

JURY-RL: Votes Propose, Proofs Dispose for Label-Free RLVR

Xinjie Chen^{1,2}, Biao Fu¹, Jing Wu¹, Guoxin Chen¹,
Xinggao Liu², Dayiheng Liu¹, Minpeng Liao^{1,*}

¹Token Foundry, Alibaba Group

²Zhejiang University

{chenxinjie.cxj, fubiaobiao.fu, minpeng.lmp}@alibaba-inc.com

Abstract

Reinforcement learning with verifiable rewards (RLVR) enhances reasoning in large language models (LLMs) but typically depends on human-annotated answers or curated reward specifications. Label-free alternatives such as majority voting or LLM-as-a-Judge remove this cost but introduce false positives that destabilize training. We introduce **JURY-RL**, a label-free RLVR framework that decouples answer proposal from reward disposal: rollouts matching the plurality-voted answer receive positive reward only when that answer is verified in Lean. When verification is inconclusive, **ResZero (Residual-Zero)** discards the unverified plurality and redistributes a zero-mean, variance-preserving signal over the residual answers, keeping the optimization gradient stable without reinforcing unverifiable consensus. Across three backbones trained on math data, JURY-RL outperforms label-free baselines on mathematical reasoning, transfers to code generation and general benchmarks, and surpasses a majority-vote baseline in the test-time RL (TTRL) setting; it matches supervised ground-truth training on pass@1 and beats it on pass@k and response diversity.

1 Introduction

Large language models (LLMs) (OpenAI, 2023; DeepSeek-AI et al., 2024; Yang et al., 2025b) continue to advance in broad capabilities, yet reliable reasoning remains a core bottleneck (Wang et al., 2023; Lightman et al., 2024; Patil and Jadon, 2025; Huang et al., 2025a). Recent progress in reinforcement learning with verifiable rewards (RLVR) (Shao et al., 2024; Lambert et al., 2024; DeepSeek-AI et al., 2025) offers a principled post-training route: rather than rewarding what merely appears plausible, RLVR aligns learning with externally

checkable correctness through signals such as program execution, unit tests, or mathematical verification (Cobbe et al., 2021; Yu et al., 2025; AI and Agentica, 2025). However, in many practical settings, RLVR still depends on human-annotated answers or carefully curated reward specifications, both of which are costly to obtain and limited in coverage (Ouyang et al., 2022; Bai et al., 2022; Shao et al., 2024).

To reduce annotation cost, recent work has explored label-free rewards, where no human ground-truth answers are provided during training. A prominent subset is self-supervised self-reward, which derives signals from the model or unlabeled data itself, including entropy minimization (Prabhudesai et al., 2025), self-certainty (Zhao et al., 2025b), and majority voting (Shafayat et al., 2025; Zhang et al., 2025; Ye et al., 2026). These approaches are attractive because they avoid per-example labels, but they often optimize agreement or confidence rather than verified correctness. As a result, they are vulnerable to false positives, reward hacking, and training collapse (Gao et al., 2022; Shafayat et al., 2025). LLM-as-a-Judge (Lee et al., 2024; Su et al., 2025; Zhao et al., 2025c) provides another label-free path, yet in practice it remains sensitive to prompting and format manipulation, incurs nontrivial compute overhead, and can inherit shared-model biases (Pang et al., 2023; Chen et al., 2024; Thakur et al., 2025; Shi et al., 2025).

The flaws of existing label-free methods, from reinforcing false consensus to reward hacking, are not isolated issues but rather symptoms of a deeper challenge. We argue that a truly robust reward signal must satisfy three essential properties: (i) **scalable** without costly human supervision, (ii) **truth-aligned**, rewarding verifiable correctness instead of error-prone consensus, and (iii) **optimization-stable**, allowing learning to proceed even when verification is inconclusive. Previous approaches have, to one degree or another, failed to meet all

* Corresponding author.




Figure 1: The **JURY-RL** workflow: Votes Propose, Proofs Dispose. For each problem, a majority vote from multiple rollouts (**Jury**) proposes a candidate answer. A Lean verifier (**Judge**) then disposes the reward. If the answer is **Verified** ($\delta = 1$), supporting rollouts receive a positive reward, directly linking the update to successful formal verification. Conversely, when verification is **Inconclusive** ($\delta = 0$), all rollouts receive the proposed ResZero (Residual-Zero) fallback reward.

three criteria. Self-supervised signals such as majority voting achieve scalability but sacrifice truth-alignment, while LLM-as-a-Judge struggles with both truth-alignment (due to high false positives) and optimization instabilities.

This motivates our core proposal, a new paradigm designed to resolve this conflict: **Votes Propose, Proofs Dispose**. Our strategy is to decouple the *proposal* of a candidate answer from the final *disposal* of its reward. To maintain scalability, votes from model rollouts first propose a single consensus candidate through a computationally cheap majority vote. A formal theorem prover (de Moura and Ullrich, 2021; Ren et al., 2025; Lin et al., 2025) then acts as a reliable judge to dispose the ultimate reward for this single candidate, thus satisfying all three principles above. This design choice avoids the prohibitive cost of formally verifying every unique answer, thereby making the entire framework viable at scale.

Our framework, **JURY-RL**, is shown in Figure 1. When the plurality-voted answer is successfully verified, only the rollouts that produced that answer receive positive reward. This suppresses the false-positive reinforcement common in consensus- or judge-based surrogates. A remaining challenge is how to learn when verification is inconclusive. Rewarding the plurality would reintroduce the very failure mode we seek to avoid, whereas assigning zero reward to the entire group would collapse the

grouped learning signal under GRPO.

To address this, we introduce **ResZero (Residual-Zero)**, a fallback reward that discards the unverified plurality proposal and assigns a zero-mean, variance-preserving reward to the residual answers. This design maintains a stable optimization gradient that allows learning to proceed without amplifying a potentially spurious consensus, ensuring both training stability and truth alignment. By combining proof-gated positive reward with a conservative residual fallback, JURY-RL remains trainable even when formal verification does not produce a decisive signal.

Contributions. (1) We introduce **JURY-RL**, a label-free RLVR framework for machine-checkable settings that decouples candidate proposal from proof-gated reward disposal, requiring only a single formal verification call per rollout group. (2) We propose **ResZero (Residual-Zero)**, a zero-mean, variance-preserving fallback reward for inconclusive verification. It preserves optimization stability while refusing to positively reinforce an unverified plurality. (3) Across mathematical reasoning, code generation, and multi-task benchmarks, JURY-RL trains more stably and achieves state-of-the-art results among label-free methods, matching a strong supervised ground-truth baseline in pass@1 while consistently surpassing it in pass@k and solution diversity. It further transfers to the test-time RL setting, surpassing a majority-vote TTRL method.

2 Related Work

LLM Reasoning. The general capabilities of LLMs have rapidly expanded (OpenAI, 2023; Dubey et al., 2024; Yang et al., 2025b), yet reliable mathematical and programmatic reasoning remains a bottleneck: models often optimize for plausibility rather than verifiable correctness (Ouyang et al., 2022; Rafailov et al., 2023; Touvron et al., 2023). Post-training techniques that elicit step-by-step reasoning (e.g., chain-of-thought and self-consistency) can raise average accuracy but also amplify confident but wrong results when no external check is available (Wei et al., 2022; Wang et al., 2023). These problems motivate recent verifiability-aligned training signals that reward what is provably correct rather than what appears correct (Shao et al., 2024; Lambert et al., 2024; Hu et al., 2025; DeepSeek-AI et al., 2025; Kimi-Team et al., 2025; Yang et al., 2025a).

Label-Free RLVR. To scale beyond labeled specifications, label-free surrogates derive rewards from the model or unlabeled data itself, e.g., via majority voting (Shafayat et al., 2025; Zhang et al., 2025), confidence (Zhao et al., 2025b), entropy (Prabhudesai et al., 2025; Agarwal et al., 2025), or LLM-as-a-Judge (Pang et al., 2023; Lee et al., 2024; Su et al., 2025; Zhao et al., 2025a). While attractive for their broad coverage and low cost, these signals are prone to false positives, prompt/format gaming (Zhao et al., 2025c), and training collapse (Zhang et al., 2025). As a result, consensus- or judge-approval-based methods risk diverging from ground truth, leading to reward hacking and instability (Shafayat et al., 2025; Zhang et al., 2025). Our work targets this conflict: retain the scalability of label-free training while removing optimism toward unverified agreement.

