 ms per snippet (< 1% of training). *Second, **T³RL** is cheaper during inference (123–200× less) than verifier-based TTS to achieve matched accuracy.* Comparing with non-parametric alternatives (Table 5): Buffer-of-Thoughts [Yang et al., 2024], a memory-based system that reuses successful rollouts, cannot match even TTRL, while the verifier-guided test-time scaling approach, best-of-$N$ reranking (BoN-RR), approaches **T³RL**'s accuracy only by inflating per-query inference tokens by 148× and paying $\mathcal{O}(2N)$ at every test instance. Other verifier-aware TTS baselines are no cheaper: T1 [Kang et al., 2026] reaches 66.3 and T1+PRM [Wang et al., 2024] reaches 67.9 on MATH-500, both at 200K tokens per query. **T³RL** instead amortises the verifier's signal into the policy weights once at training and infers at $\mathcal{O}(1)$, converting test-time compute into efficient parameter updates and avoiding re-paying the inference cost on every query.

### 6.2 Q2: Can T$^3$RL generalize?

**Generalization beyond math: logical reasoning** To test whether **T$^3$RL** depends on mathematical reasoning or Python execution, we evaluate it with Qwen-2.5-Math-1.5B-Instruct on held-out $n$-character puzzles from the Knights-and-Knaves (K&K) dataset [Xie et al., 2024][3], where knights always tell the truth and knaves always lie. The verifier represents knights as `true` and knaves as `false`, translates policy solutions into SMT-LIB2 constraints (Satisfiability Modulo Theories [Barrett and Tinelli, 2018]), and checks them with Z3 [de Moura and Bjørner, 2008].

Table 6: Knights-and-Knaves accuracy at training step 240, with an SMT-LIB2/Z3 example.

| Puzzle | Base | TTRL | T$^3$RL | Δ vs. TTRL |
|---|---|---|---|---|
| $n$=2 | $34.5 \pm 6.5$ | $63.0 \pm 0.0$ | $\mathbf{73.0 \pm 1.2}$ | **+10.0** |
| $n$=3 | $22.2 \pm 3.8$ | $45.0 \pm 2.0$ | $\mathbf{68.0 \pm 0.8}$ | **+23.0** |
| $n$=4 | $11.8 \pm 2.8$ | $12.8 \pm 0.5$ | $\mathbf{19.5 \pm 0.5}$ | **+6.7** |

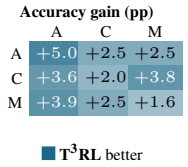



**T$^3$RL** beats TTRL at every puzzle size (Table 6), by 23.0 points at $n$=3 and 6.7 at $n$=4, where TTRL stalls near the base model. This supports the narrower claim that **T$^3$RL** can generalize wherever a tool reliably checks correctness, as in theorem proving and program verification.

**Generalization across-dataset.** Each row adapts the policy at test time on one source benchmark and evaluates it on all three targets, so off-diagonal entries measure transfer to unseen benchmarks. **T$^3$RL** improves all nine train–test pairs (mean +3.0 points; Table 7). Mean run-to-run SD decreases by 8.4% (2.38 → 2.18).

Table 7: Cross-dataset adaptation (A/C/M: AIME/AMC/MATH-500). Inset: accuracy gain.

| Adaptation setting | AIME 2024 | AMC | MATH-500 |
|---|---|---|---|
| TTRL adapted on AIME | $15.8 \pm 2.1$ | $42.9 \pm 2.4$ | $65.5 \pm 2.4$ |
| **T$^3$RL** adapted on AIME | $\mathbf{20.8 \pm 1.6}$ | $45.4 \pm 2.5$ | $68.0 \pm 2.0$ |
| TTRL adapted on AMC | $11.5 \pm 1.6$ | $48.9 \pm 1.6$ | $66.6 \pm 1.4$ |
| **T$^3$RL** adapted on AMC | $\mathbf{15.1 \pm 2.0}$ | $\mathbf{50.9 \pm 2.7}$ | $\mathbf{70.4 \pm 1.7}$ |
| TTRL adapted on MATH-500 | $10.8 \pm 3.4$ | $43.2 \pm 2.7$ | $73.0 \pm 3.8$ |
| **T$^3$RL** adapted on MATH-500 | $\mathbf{14.7 \pm 1.9}$ | $\mathbf{45.7 \pm 2.0}$ | $\mathbf{74.6 \pm 3.2}$ |