Lean and Other Verifiers. Verification-based training employs externally checkable signals such as program execution and unit tests, SMT solvers, or formal proof assistants such as Lean/Coq to couple reward with correctness (P. Huet et al., 1997; C. Blanchette et al., 2011; de Moura and Ullrich, 2021; Cobbe et al., 2021; AI and Agentica, 2025). Previous generate-then-verify pipelines typically provide no learning signal when verification fails, limiting stability and sample efficiency. JURY-RL decouples proposal from disposal to preserve label-free scalability while maintaining optimization alignment with provable correctness in the

positive-reward branch. Unlike process-reward approaches or hybrid verifiers (Huang et al., 2025b) that rely on learned judges, JURY-RL uses successful Lean verification as the final acceptance mechanism for positive reward, conditional on semantically faithful formalization. The autoformalization pipeline is therefore part of the deployed verification system rather than an independent learned supervisor, and its practical imperfections primarily affect whether a candidate can be certified.

3 Preliminaries

Problem Setup. Let π_θ denote a policy LLM with parameters θ . Given a problem x , the model generates a token sequence $y = (y_1, \dots, y_n) \sim \pi_\theta(\cdot | x)$ and a deterministic extractor $\text{ans}(\cdot)$ parses a candidate answer $a = \text{ans}(y)$. In the *label-free* setting, ground-truth answers are unavailable during training. Instead, each x can be associated with a machine-checkable specification $\text{spec}(x)$, and an external verifier (e.g., a Lean-based checker) exposes a binary oracle

$$\text{verify}(x, a) \in \{0, 1\},$$

which returns 1 when the candidate is accepted by the deployed verification pipeline. In our setting, this means successful Lean verification of a semantically faithful formalization; $\text{verify}(x, a) = 0$ may arise either because the answer is wrong or because upstream formalization or proof-search steps fail to certify it.

We train with Group Relative Policy Optimization (GRPO) (Shao et al., 2024), using a KL penalty to a reference policy. For each problem x , sample G rollouts $\mathcal{G}_x = \{y_i\}_{i=1}^G$ from π_{old} and compute grouped rewards $r_i := r(x, y_i; \mathcal{G}_x)$. The group-normalized scalar advantage is

$$\hat{A}_i = \frac{r_i - \bar{r}}{\text{std}(\{r_j\}_{j=1}^G) + \varepsilon}, \quad \bar{r} = \frac{1}{G} \sum_{j=1}^G r_j. \quad (1)$$

GRPO then broadcasts $\hat{A}_i$ token-wise in the standard PPO-clipped update with KL regularization; our analysis only depends on the normalized group advantage above.

Label-free Self-Reward in Reinforcement Learning. Existing label-free reward signals can be broadly grouped by where the supervision comes from. Model-internal signals reward confidence, low entropy, or agreement among sampled answers,

including single-view majority voting and multi-view consistency (Prabhudesai et al., 2025; Agarwal et al., 2025; Shafayat et al., 2025; Zhang et al., 2025). External-judge signals instead use another model, such as an LLM-as-a-Judge or an LLM-based knowledge distillation (LLM-KD) teacher, to score or supervise outputs (Lee et al., 2024; Su et al., 2025; Zhao et al., 2025a; Gu et al., 2024b). These approaches are scalable but can still reinforce confident or judge-approved false positives, motivating our proof-gated design.

4 JURY-RL

To operationalize the principles of scalability, truth alignment, and optimization stability, JURY-RL decouples the answer proposal process from reward disposal. In this framework, votes from model rollouts serve as a scalable proposal mechanism, while a formal theorem prover acts as the definitive judge for reward disposal. By restricting formal verification to *only the single* majority-voted candidate rather than evaluating every unique response, we maintain computational tractability while ensuring the reward signal remains anchored in provable correctness.

This section details the two key components of our framework: the overarching **proof-gated reward mechanism** that enforces truth alignment when verification succeeds, and the **Residual-Zero (ResZero) fallback** designed to maintain optimization stability when verification is inconclusive.

4.1 The Proof-Gated Reward Framework

The JURY-RL workflow follows the decoupled paradigm introduced above. First, in the **proposal stage**, given a problem x , we generate G trajectories $\{y_i\}_{i=1}^G \sim \pi_\theta(\cdot|x)$ and parse their corresponding answers $a_i = \text{ans}(y_i)$. The policy proposes a single candidate $\hat{a}$ via majority voting:

$$\hat{a} = \arg \max_a \sum_{i=1}^G \mathbb{I}[a_i = a]. \quad (2)$$

Next, in the **disposal stage**, a single call to the external Lean verifier evaluates $\hat{a}$ against the formal specification of x . This yields a binary *proof-gate*, $\delta = \text{verify}(x, \hat{a}) \in \{0, 1\}$, which governs the reward assignment.

The final reward r_i for each trajectory y_i is determined by a conditional function gated by δ :

$$r_i = \underbrace{\delta \cdot \mathbb{I}[a_i = \hat{a}]}_{\text{Verified Correctness}} + \underbrace{(1 - \delta) \cdot r_i^{\text{ResZero}}}_{\text{Inconclusive Fallback}}, \quad (3)$$

where r_i^{ResZero} is the residual reward defined in Section 4.2.

The stability of this proof-gated design stems from its principled handling of both successful and inconclusive verification. First, when verification succeeds ($\delta = 1$), a positive reward is granted exclusively to the trajectories that produced the proven-correct answer. This approach binds the learning signal to hard evidence, which, under standard soundness assumptions, suppresses the false positives that plague self-reward or judge-based surrogates (Zhang et al., 2025; Cobbe et al., 2021). Second, when verification is inconclusive ($\delta = 0$), the system defaults to ResZero fallback rather than rewarding the majority consensus. This centered substitute maintains a stable optimization gradient by preserving group-wise variance, which prevents learning from stalling or oscillating, a common situation when verification fails for reasons like search limits or incomplete proofs. By paying only for verifiable correctness, this framework narrows the attack surface for prompt and format hacking, a vulnerability in other label-free systems.

4.2 ResZero Reward

When formal verification is inconclusive, a learning signal is still needed to maintain optimization stability. However, a naive choice like rewarding the majority vote (MV) is brittle and risks training collapse. Using MV as a reward conflates agreement with correctness and can induce entropy collapse under GRPO, as the spurious consensus strengthens. A simple zero-reward fallback is also suboptimal, as it would lead to zero group-wise advantage and stall the learning process.

To address this, we introduce the **ResZero (Residual-Zero) Reward**. Its principle is to **penalize the unverifiable majority proposal and construct a meaningful, zero-mean reward from the remaining (residual) answers**. This design preserves a useful learning signal by maintaining variance among minority opinions without amplifying a potentially false consensus. Furthermore, we propose an adaptive variant that strengthens this signal precisely when the model is most confidently wrong. The intuition is that a strong but unverified consensus (indicated by a high majority share, α) requires a stronger corrective signal. ResZero operationalizes this by using α to simultaneously amplify the reward signal for residual answers and suppress the majority answer.

Let $M = \{i : a_i = \hat{a}\}$ be the set of rollouts

supporting the majority answer and $R = \{i : a_i \neq \hat{a}\}$ be the set of residual rollouts. The majority share is $\alpha = |M|/G$. We first define the leave-one-out residual share for an answer b within the residual group:

$$u^{(-i)}(b) = \frac{1}{|R| - 1} \sum_{\substack{j \in R \\ j \neq i}} \mathbb{I}[a_j = b].$$

Let $z_i = u^{(-i)}(a_i)$ if $i \in R$ (the relative support for a residual answer a_i within its peer group) and $z_i = 0$ if $i \in M$. The **ResZero** reward is then assigned as:

$$r_i^{\text{ResZero}} = \underbrace{\alpha \cdot \mathbb{I}[i \in R] \cdot (z_i - \bar{u})}_{\text{Amplify residual}} - \underbrace{c\alpha \cdot \mathbb{I}[i \in M]}_{\text{Penalize majority}} + \underbrace{\gamma}_{\text{Re-center}}, \quad (4)$$

$$\text{where } \bar{u} = \frac{1}{|R|} \sum_{j \in R} z_j, \quad \gamma = c\alpha^2.$$

Here, c is a positive hyperparameter controlling the penalty strength. γ is a global re-centering term to ensure the total reward sums to zero. The term $(z_i - \bar{u})$ creates a zero-mean signal *within the residual group*, rewarding answers that are more popular among the minorities and penalizing those that are less so. This entire residual signal is then scaled by α . By construction, the total reward sums to zero ($\sum_i r_i^{\text{ResZero}} = 0$), preserving the zero-mean property crucial for GRPO stability.¹

The design of our ResZero reward ensures robust optimization through three properties. (i) **Variance preservation**: It maintains non-zero variance among differing answers, which is critical for variance-normalized optimizers like GRPO to prevent vanishing gradients. (ii) **Zero-mean construction**: Its strictly zero-mean property makes it a principled, optimizer-agnostic signal, ensuring portability beyond GRPO to any general RL paradigm. (iii) **Adaptive economy**: The corrective signal dynamically scales with the majority share α , applying maximum pressure when the model is most confidently wrong, with this entire behavior governed by only a single hyperparameter c . These properties prevent collapse into spurious consensus and foster robust, exploratory learning.

This dual-reward strategy hinges on the verifier’s outcome. If verification succeeds ($\delta=1$), the policy

update is guided by verifiable correctness. Conversely, if inconclusive ($\delta=0$), our ResZero fallback penalizes the unverified majority, which is critical for mitigating entropy collapse caused by spurious consensus. This design maintains the scalability of label-free training by requiring only a single verification per step. The full procedure is detailed in Appendix B.

5 Experiments

5.1 Setting

Backbone Models. We evaluate JURY-RL on three open-source backbones spanning base and instruction-tuned models: Qwen3-1.7B-Base (Yang et al., 2025a), Llama-3.2-3B-Instruct (Dubey et al., 2024) and Qwen2.5-7B (Yang et al., 2025b), to cover different sizes and instruction-following behaviors.

Baselines. We compare JURY-RL with several established label-free and supervised reward baselines. The self-supervised baselines include **Majority-Voting** (Shafayat et al., 2025), **Self-Certainty** (Zhao et al., 2025b), **Entropy** minimization (Prabhudesai et al., 2025) and **CoReward** (Zhang et al., 2025). To further contextualize the comparison, we include **ground-truth supervised oracle** and **LLM-KD** as baselines. Additionally, we benchmark against the judge-based method, **LLM-as-a-Judge** (Pang et al., 2023; Zhao et al., 2025a), to provide a comprehensive evaluation. Additional details are provided in Appendix D.

Implementation Details. All methods use the VeRL framework (Sheng et al., 2025) on $8 \times$ NVIDIA A100 GPUs, train on 7,500 MATH problems (Hendrycks et al., 2021), and evaluate on a held-out 5,000-problem split (MATH5000). We use the officially released chat-based prompting formats for fair comparison; GRPO hyperparameters and per-method settings are in Appendix E.

Evaluation Datasets and Metrics. To comprehensively assess model capabilities, we evaluate on a suite of benchmarks covering mathematical reasoning, code generation, and general abilities. **Mathematical Reasoning:** We evaluate on the AIME24/25 (Hugging Face H4, 2024; OpenCompass, 2025), MATH500 (Lightman et al., 2024), GSM8K (Cobbe et al., 2021), and the competition-level AMC datasets (math-ai Team, 2024). **Code Generation:** We assess coding proficiency using

¹See App. A.2 for a formal proof and App. B.2 for a worked example. $c = 0.01$ in our experiments.