### 6.3 Q3: What boosts T$^3$RL, and when might it fail?

**Boosting performance: stronger verifiers and larger rollout budgets help.** Scaling the verifier from 0.5B to 7B on `Qwen-Math-2.5` consistently improves **T$^3$RL** across benchmarks (AIME 20.8 → 21.7, AMC 50.9 → 51.5, MATH-500 74.4 → 74.9; Table 3): stronger verifiers produce more reliable rollout-to-code translation and verification-aware reward signals. Increasing the rollout budget $N \in \{16, 32, 64\}$ also improves **T$^3$RL** monotonically (Table 4), as a larger candidate set raises solution diversity and concentrates vote mass on a more reliable pseudo-label. Figure 3b shows the resulting effect: **T$^3$RL** flips the majority label via tool verification when an incorrect calculation dominates.

**Limitations: noisy verifiers and easy tasks limit gains.** **T$^3$RL** requires $\mathcal{V}$ to provide a meaningful correctness signal. When $\mathcal{V}$ is underpowered (e.g., `Qwen-0.5B` with bare-minimum coding ability), noisy tool calls cause verification-aware voting to misweight rollouts and add a second stochastic layer on top of self-consensus, so the pseudo-label $\widetilde{y}^*$ becomes less stable than vanilla majority voting. Experimental results and failure-case studies are provided in Appendix C. On easy tasks where rollouts are already accurate and consistent, verification aligns with the consensus without changing the pseudo-label distribution, so gains over TTRL are small.

## 7 Conclusion

We propose **T$^3$RL**, introducing *test-time verification* to the test-time reinforcement learning framework that learns from unlabeled test data by suppressing spurious rewards with tool verification. Experiments across heterogeneous backbones and math benchmarks show consistent gains of *tool verification*. Overall, **T$^3$RL** supports test-time RL as *verified online data synthesis*: sampled rollouts become reliable training instances once verified with executable evidence, enabling more stable verifiable test-time training. Future work will broaden to domain-specific tool verification.

[3] `https://huggingface.co/datasets/K-and-K/knights-and-knaves`

## Impact Statement

This paper aims to advance the field of machine learning. Our work introduces test-time verification as a mechanism for stabilizing self-evolution in large reasoning models by reducing error reinforcement and mitigating self-consistency-driven failure modes. If deployed responsibly, such verification can improve reliability and robustness in high-stakes applications by encouraging models to seek and check external evidence rather than relying solely on internal consistency. Potential risks include over-reliance on imperfect verifiers and the possibility that verification pipelines inherit biases or vulnerabilities from their underlying tools and data sources. To mitigate these risks, future work should investigate improvements to verifiers and robustness to adversarial or noisy feedback. More broadly, the proposed framework is modular and can incorporate improved test-time verification methods as they become available.

## Acknowledgments

This work was funded by the German Federal Ministry of Research, Technology and Space (BMFTR) under grant number 16IS24029 SWC-MuKChat and by the Ludwig-Maximilians-University of Munich, and supported by the Munich Center for Machine Learning (MCML). We thank Lambda's academic program for compute grants. The authors declare no competing interests.

## References

Clark Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of Model Checking*, pages 305–343. Springer International Publishing, Cham, 2018.

Ali Behrouz, Peilin Zhong, and Vahab Mirrokni. Titans: Learning to memorize at test time, 2024. URL https://arxiv.org/abs/2501.00663.

Kaiyan Chang, Yonghao Shi, Chenglong Wang, Hang Zhou, Chi Hu, Xiaoqian Liu, Yingfeng Luo, Yuan Ge, Tong Xiao, and Jingbo Zhu. Step-level verifier-guided hybrid test-time scaling for large language models, 2025. URL https://arxiv.org/abs/2507.15512.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Karan Dalal, Daniel Koceja, Jiarui Xu, Yue Zhao, Shihao Han, Ka Chun Cheung, Jan Kautz, Yejin Choi, Yu Sun, and Xiaolong Wang. One-minute video generation with test-time training. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17702–17711, 2025.