Methods	Mathematics					Code		Instruction	Multi-Task	Average
	AIME24	AIME25	MATH-500	GSM8K	AMC	LiveCode	CRUX	IFEval	MMLU-Pro	
Qwen3-1.7B-Base										
Before RL GT-Reward	3.12 $\pm$ 1.8	1.25 $\pm$ 1.1	46.85 $\pm$ 2.9	64.20 $\pm$ 1.8	20.78 $\pm$ 2.5	4.59 $\pm$ 0.7	7.52 $\pm$ 1.1	33.66 $\pm$ 1.6	33.60 $\pm$ 0.7	23.95
	6.88 $\pm$ 2.1	5.21 $\pm$ 2.0	68.35 $\pm$ 2.3	82.01 $\pm$ 1.1	30.87 $\pm$ 3.3	14.84 $\pm$ 0.5	33.73 $\pm$ 0.6	38.70 $\pm$ 1.6	39.43 $\pm$ 0.7	35.56
LLM-KD	6.67 $\pm$ 1.1	3.54 $\pm$ 1.2	67.8 $\pm$ 1.2	81.97 $\pm$ 1.3	35.69 $\pm$ 2.2	14.05 $\pm$ 0.9	33.80 $\pm$ 1.5	34.65 $\pm$ 1.6	37.97 $\pm$ 0.7	35.13
Self-Certainty Entropy	3.54 $\pm$ 1.6	3.12 $\pm$ 1.8	57.05 $\pm$ 2.7	73.71 $\pm$ 1.4	27.41 $\pm$ 2.3	10.37 $\pm$ 0.9	18.73 $\pm$ 1.3	38.58 $\pm$ 1.6	35.28 $\pm$ 0.7	29.75
	6.67 $\pm$ 1.6	4.17 $\pm$ 1.5	65.0 $\pm$ 1.7	79.45 $\pm$ 1.6	31.17 $\pm$ 3.0	12.64 $\pm$ 1.0	30.38 $\pm$ 0.5	35.00 $\pm$ 1.6	37.01 $\pm$ 0.7	33.50
Majority-Voting	2.08 $\pm$ 1.3	2.08 $\pm$ 1.3	59.6 $\pm$ 0.9	81.08 $\pm$ 0.8	29.82 $\pm$ 1.7	14.29 $\pm$ 0.5	32.00 $\pm$ 2.3	37.16 $\pm$ 1.6	35.68 $\pm$ 0.7	32.64
CoReward	6.88 $\pm$ 1.5	3.12 $\pm$ 1.5	65.4 $\pm$ 1.2	81.07 $\pm$ 0.8	33.13 $\pm$ 3.0	14.27 $\pm$ 0.5	34.10 $\pm$ 1.1	37.28 $\pm$ 1.6	37.39 $\pm$ 0.7	34.74
LLM-as-a-Judge	5.83 $\pm$ 1.5	1.88 $\pm$ 0.9	63.0 $\pm$ 2.4	81.84 $\pm$ 1.4	33.89 $\pm$ 1.1	14.33 $\pm$ 0.8	34.98 $\pm$ 1.5	36.28 $\pm$ 1.6	33.83 $\pm$ 0.7	33.98
JURY-RL (Ours)	7.71 $\pm$ 1.8	4.58 $\pm$ 2.0	66.4 $\pm$ 2.0	82.58 $\pm$ 0.3	34.19 $\pm$ 3.1	14.54 $\pm$ 0.4	33.98 $\pm$ 2.5	36.55 $\pm$ 1.6	38.11 $\pm$ 0.7	35.40
Llama-3.2-3B-Instruct										
Before RL GT-Reward	6.88 $\pm$ 2.1	0.42 $\pm$ 0.6	44.4 $\pm$ 2.1	68.78 $\pm$ 1.2	17.47 $\pm$ 3.3	3.00 $\pm$ 0.4	24.70 $\pm$ 1.2	54.41 $\pm$ 1.8	32.01 $\pm$ 0.7	28.01
	10.62 $\pm$ 2.1	0.21 $\pm$ 0.4	47.5 $\pm$ 1.4	78.60 $\pm$ 0.8	22.89 $\pm$ 1.9	7.05 $\pm$ 0.6	32.48 $\pm$ 0.8	50.16 $\pm$ 1.7	34.26 $\pm$ 0.7	31.53
LLM-KD	10.21 $\pm$ 1.9	0.42 $\pm$ 0.6	51.25 $\pm$ 4.5	79.72 $\pm$ 1.2	24.55 $\pm$ 3.4	7.55 $\pm$ 0.6	31.20 $\pm$ 1.6	49.06 $\pm$ 1.7	34.04 $\pm$ 0.7	32.00
Self-Certainty Entropy	2.71 $\pm$ 1.3	0.62 $\pm$ 0.7	42.4 $\pm$ 0.9	73.52 $\pm$ 1.2	17.62 $\pm$ 1.8	5.57 $\pm$ 0.9	24.55 $\pm$ 1.2	54.11 $\pm$ 1.8	34.42 $\pm$ 0.7	28.39
	3.54 $\pm$ 1.5	0.21 $\pm$ 0.4	40.5 $\pm$ 2.9	68.57 $\pm$ 1.7	17.17 $\pm$ 2.6	5.38 $\pm$ 0.8	26.30 $\pm$ 1.7	54.24 $\pm$ 1.8	33.54 $\pm$ 0.7	27.72
Majority-Voting	8.33 $\pm$ 1.3	0.00 $\pm$ 0.0	48.4 $\pm$ 3.0	78.64 $\pm$ 1.4	21.69 $\pm$ 2.8	8.11 $\pm$ 0.5	30.83 $\pm$ 1.4	48.36 $\pm$ 1.7	33.94 $\pm$ 0.7	30.92
CoReward	7.71 $\pm$ 1.4	0.00 $\pm$ 0.0	47.2 $\pm$ 1.9	80.00 $\pm$ 1.7	19.58 $\pm$ 3.6	5.57 $\pm$ 0.4	30.87 $\pm$ 1.2	50.46 $\pm$ 1.7	32.95 $\pm$ 0.7	30.48
LLM-as-a-Judge	9.38 $\pm$ 1.9	0.62 $\pm$ 0.7	47.35 $\pm$ 2.5	78.51 $\pm$ 2.2	20.93 $\pm$ 3.4	3.96 $\pm$ 0.5	31.45 $\pm$ 1.4	51.46 $\pm$ 1.7	34.22 $\pm$ 0.7	30.88
JURY-RL (Ours)	9.17 $\pm$ 2.0	1.25 $\pm$ 0.9	48.8 $\pm$ 2.6	78.96 $\pm$ 0.5	26.05 $\pm$ 1.6	6.16 $\pm$ 0.6	31.60 $\pm$ 1.2	50.09 $\pm$ 1.8	34.54 $\pm$ 0.7	31.85
Qwen2.5-7B										
Before RL GT-Reward	4.38 $\pm$ 1.4	1.67 $\pm$ 1.4	49.2 $\pm$ 2.6	71.00 $\pm$ 1.9	22.44 $\pm$ 3.7	4.57 $\pm$ 0.5	27.38 $\pm$ 2.3	40.61 $\pm$ 1.7	44.03 $\pm$ 0.8	29.48
	18.33 $\pm$ 2.3	11.25 $\pm$ 1.7	75.1 $\pm$ 1.4	90.83 $\pm$ 0.5	46.54 $\pm$ 2.5	12.78 $\pm$ 1.6	53.67 $\pm$ 1.0	41.50 $\pm$ 1.7	45.09 $\pm$ 0.7	43.90
LLM-KD	12.08 $\pm$ 2.0	7.92 $\pm$ 1.7	72.7 $\pm$ 1.2	90.96 $\pm$ 0.9	44.28 $\pm$ 2.1	8.09 $\pm$ 0.5	53.57 $\pm$ 0.5	43.09 $\pm$ 1.7	43.73 $\pm$ 0.7	41.82
Self-Certainty Entropy	9.79 $\pm$ 2.4	8.96 $\pm$ 1.8	72.45 $\pm$ 2.8	88.21 $\pm$ 0.5	40.51 $\pm$ 2.1	11.87 $\pm$ 0.9	53.18 $\pm$ 1.4	39.66 $\pm$ 1.7	43.89 $\pm$ 0.7	40.95
	11.04 $\pm$ 2.3	8.33 $\pm$ 2.2	73.15 $\pm$ 1.7	87.85 $\pm$ 0.6	40.51 $\pm$ 3.1	15.68 $\pm$ 0.9	51.08 $\pm$ 1.3	40.25 $\pm$ 1.7	42.61 $\pm$ 0.7	41.17
Majority-Voting	11.25 $\pm$ 2.1	4.17 $\pm$ 1.2	71.0 $\pm$ 0.7	90.52 $\pm$ 0.7	38.70 $\pm$ 1.3	18.37 $\pm$ 0.9	52.20 $\pm$ 1.1	42.72 $\pm$ 1.7	43.83 $\pm$ 0.8	41.42
CoReward	11.25 $\pm$ 1.8	3.96 $\pm$ 1.9	73.85 $\pm$ 1.1	90.16 $\pm$ 0.9	40.81 $\pm$ 1.8	10.37 $\pm$ 0.5	55.08 $\pm$ 2.4	43.75 $\pm$ 1.7	42.08 $\pm$ 0.7	41.26
LLM-as-a-Judge	10.42 $\pm$ 2.1	5.62 $\pm$ 1.2	72.3 $\pm$ 1.2	89.73 $\pm$ 0.7	41.11 $\pm$ 3.1	10.58 $\pm$ 0.3	51.90 $\pm$ 1.6	44.17 $\pm$ 1.7	48.67 $\pm$ 0.7	41.61
JURY-RL (Ours)	13.75 $\pm$ 1.8	7.50 $\pm$ 1.6	75.65 $\pm$ 0.7	90.35 $\pm$ 0.9	47.29 $\pm$ 3.2	14.69 $\pm$ 1.1	54.37 $\pm$ 0.4	41.51 $\pm$ 1.7	50.00 $\pm$ 0.8	43.90

Table 1: RL performance (%) on math reasoning benchmarks. Cell background colors indicate relative performance: darker colors denote better results within each model group.

LiveCodeBench (Jain et al., 2024) and CRUX (Gu et al., 2024a). **Instruction-Following and Multi-Task:** General abilities are measured using IFEval (Zhou et al., 2023) for instruction following and MMLU-Pro (Wang et al., 2024) for multi-task understanding. We report avg@k accuracy with k varying by benchmark difficulty: $k=16$ for AIME24/25, $k=4$ for MATH500 and GSM8K, $k=8$ for AMC, and $k=5$ for code generation benchmarks. More details in Appendix F.

5.2 Main Results

We evaluate **JURY-RL** on both in- and out-of-domain tasks, as shown in Table 1 under avg@k settings². A key finding is that JURY-RL not only outperforms all label-free RL baselines but also consistently matches the supervised GRPO with ground-truth rewards (GT), suggesting that proof-gated rewards can serve as a promising alternative to direct supervision.

Mathematical Reasoning and In-Domain Generalization. Across three backbones, JURY-RL attains overall performance comparable to the supervised GT baseline, with the clearest gains on harder competition-style evaluations and under multi-sample decoding. The gains on the **in-domain** MATH500 set are more modest, which is unsurprising given that training is performed on MATH. In contrast, the improvements on **out-of-domain** math benchmarks—most consistently on AMC, and on GSM8K for smaller backbones—indicate that the proof-gated objective is not merely memorizing dataset-specific answer patterns. Table 2 further shows that JURY-RL improves average pass@1 over both GT-Reward and LLM-as-a-Judge, while achieving even larger gains in average pass@k, suggesting that the method improves both single-sample competence and the breadth of useful candidate solutions.

Out-of-Domain Generalization. The model exhibits robust out-of-domain performance across code generation, instruction following, and multi-

²Values are reported as mean $\pm$ 95% confidence intervals.

Model	Average (pass@k)					Average (pass@1)				
	GT-Reward	LLM-as-a-judge	JURY-RL	Δ_{GT} (pp)	Δ_{LJ} (pp)	GT-Reward	LLM-as-a-judge	JURY-RL	Δ_{GT} (pp)	Δ_{LJ} (pp)
Qwen3-1.7B-Base	55.36	50.35	59.41	+4.05	+9.06	39.25	37.60	41.57	+2.32	+3.97
Llama-3.2-3B-Instruct	45.46	45.93	48.48	+3.02	+2.55	32.19	30.65	34.10	+1.91	+3.45
Qwen2.5-7B	62.48	53.99	64.04	+1.56	+10.05	46.60	44.13	48.13	+1.53	+4.00

Table 2: Performance of JURY-RL vs. GT-Reward and LLM-as-a-judge on math reasoning tasks. $k=16$ for AIME, $k=4$ for MATH500 and GSM8K, and $k=8$ for AMC. Δ_{GT} and Δ_{LJ} denote JURY-RL’s improvements (in percentage points) over GT-Reward and LLM-as-a-judge, respectively.

Methods	Mathematics					Code		Instruction	Multi-Task	Average
	AIME24	AIME25	MATH-500	GSM8K	AMC	LiveCode	CRUX	IFEval	MMLU-Pro	
Qwen3-1.7B-Base										
GT-Reward	6.88 $\pm$ 2.1	5.21 $\pm$ 2.0	68.35 $\pm$ 2.3	82.01 $\pm$ 1.1	30.87 $\pm$ 3.3	14.84 $\pm$ 0.5	33.73 $\pm$ 0.6	38.70 $\pm$ 1.6	39.43 $\pm$ 0.7	35.56
Majority-Voting	2.08 $\pm$ 1.3	2.08 $\pm$ 1.3	59.6 $\pm$ 0.9	81.08 $\pm$ 0.8	29.82 $\pm$ 1.7	14.29 $\pm$ 0.5	32.00 $\pm$ 2.3	37.16 $\pm$ 1.6	35.68 $\pm$ 0.7	32.64
Proof-Gate + Zero Reward	3.54 $\pm$ 1.6	2.92 $\pm$ 1.4	59.85 $\pm$ 1.2	83.21 $\pm$ 0.5	30.87 $\pm$ 2.4	15.00 $\pm$ 0.8	34.10 $\pm$ 1.4	37.40 $\pm$ 1.6	34.43 $\pm$ 0.7	33.48
Proof-Gate + Random Reward	4.58 $\pm$ 1.4	2.50 $\pm$ 1.2	59.75 $\pm$ 2.0	77.58 $\pm$ 0.8	26.96 $\pm$ 1.3	9.14 $\pm$ 1.3	18.48 $\pm$ 1.6	37.19 $\pm$ 1.6	35.12 $\pm$ 0.7	30.14
Proof-Gate + MV Reward	0.00 $\pm$ 0.0	0.00 $\pm$ 0.0	41.2 $\pm$ 1.0	82.47 $\pm$ 0.4	9.64 $\pm$ 0.8	14.48 $\pm$ 1.1	34.40 $\pm$ 2.1	36.78 $\pm$ 1.6	35.33 $\pm$ 0.7	28.26
JURY-RL (Proof-Gate + ResZero)	7.71 $\pm$ 1.8	4.58 $\pm$ 2.0	66.4 $\pm$ 2.0	82.58 $\pm$ 0.3	34.19 $\pm$ 3.1	14.54 $\pm$ 0.4	33.98 $\pm$ 2.5	36.55 $\pm$ 1.6	38.11 $\pm$ 0.7	35.40
Llama-3.2-3B-Instruct										
GT-Reward	10.62 $\pm$ 2.1	0.21 $\pm$ 0.4	47.5 $\pm$ 1.4	78.60 $\pm$ 0.8	22.89 $\pm$ 1.9	7.05 $\pm$ 0.6	32.48 $\pm$ 0.8	50.16 $\pm$ 1.8	34.26 $\pm$ 0.7	31.53
Majority-Voting	8.33 $\pm$ 1.3	0.00 $\pm$ 0.0	48.4 $\pm$ 3.0	78.64 $\pm$ 1.4	21.69 $\pm$ 2.8	8.11 $\pm$ 0.5	30.83 $\pm$ 1.4	48.36 $\pm$ 1.8	33.94 $\pm$ 0.7	30.92
Proof-Gate + Zero Reward	8.33 $\pm$ 1.6	0.42 $\pm$ 0.6	47.8 $\pm$ 2.2	78.54 $\pm$ 0.8	20.48 $\pm$ 3.2	4.36 $\pm$ 0.5	30.70 $\pm$ 1.1	49.98 $\pm$ 1.7	35.06 $\pm$ 0.7	30.63
Proof-Gate + Random Reward	2.92 $\pm$ 1.4	1.25 $\pm$ 1.1	41.85 $\pm$ 2.2	72.71 $\pm$ 2.9	17.32 $\pm$ 2.1	3.03 $\pm$ 0.8	26.88 $\pm$ 1.1	52.30 $\pm$ 1.8	31.90 $\pm$ 0.7	27.80
Proof-Gate + MV Reward	7.29 $\pm$ 1.6	0.00 $\pm$ 0.0	47.25 $\pm$ 3.3	78.28 $\pm$ 0.6	21.99 $\pm$ 2.6	7.07 $\pm$ 1.0	30.87 $\pm$ 0.7	48.49 $\pm$ 1.7	33.68 $\pm$ 0.7	30.55
JURY-RL (Proof-Gate + ResZero)	9.17 $\pm$ 2.0	1.25 $\pm$ 0.9	48.8 $\pm$ 2.6	78.96 $\pm$ 0.5	26.05 $\pm$ 1.6	6.16 $\pm$ 0.6	31.60 $\pm$ 1.2	50.09 $\pm$ 1.7	34.54 $\pm$ 0.7	31.85
Qwen2.5-7B										
GT-Reward	18.33 $\pm$ 2.3	11.25 $\pm$ 1.7	75.1 $\pm$ 1.4	90.83 $\pm$ 0.5	46.54 $\pm$ 2.5	12.78 $\pm$ 1.6	53.67 $\pm$ 1.0	41.50 $\pm$ 1.7	45.09 $\pm$ 0.7	43.90
Majority-Voting	11.25 $\pm$ 2.1	4.17 $\pm$ 1.2	71.0 $\pm$ 0.7	90.52 $\pm$ 0.7	38.70 $\pm$ 1.3	18.37 $\pm$ 0.9	52.20 $\pm$ 1.1	42.72 $\pm$ 1.7	43.83 $\pm$ 0.7	41.42
Proof-Gate + Zero Reward	11.67 $\pm$ 1.7	10.21 $\pm$ 2.9	75.15 $\pm$ 1.8	90.13 $\pm$ 0.5	41.11 $\pm$ 3.0	18.10 $\pm$ 1.2	52.02 $\pm$ 1.8	43.67 $\pm$ 1.7	46.97 $\pm$ 0.8	43.23
Proof-Gate + Random Reward	8.75 $\pm$ 1.9	2.50 $\pm$ 1.2	60.7 $\pm$ 2.0	82.66 $\pm$ 1.7	30.72 $\pm$ 4.5	21.97 $\pm$ 1.2	43.50 $\pm$ 1.1	38.56 $\pm$ 1.7	43.90 $\pm$ 0.7	37.03
Proof-Gate + MV Reward	0.00 $\pm$ 0.0	0.00 $\pm$ 0.0	51.4 $\pm$ 1.1	90.54 $\pm$ 1.0	14.76 $\pm$ 1.0	17.29 $\pm$ 1.0	54.83 $\pm$ 1.5	43.64 $\pm$ 1.7	35.01 $\pm$ 0.7	34.16
JURY-RL (Proof-Gate + ResZero)	13.75 $\pm$ 1.8	7.50 $\pm$ 1.6	75.65 $\pm$ 0.7	90.35 $\pm$ 0.9	47.29 $\pm$ 3.2	14.69 $\pm$ 1.1	54.37 $\pm$ 0.4	41.51 $\pm$ 1.7	50.00 $\pm$ 0.7	43.90

Table 3: Ablation results for the proposed ResZero reward ($\delta = 0$) on reasoning benchmark.

task knowledge benchmarks. On Qwen2.5-7B, JURY-RL again surpasses the GT baseline on these out-of-domain benchmarks, achieving an average of **40.14%** (+1.88 pts, +4.92% rel.). On the other two backbone models, its performance remains close to GT while consistently ranking as the best-performing label-free method. Across both in- and out-of-domain settings, these transferable gains and the aggregated averages in Table 1 show that JURY-RL is the strongest label-free method while remaining competitive with supervised or teacher-based baselines such as GT and LLM-KD.