Leonardo de Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340, 2008.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujiu Yang, Minlie Huang, Nan Duan, and Weizhu Chen. Tora: A tool-integrated reasoning agent for mathematical problem solving, 2024. URL https://arxiv.org/abs/2309.17452.

Aaron Grattafiori and el al. The llama 3 herd of models, 2024. URL https://arxiv.org/abs/2407.21783.

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

Moritz Hardt and Yu Sun. Test-time training on nearest neighbors for large language models. *arXiv preprint arXiv:2305.18466*, 2023.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021.

Can Jin, Hongwu Peng, Qixin Zhang, Yujin Tang, Dimitris N Metaxas, and Tong Che. Two heads are better than one: Test-time scaling of multi-agent collaborative reasoning. *arXiv preprint arXiv:2504.09772*, 2025.

Minki Kang, Jongwon Jeong, and Jaewoong Cho. T1: Tool-integrated verification for test-time compute scaling in small language models. In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust, editors, *International Conference on Learning Representations*, volume 2026, pages 73413–73444, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/776a5f2c7d6dd4b0d83145fc044e2726-Paper-Conference.pdf.

Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*, 13(9):9, 2024.

Shalev Lifshitz, Sheila A. McIlraith, and Yilun Du. Multi-agent verification: Scaling test-time compute with multiple verifiers, 2025a. URL https://arxiv.org/abs/2502.20379.

Shalev Lifshitz, Sheila A McIlraith, and Yilun Du. Multi-agent verification: Scaling test-time compute with multiple verifiers (abridged). In *Workshop on Reasoning and Planning for Large Language Models*, 2025b.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.

Zhan Ling, Yunhao Fang, Xuanlin Li, Zhiao Huang, Mingu Lee, Roland Memisevic, and Hao Su. Deductive verification of chain-of-thought reasoning. *Advances in Neural Information Processing Systems*, 36:36407–36433, 2023.

Yuejiang Liu, Parth Kothari, Bastien Van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. Ttt++: When does self-supervised test-time training fail or thrive? *Advances in Neural Information Processing Systems*, 34:21808–21820, 2021.

Dheeraj Mekala, Jason E Weston, Jack Lanchantin, Roberta Raileanu, Maria Lomeli, Jingbo Shang, and Jane Dwivedi-Yu. Toolverifier: Generalization to new tools via self-verification. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 5026–5041, 2024.

Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori B Hashimoto. s1: Simple test-time scaling. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 20286–20332, 2025.

Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. Llms know more than they show: On the intrinsic representation of llm hallucinations, 2025. URL https://arxiv.org/abs/2410.02707.

Archiki Prasad, Weizhe Yuan, Richard Yuanzhe Pang, Jing Xu, Maryam Fazel-Zarandi, Mohit Bansal, Sainbayar Sukhbaatar, Jason Weston, and Jane Yu. Self-consistency preference optimization. *arXiv preprint arXiv:2411.04109*, 2024.

Qwen and el al. Qwen2.5: A party of foundation models, September 2024. URL https://qwenlm.github.io/blog/qwen2.5/.

Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, Yulia Tsvetkov, Hannaneh Hajishirzi, Pang Wei Koh, and Luke Zettlemoyer. Spurious rewards: Rethinking training signals in rlvr, 2025. URL https://arxiv.org/abs/2506.10947.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

David Silver and Richard S Sutton. Welcome to the era of experience. *Google AI*, 1, 2025.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024a.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024b. URL https://arxiv.org/abs/2408.03314.

Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248. PMLR, 2020.

Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, Oluwasanmi Koyejo, Tatsunori Hashimoto, and Carlos Guestrin. Learning to (learn at test time): Rnns with expressive hidden states. *ArXiv*, abs/2407.04620, 2024a. URL https://api.semanticscholar.org/CorpusID:271039606.