Does JURY-RL Enhance Diversity? Table 2 shows that JURY-RL improves both pass@k

and pass@1 over GT-Reward and LLM-as-a-Judge on all three backbones. On average, the gains in pass@k under multi-sample decoding are substantially larger than those in pass@1, especially compared to LLM-as-a-Judge (e.g., +9.06 and +10.05 percentage points in pass@k on Qwen3-1.7B-Base and Qwen2.5-7B). This pattern indicates that JURY-RL increases the number of distinct high-quality solutions rather than only refining the




Figure 2: Average unique answers per sample over training steps on Qwen3-1.7B-Base.

single best prediction. Our **ResZero** reward contributes to this behavior by down-weighting unverifiable majority answers and redistributing reward to residual trajectories, which promotes exploration of alternative reasoning paths and mitigates mode collapse. Consistently, Figure 2 shows that the average number of unique answers per problem remains high throughout training for JURY-RL, whereas baselines such as Majority-Voting quickly collapse to a single dominant answer. Per-benchmark pass@k breakdowns are provided in Table 11, Appendix G.2.




Figure 3: Accuracy on MATH5000 validation set over training steps.

5.3 Analysis

Ablation Studies of ResZero. We compare four fallback designs under the same proof-gated framework when $\delta=0$: Zero Reward (no signal), Random Reward (rewarding a randomly selected candidate), MV Reward, and our proposed ResZero, with GT and Majority-Voting as references. As shown in Table 3, among all proof-gated variants ResZero consistently achieves the highest *Average*. On average, it outperforms Zero Reward by **+1.3** pts, MV Reward by **+6.1** pts, and essentially matches GT to within **0.1** pts, while surpassing the Random Reward baseline by **+5.4** pts. Notably, *Proof-Gate + MV Reward* performs worse than the plain Majority-Voting baseline. This highlights that rewarding an unverified majority can reinforce spurious consensus even after verification fails, accelerating mode collapse. A simple Zero Reward achieves suboptimal performance by taking a safe but inefficient path: it avoids reinforcing errors, but at the cost of stalling the learning process for inconclusive samples. Meanwhile, the naive Random Reward baseline, which injects unstructured reward noise when verification fails, underperforms Zero Reward by **4.1** pts on average, indicating that random reinforcement alone cannot provide a useful learning signal. ResZero provides a robust solution by navigating this trade-off between dangerous reinforcement and inefficient stagnation. Per-benchmark pass@k breakdowns can be found in Table 12, Appendix G.2, where ResZero significantly outperforms other variants and even surpasses the GT-Reward baseline across all models.

JURY-RL Achieves Stable Training and Avoids Collapse. To evaluate training stability, we tracked the validation accuracy of JURY-RL against label-free baselines on the MATH5000 validation set throughout the training process. As illustrated in Figure 3, Entropy and Self-Certainty

Verifier	Prec.	Rec.	F1
LLM-as-a-Judge	75.9	96.1	84.8
Lean Verifier (Ours)	84.5	88.0	86.2

Table 4: Verifier signal quality on training set. All metrics reported in percent (%).

Method	Average		AIME24		AIME25	
	pass@1	pass@k	pass@1	pass@k	pass@1	pass@k
MV-TTRL	46.23	57.38	10.00	20.00	3.33	23.33
JURY-RL	49.76	62.54	16.67	26.67	10.00	30.00

Table 5: Test-time RL on Qwen2.5-14B.

show collapse after an initial performance gain, as the model begins to reinforce spurious consensus. LLM-as-a-Judge, LLM-KD, and Majority-Voting exhibit noisy and suboptimal convergence. In contrast, JURY-RL demonstrates stable, monotonic improvement, confirming that its proof-gated reward mechanism effectively prevents the mode collapse common in self-supervised methods.

Analysis of Verifier Signal Quality. While Lean’s kernel is sound by construction, our deployed pipeline also runs upstream auto-formalization and consistency checks that can fail, so the practical question is whether the resulting signal is still cleaner than a learned judge. Table 4 answers this: our Lean verifier achieves substantially higher **precision** than the LLM-as-a-Judge (84.5% vs. 75.9%) at a modest cost in recall (88.0% vs. 96.1%), and a higher F1 (86.2% vs. 84.8%)—exactly the trade-off our “Proofs Dispose” principle prescribes, since false positives are what drive reward hacking under RL. We defer a breakdown of where the pipeline’s imperfections come from and how they shape training dynamics to Appendix G.5.

Test-Time RL Transfer. The proof-gated reward also generalizes to the *test-time* RLVR setting

(TTRL) (Zuo et al., 2025), where the model is adapted on the evaluation distribution itself. On Qwen2.5-14B, using each benchmark’s validation split as the TTRL environment, Table 5 shows that JURY-RL improves over the majority-vote TTRL baseline by **+3.53** pass@1 and **+5.16** pass@k on average, with the largest gains on the hardest competition benchmarks (AIME24 +6.67/ + 6.67, AIME25 +6.67/ + 6.67 in pass@1 / pass@k). The training dynamics mirror the offline regime: MV-TTRL first gains and then degrades as spurious consensus drives entropy collapse, whereas JURY-RL exhibits smoother and sustained improvement. Full per-benchmark numbers, training curves, and 7B results are in Appendix G.3.

6 Conclusion

We introduced **JURY-RL**, a label-free RLVR framework that decouples answer proposal from reward disposal: majority voting proposes a candidate, a Lean verifier gates positive reward, and **ResZero** supplies a zero-mean, variance-preserving signal when verification is inconclusive. Across math, code, and multi-task benchmarks, plus a test-time RL setting, JURY-RL trains more stably than label-free baselines, matches supervised ground-truth training on pass@1, and surpasses it on pass@k and diversity. Sparse but formally verified rewards can drive label-free RL without the false-positive failure modes that plague consensus- or judge-based surrogates.

Ethics Statement

This work studies label-free RLVR for mathematical and programmatic reasoning, and involves no human subjects or personally identifiable information. All benchmarks (MATH, GSM8K, AMC, AIME, LiveCodeBench, CRUX, MMLU-Pro, IFEval) are publicly available under their respective licenses. A potential concern is that stronger math reasoners can also be misused on graded coursework or contests.

Limitations

JURY-RL applies only to domains with a faithful machine-checkable verifier; we instantiate it on math via Lean. This is more reliable than majority voting or LLM judges, but rules out open-ended tasks where no such verifier exists. The verifier itself is a pipeline, not just the Lean kernel. Autoformalization and consistency checking sit upstream,

and their failures show up as inconclusive verification in our reward signal.

References

- Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, and Hao Peng. 2025. [The unreasonable effectiveness of entropy minimization in llm reasoning](#). *CoRR*, abs/2505.15134.
- Together AI and Agentica. 2025. [Deepcoder: A fully open-source 14b coder at o3-mini level](#). Blog post. Accessed: 2025-09-10.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, and 12 others. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.
- Jasmin C. Blanchette, Sascha Böhme, and Lawrence C. Paulson. 2011. [Extending Sledgehammer with SMT Solvers](#). In *Proceedings of the 23rd International Conference on Automated Deduction (CADE23)*, volume 6803 of *Lecture Notes in Artificial Intelligence*, pages 116–130. Springer.
- Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. [Humans or LLMs as the judge? a study on judgement bias](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8301–8327, Miami, Florida, USA. Association for Computational Linguistics.
- Guoxin Chen, Jing Wu, Xinjie Chen, Wayne Xin Zhao, Ruihua Song, Chengxi Li, Kai Fan, Dayiheng Liu, and Minpeng Liao. 2025. [Reform: Reflective autoformalization with prospective bounded sequence optimization](#). *CoRR*, abs/2510.24592.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *CoRR*, abs/2110.14168.
- Leonardo de Moura and Sebastian Ullrich. 2021. [The Lean4 Theorem Prover and Programming Language](#). In *Automated Deduction—CADE28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 81 others.