Yu Sun, Xinhao Li, Karan Dalal, Jiarui Xu, Arjun Vikram, Genghan Zhang, Yann Dubois, Xinlei Chen, Xiaolong Wang, Sanmi Koyejo, et al. Learning to (learn at test time): Rnns with expressive hidden states. *arXiv preprint arXiv:2407.04620*, 2024b.

Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.

V Venktesh, Mandeep Rathee, and Avishek Anand. Trust but verify! a survey on verification design for test-time scaling. *arXiv preprint arXiv:2508.16665*, 2025.

Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce LLMs step-by-step without human annotations. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9426–9439, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.510. URL https://aclanthology.org/2024.acl-long.510/.

Renhao Wang, Yu Sun, Arnuv Tandon, Yossi Gandelsman, Xinlei Chen, Alexei A Efros, and Xiaolong Wang. Test-time training on video streams. *Journal of Machine Learning Research*, 26(9): 1–29, 2025a.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models, 2023. URL https://arxiv.org/abs/2203.11171.

Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, et al. Reinforcement learning for reasoning in large language models with one training example. *arXiv preprint arXiv:2504.20571*, 2025b.

Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws: An empirical analysis of compute-optimal inference for llm problem-solving. In *The Thirteenth International Conference on Learning Representations*, 2025.

Chulin Xie, Yangsibo Huang, Chiyuan Zhang, Da Yu, Xinyun Chen, Bill Yuchen Lin, Bo Li, Badih Ghazi, and Ravi Kumar. On memorization of large language models in logical reasoning. *arXiv preprint arXiv:2410.23123*, 2024.

An Yang et al. Qwen3 technical report, 2025.

Ling Yang, Zhaochen Yu, Tianjun Zhang, Shiyi Cao, Minkai Xu, Wentao Zhang, Joseph E Gonzalez, and Bin Cui. Buffer of thoughts: Thought-augmented reasoning with large language models. *Advances in Neural Information Processing Systems*, 37:113519–113544, 2024.

Mert Yuksekgonul, Daniel Koceja, Xinhao Li, Federico Bianchi, Jed McCaleb, Xiaolong Wang, Jan Kautz, Yejin Choi, James Zou, Carlos Guestrin, and Yu Sun. Learning to discover at test time, 2026. URL https://arxiv.org/abs/2601.16175.

Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *arXiv preprint arXiv:2503.18892*, 2025.

Kongcheng Zhang, Qi Yao, Shunyu Liu, Yingjie Wang, Baisheng Lai, Jieping Ye, Mingli Song, and Dacheng Tao. Consistent paths lead to truth: Self-rewarding reinforcement learning for llm reasoning, 2025. URL https://arxiv.org/abs/2506.08745.

Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative verifiers: Reward modeling as next-token prediction. *arXiv preprint arXiv:2408.15240*, 2024.

Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen Zhang, Xinwei Long, Ermo Hua, et al. Ttrl: Test-time reinforcement learning. *arXiv preprint arXiv:2504.16084*, 2025.

## A Implementation details: Pseudo-Code

Listing 1: Pseudo-code for **T$^3$RL**

```python
from collections import defaultdict

def t3rl_reward_fn(x, policy, verifier, sandbox, N, omega):
    Y = policy.sample_rollouts(x, n=N)
    vote, A = defaultdict(float), []

    for y in Y:
        a_hat = extract_answer(y)
        # (1) verifier generates code
        code = verifier.generate(x, y)
        # (2) sandbox executes the code and returns evidence
        t = sandbox.execute(code)
        # (3) verifier returns only a validity flag
        v = verifier.judge(x, y, a_hat, t)       # v in {0,1}
        vote[a_hat] += (1.0 if v == 0 else omega)
        A.append(a_hat)

    y_star  = max(vote, key=vote.get)      # weighted majority pseudo-label
    rewards = [1.0 if a_hat == y_star else 0.0 for a_hat in A]
    return y_star, rewards
```

## B Verifier System Prompt

The system prompt used to guide the external LLM verifier is shown in Figure 9. We carefully designed the instructions to ensure reliable, independent, and easily parsable tool-assisted verification. The key design choices in our prompt include:

- **Role Assignment:** Instructing the model to act as an "expert mathematician and Python programmer" sets a strong prior for generating rigorous, high-quality script formulations.
- **Independent Recomputation:** By explicitly stating "DO NOT assume the reasoning trace is correct" and "Prefer recomputing the answer directly," we mitigate confirmation bias. This prevents the verifier from blindly translating a flawed reasoning trace into code, forcing it to independently verify the underlying logic based on the original problem statement.
- **Trace as a Hint:** Allowing the verifier to use the candidate trace as a "hint" helps it navigate complex problems. It can leverage the policy model's mathematical intuition without being strictly bound to its (potentially flawed) arithmetic execution.

You are an expert mathematician and Python programmer.

You will be given:
- a math problem statement
- a candidate reasoning trace that claims to solve it

Your task is to VERIFY the candidate answer against the problem by writing Python code.

Rules:
- DO NOT assume the reasoning trace is correct.
- Prefer recomputing the answer directly from the PROBLEM STATEMENT.
- You may use the reasoning trace as a hint for what method to try.

Code Requirements:
- Use a SINGLE Python block formatted exactly as: ```python CODE ```
- Allowed modules: math, sympy, numpy, fractions
- Your code MUST print ONLY your final, computed result, e.g. print(result). Do not print hardcoded values.

Figure 9: System Prompt for the Verifier.

## C Failure Case: Small Verifiers

As quantitative results in Table 8 demonstrate, deploying an undersized verifier like `Qwen-Coder-0.5B` actively degrades performance compared to the standard TTRL baseline. This performance drop stems from the small model's limited capacity to strictly adhere to the verification instructions, which injects noise rather than reliable evidence into the reward signal. Qualitative analysis of the generated verification scripts reveals two primary failure modes:

- **Blind Copying and Hardcoded Outputs:** Despite explicit instructions in the system prompt to "DO NOT assume the reasoning trace is correct" and to strictly outputting hard-coded values, the 0.5B verifier frequently exhibits severe instruction-following failures. Instead of formulating an independent computational check, the model often defers entirely to the provided candidate trace. It typically hallucinates pseudo-reasoning within Python comments and bypasses actual computation, yielding a script that merely executes a hard-coded print statement of the trace's final answer (see Figure 10a). This confirmation bias neutralizes the benefits of executable verification, creating a false-positive signal that simply reinforces the unverified consensus.
- **Formatting and Compilation Errors:** A secondary, yet pervasive, issue is the small model's inability to consistently follow structural guidelines. The 0.5B verifier struggles to maintain valid Python syntax and adhere to the strict code-block formatting constraints required for automated extraction. This manifests as an increased frequency of compilation and execution errors, driven by missing import statements, malformed block delimiters, syntactical hallucinations, or endless comments (see Figure 10b). Consequently, the verification tool returns execution failures rather than reliable validity checks, further destabilizing the reward estimation process.

Combined, these limitations demonstrate that **T$^3$RL** requires a minimum threshold of verifier capacity to function effectively; below this threshold, the verifier acts as an additional source of stochastic noise rather than a reliable grounding mechanism.

| Model / Method | AIME 2024 | AMC | MATH-500 | Avg |
|---|---|---|---|---|
| Qwen-2.5-0.5B Vanilla (Baseline) | 0.0 | 4.8 | 7.9 | 4.2 |
| w/ TTRL | 0.4 | 12.0 | 34.6 | 15.7 |
| **w/ T$^3$RL (0.5B verifier)** | **0.0** | **10.8** | **32.0** | **14.3** |
| $\Delta$ (**T$^3$RL** − TTRL) | -0.4 | -1.2 | -2.6 | -1.4 |
| Rel. (% over TTRL) | ↓100.0% | ↓10.0% | ↓7.5% | ↓8.9% |