2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *CoRR*, abs/2501.12948.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 80 others. 2024. [Deepseek-v3 technical report](#). *CoRR*, abs/2412.19437.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and et al. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Leo Gao, John Schulman, and Jacob Hilton. 2022. [Scaling laws for reward model overoptimization](#). *CoRR*, abs/2210.10760.
- Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I. Wang. 2024a. [Cruxeval: A benchmark for code reasoning, understanding and execution](#). *CoRR*, abs/2401.03065.
- Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2024b. [Minillm: Knowledge distillation of large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the MATH dataset. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum. 2025. [Openreasoner-zero: An open source approach to scaling up reinforcement learning on the base model](#). *CoRR*, abs/2503.24290.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025a. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Transactions on Information Systems*, 43(2):42:1–42:55.
- Yuzhen Huang, Weihao Zeng, Xingshan Zeng, Qi Zhu, and Junxian He. 2025b. [Pitfalls of rule- and model-based verifiers - A case study on mathematical reasoning](#). *CoRR*, abs/2505.22203.
- Hugging Face H4. 2024. AIME 2024 Benchmark. https://huggingface.co/datasets/HuggingFaceH4/aime_2024.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. [Livecodebench: Holistic and contamination free evaluation of large language models for code](#). *CoRR*, abs/2403.07974.
- Kimi-Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, and 75 others. 2025. [Kimi k1.5: Scaling reinforcement learning with llms](#). *CoRR*, abs/2501.12599.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Taffjord, Chris Wilhelm, Luca Soldaini, and 4 others. 2024. [Tulu 3: Pushing frontiers in open language model post-training](#). *CoRR*, abs/2411.15124.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, and Sushant Prakash. 2024. [Rlaif vs. rlhf: Scaling reinforcement learning from human feedback with ai feedback](#). In *Proceedings of the 41st International Conference on Machine Learning*, pages 26874–26901. PMLR.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. [Let’s verify step by step](#). *CoRR*, abs/2305.20050.
- Yong Lin, Shange Tang, Bohan Lyu, Ziran Yang, Jui-Hui Chung, Haoyu Zhao, Lai Jiang, Yihan Geng, Jiawei Ge, Jingruo Sun, Jiayun Wu, Jiri Gesi, Ximing Lu, David Acuna, Kaiyu Yang, Hongzhou Lin, Yejin Choi, Danqi Chen, Sanjeev Arora, and Chi Jin. 2025. [Goedel-prover-v2: Scaling formal theorem proving with scaffolded data synthesis and self-correction](#). *Preprint*, arXiv:2508.03613.
- math-ai Team. 2024. Amc23: American mathematics competitions 2023 test set. <https://huggingface.co/datasets/math-ai/amc23>. Hugging Face dataset; 40 AMC 2023 problems; accessed 2025-09-16.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- OpenCompass. 2025. AIME 2025 Benchmark. <https://huggingface.co/datasets/opencompass/AIME2025>.
- 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 Miller, Maddie Simens, Amanda Askell, Peter Welinder,

- Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *CoRR*, abs/2203.02155.
- Jing-Cheng Pang, Pengyuan Wang, Kaiyuan Li, Xiong-Hui Chen, Jiacheng Xu, Zongzhang Zhang, and Yang Yu. 2023. [Language model self-improvement by reinforcement learning contemplation](#). *CoRR*, abs/2305.14483.
- Avinash Patil and Aryan Jadon. 2025. [Advancing reasoning in large language models: Promising methods and approaches](#). *arXiv preprint arXiv:2502.03671*.
- Zhongyuan Peng, Yifan Yao, Kaijing Ma, Shuyue Guo, Yizhe Li, Yichi Zhang, Chenchen Zhang, Yifan Zhang, Zhouliang Yu, Luming Li, Minghao Liu, Yihang Xia, Jiawei Shen, Yuchen Wu, Yixin Cao, Zhaoxiang Zhang, Wenhao Huang, Jiaheng Liu, and Ge Zhang. 2025. [Criticlean: Critic-guided reinforcement learning for mathematical formalization](#). *CoRR*, abs/2507.06181.
- G rard P.Huet, Gilles Kahn, and Christine Paulin-Mohring. 1997. [The Coq Proof Assistant: A Tutorial](#). Technical report, INRIA.
- Mihir Prabhudesai, Lili Chen, Alex Ippoliti, Katerina Fragkiadaki, Hao Liu, and Deepak Pathak. 2025. [Maximizing confidence alone improves reasoning](#). *CoRR*, abs/2505.22660.
- QwenTeam. 2025. [Qwq-32b: Embracing the power of reinforcement learning](#).
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Z. Z. Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanxia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, Z. F. Wu, Zhibin Gou, Shirong Ma, Hongxuan Tang, Yuxuan Liu, Wenjun Gao, Daya Guo, and Chong Ruan. 2025. [Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition](#). *CoRR*, abs/2504.21801.
- Sheikh Shafayat, Fahim Tajwar, Ruslan Salakhutdinov, Jeff Schneider, and Andrea Zanette. 2025. [Can large reasoning models self-train?](#) *CoRR*, abs/2505.21444.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *CoRR*, abs/2402.03300.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. [Hybridflow: A flexible and efficient RLHF framework](#). In *Proceedings of the Twentieth European Conference on Computer Systems, EuroSys 2025, Rotterdam, The Netherlands, 30 March 2025 - 3 April 2025*, pages 1279–1297. ACM.
- Jiawen Shi, Zenghui Yuan, Yinyao Liu, Yue Huang, Pan Zhou, Lichao Sun, and Neil Zhenqiang Gong. 2025. [Optimization-based prompt injection attack to llm-as-a-judge](#). *Preprint*, arXiv:2403.17710.
- Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu. 2025. [Crossing the reward bridge: Expanding RL with verifiable rewards across diverse domains](#). *CoRR*, abs/2503.23829.
- Aman Singh Thakur, Kartik Choudhary, Venkat Srinik Ramayapally, Sankaran Vaidyanathan, and Dieuwke Hupkes. 2025. [Judging the judges: Evaluating alignment and vulnerabilities in llms-as-judges](#). *Preprint*, arXiv:2406.12624.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *International Conference on Learning Representations*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhao Chen. 2024. [Mmlu-pro: A more robust and challenging multi-task language understanding benchmark](#). *CoRR*, abs/2406.01574.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *CoRR*, abs/2201.11903.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, and et al. 2025a. [Qwen3 technical report](#). *CoRR*, abs/2505.09388.
- 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, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 24 others. 2025b. [Qwen2.5 technical report](#). *CoRR*, abs/2412.15115.

Yongshi Ye, Liang Zhang, Yidong Chen, Xiaodong Shi, and Biao Fu. 2026. [Don’t peek at the answer: Outcome-masked group relative policy optimization for label-free rlvr](#). *Preprint*, arXiv:2608.03119.

Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Hongchang Gao, Arman Kulkarni, Binyuan Hui, Lei Li, Zhengyuan Xu, Minhao Jiang, Qi Liu, Xiang Zhu, Di Chen, Longjun Li, Wensen Cheng, Yu Zhang, Tianqi Zhang, Hao Zhang, Chao Shang, and 15 others. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). *CoRR*, abs/2503.14476.

Zizhuo Zhang, Jianing Zhu, Xinmu Ge, Zihua Zhao, Zhanke Zhou, Xuan Li, Xiao Feng, Jiangchao Yao, and Bo Han. 2025. [Co-reward: Self-supervised reinforcement learning for large language model reasoning via contrastive agreement](#). *CoRR*, abs/2508.00410.

Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. 2025a. [Absolute zero: Reinforced self-play reasoning with zero data](#). *CoRR*, abs/2505.03335.

Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. 2025b. [Learning to reason without external rewards](#). *CoRR*, abs/2505.19590.

Yulai Zhao, Haolin Liu, Dian Yu, S. Y. Kung, Haitao Mi, and Dong Yu. 2025c. [One token to fool llm-as-a-judge](#). *Preprint*, arXiv:2507.08794.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-following evaluation for large language models](#). *CoRR*, abs/2311.07911.

Yuxin Zuo, Kaiyan Zhang, Shang Qu, Li Sheng, Xuekai Zhu, Biqing Qi, Youbang Sun, Ganqu Cui, Ning Ding, and Bowen Zhou. 2025. [TTRL: Test-time reinforcement learning](#). *CoRR*, abs/2504.16084.