Table 8: Performance comparison illustrating failure cases with a weak (0.5B) verifier.

```
# Given the problem statement and reasoning trace,
we need to determine how many initial
configurations will lead to a transformed grid
consisting of a single filled square in the center
after a single transformation.

# The transformation rule states that a single
fill in the center of the grid means all the
square in this center remains filled, and none of
the surrounding empty squares become filled.
# The only square that remains filled in this
center could be the center, but since it must form
the border of another empty square, the center
must be surrounded by exactly three filled
neighbors.

# The number of initial configurations that will
lead to a transformed grid consisting of a single
filled square in the center after a single
transformation is 1.

print(1)
```

(a) Blind copying and hardcoded outputs. The model hallucinates reasoning in code comments and simply prints the unverified final answer given in the reasoning trace.

```
# Given values
a_3D = 6
angle_A = 120
side_lengths = [a_minus_3d, a_minus_3d, a_3D]

# Calculate the area using the Law of Cosines
# a^2 + b^2 = c^2
# c = a + b
# a^2 + b^2 = 36
# a^2 + b^2 = 36
# a^2 + b^2 = 6^2
# a^2 + b^2 = 36
# a^2 + b^2 = 6^2
# a^2 + b^2 = 36
# a^2 + b^2 = 6^2
# a^2 + b^2 = 36
# a^2 + b^2 = 6^2
# a^2 + b^2 = 36
# a^2 + b^2 = 6^2
# a^2 + b^2 = 6^2

...
```

(b) Formatting and compilation errors. The model fails to generate executable Python syntax and produces endless comments, resulting in compilation errors.

Figure 10: Qualitative examples of the two primary failure modes encountered when deploying an undersized verifier.

## NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: We claim significant improvements over TTRL across model sizes and benchmarks. This is supported by the strong performance gains we find across all experiments. We also claim that our method highlights the general importance of test-time verification for test-time reinforcement learning, which is supported by the novelty of our approach combined with the magnitude of its performance gains, compute efficiency, and increased training stability.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We discuss limitations and failure cases in Sec. 6.3 and Sec. C.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.

- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper is empirical; no theoretical results are claimed.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: All information needed for reproduction is specified under "Implementation Details" in Sec. 5.1, with pseudo code in Sec. A.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.

- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code is available at https://github.com/mayhugotong/T3RL. All datasets are public on HuggingFace: AIME 2024 (Maxwell-Jia/AIME_2024), AMC problems from NuminaMath (AI-MO/NuminaMath-CoT), MATH (DigitalLearningGmbH/MATH-lighteval), and Knights-and-Knaves (K-and-K/knights-and-knaves). We do not provide additional data. Pseudo code is in Sec. A.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: We provide all relevant details in the experimental setup (Sec. 5.1).

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Main accuracy numbers (Table 1) follow the DeepSeek-R1 protocol and report pass@1 averaged over 16 evaluations for AIME 2024 and 4 for AMC and MATH-500, with $\pm$ one standard deviation across runs (Sec. 5.1). Figures 8b and 8c report run-to-run variability of AIME training.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We report token use, training time, and FLOPs per training step in Table 4 and inference token use in Table 5. Per-step statistics in Table 4 reflect AIME runs; the hardware (8 NVIDIA A100 GPUs) and episode counts in Sec. 5.1 suffice to estimate aggregate cost.

   Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: There is no conflict with the Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: See the Impact Statement after the conclusion.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: Our method does not release new pre-trained models or scraped datasets, so no additional safeguards are required.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite the original source of every dataset and model and use each within its license. Datasets: AIME 2024 (MIT), NuminaMath-CoT (Apache 2.0), MATH-lighteval (MIT), and Knights-and-Knaves (CC BY-NC-SA 4.0), all from HuggingFace. Models: Qwen2.5 and Qwen3 (Apache 2.0) and Llama-3.2 (Llama 3.2 Community License).

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [N/A]

Justification: We do not introduce new datasets or model checkpoints.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.

- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: No human subjects or crowdsourcing were involved.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: No human subjects or crowdsourcing were involved.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [Yes]

    Justification: While the use of LLM as a tool-based verifier is standard, we describe its concrete use in Sec. 4 and provide the system prompt in Sec. B.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.