Organization of the Appendix. For navigation, Appendix A presents the theoretical analysis (stability of JURY-RL and the zero-mean property of ResZero). Appendix B gives the full algorithm and a worked example of reward computation. Appendix C details the Lean Verifier (components, cost trade-offs, and full prompt templates). Appendices D–F cover baselines, hyperparameters, and evaluation metrics. Appendix G provides extended empirical analyses (sensitivity, pass@k, and test-time training results). Appendix H presents a qualitative case study comparing an LLM-as-a-Judge with the formal verifier.

A Theoretical Analysis

This section gives a minimal, checkable account of why Jury-RL stabilizes policy optimization.

- **§A.1 — Why naive majority voting is brittle.** We show that, under GRPO, majority-vote rewards conflate agreement with correctness: as the majority share increases, supporters’ advantages go to zero while dissenters’ penalties blow up, driving entropy collapse.
- **§A.2 — ResZero is zero-mean with non-degenerate variance.** We prove the group reward strictly sums to zero with the choice $\gamma = c\alpha^2$, and that the group variance is non-zero whenever residual answers are not all identical, so gradients remain informative when verification fails (Eq. 4).
- **§A.3 — Fallback comparison (MV / Zero / ResZero).** We derive group-normalized advantages for each fallback and show that only ResZero yields a corrective (negative for unverifiable majorities) yet exploratory (variance preserved on residuals) update, aligning with the stability observed in §5.3.

A.1 Why naive majority voting is brittle.

We retain the previously derived analysis showing that majority voting (MV) conflates agreement with correctness and induces entropy collapse under GRPO as consensus strengthens. Using MV as a label-free reward conflates single-view agreement with correctness and induces entropy collapse. For a question x , sample G rollouts $y_i \sim \pi_\theta(\cdot | x)$ and parse answers $a_i = \text{ans}(y_i)$. Let $\hat{a} = \arg \max_a \sum_{j=1}^G \mathbb{I}[a_j = a]$ with vote share $\hat{v} = \frac{1}{G} \sum_{j=1}^G \mathbb{I}[a_j = \hat{a}]$. MV assigns binary rewards $r_i^{\text{MV}} = \mathbb{I}[a_i = \hat{a}]$. Under GRPO, group-normalized advantages satisfy $\bar{r} = \hat{v}$, $\text{std} = \sqrt{\hat{v}(1 - \hat{v})}$, hence

$$\hat{A}_i = \frac{r_i^{\text{MV}} - \bar{r}}{\text{std}} = \begin{cases} \sqrt{\frac{1-\hat{v}}{\hat{v}}}, & a_i = \hat{a}, \\ -\sqrt{\frac{\hat{v}}{1-\hat{v}}}, & a_i \neq \hat{a}. \end{cases}$$

As a spurious consensus strengthens ($\hat{v} \rightarrow 1$), supporters receive vanishing signal ($\hat{A}_i^+ \rightarrow 0$) while dissenters incur diverging penalties ($\hat{A}_i^- \rightarrow -\infty$), suppressing exploration and shrinking the token-level entropy toward a single mode; ratio clipping preserves the sign of $\hat{A}_i$, so the collapse persists.

Because r_i^{MV} ignores ground truth, MV is also vulnerable to formatting hacks: repeated insertion of a frequent symbol in the “answer box” can maximize agreement without correctness.

A.2 Proof of Zero-Mean Property for ResZero Reward

Proposition (Zero-mean property of ResZero). Consider any group of $G \geq 1$ trajectories, with majority set M , residual set R , and majority share $\alpha = |M|/G$. Let the ResZero reward r_i^{ResZero} be defined as in Equation (4) with $\gamma = c\alpha^2$. Then, for all possible vote patterns (including ties and degenerate cases),

$$\sum_{i=1}^G r_i^{\text{ResZero}} = 0.$$

This property is critical for ensuring stable optimization for different RL methods.

Edge cases. (i) *Tie-breaking* When multiple answers tie for the most frequent, we select a single $\hat{a}$ by a fixed rule, so $M = \{i : a_i = \hat{a}\}$ and $R = \{i : a_i \neq \hat{a}\}$ are uniquely defined. (ii) *No residuals* If all G trajectories coincide, then $|R| = 0$, $|M| = G$, and the residual sum is zero. The total reward becomes $\sum_{i \in M} r_i^{\text{ResZero}} = |M|(-c\alpha + \gamma)$. With $\alpha = |M|/G = 1$ and $\gamma = c\alpha^2 = c$, this equals $G(-c + c) = 0$. (iii) *Single dissenter* ($|R| = 1$) With a single residual trajectory, $\alpha = (G - 1)/G$ and $\bar{u}$ coincides with the unique z_i , so the dissenter receives $r = \alpha(z_i - \bar{u}) + \gamma = \gamma = c\alpha^2 > 0$, while each of the $|M| = G - 1$ majority members receives $-c\alpha + \gamma = c\alpha(\alpha - 1) < 0$. The total sum is $c\alpha^2 + (G - 1)c\alpha(\alpha - 1) = c\alpha[\alpha + (G - 1)(\alpha - 1)] = c\alpha[G\alpha - (G - 1)] = 0$, using $G\alpha = G - 1$. Thus the zero-mean property also holds in the minimal-residual case, in which a lone dissenter is positively credited while the over-represented majority is uniformly down-weighted.

We now present the proof for the general case where $|R| > 0$. The total sum of rewards is decomposed as:

$$\sum_{i=1}^G r_i^{\text{ResZero}} = \sum_{i \in M} r_i^{\text{ResZero}} + \sum_{i \in R} r_i^{\text{ResZero}}$$

Step 1: Calculate the sum of rewards for the residual group (R). For any trajectory $i \in R$, the reward is $r_i^{\text{ResZero}} = \alpha \cdot (z_i - \bar{u}) + \gamma$. Summing

over all members of the residual group:

$$\begin{aligned} \sum_{i \in R} r_i^{\text{ResZero}} &= \sum_{i \in R} [\alpha(z_i - \bar{u}) + \gamma] \\ &= \alpha \sum_{i \in R} (z_i - \bar{u}) + \sum_{i \in R} \gamma \end{aligned}$$

By the definition of a mean, the sum of deviations from the mean is zero, i.e., $\sum_{i \in R} (z_i - \bar{u}) = 0$. This simplifies the total reward for the residual group to:

$$\sum_{i \in R} r_i^{\text{ResZero}} = \alpha \cdot 0 + |R| \cdot \gamma = |R|\gamma$$

Step 2: Calculate the sum of rewards for the majority group (M). For any trajectory $i \in M$, the reward is $r_i^{\text{ResZero}} = -c\alpha + \gamma$. Since the reward is identical for all members of the majority group, the sum is:

$$\sum_{i \in M} r_i^{\text{ResZero}} = |M| \cdot (-c\alpha + \gamma)$$

Step 3: Combine the sums from both groups. We add the sums from Step 1 and Step 2 to find the total sum:

$$\begin{aligned} \sum_{i=1}^G r_i^{\text{ResZero}} &= |M|(-c\alpha + \gamma) + |R|\gamma \\ &= -c\alpha|M| + (|M| + |R|)\gamma \end{aligned}$$

Since $|M| + |R| = G$, the equation becomes:

$$\sum_{i=1}^G r_i^{\text{ResZero}} = -c\alpha|M| + G\gamma$$

Step 4: Derivation of the Global Re-centering Term γ . To enforce the zero-mean property, we design the global re-centering term γ such that the total sum from Step 3 is identically zero.

$$\begin{aligned} -c\alpha|M| + G\gamma &= 0 \\ \gamma &= \frac{c\alpha|M|}{G} \end{aligned}$$

By substituting the definition of the majority share, $\alpha = |M|/G$, we arrive at the required form for γ :

$$\gamma = c\alpha \left(\frac{|M|}{G} \right) = c\alpha^2$$

This derivation shows that setting $\gamma = c\alpha^2$ is the precise design choice required to make the ResZero reward strictly zero-centered. This term acts as a global offset that exactly balances the penalties applied to the majority group and the rewards distributed among the residual group.

A.3 Fallback Rewards for Inconclusive Verification

This section provides a theoretical analysis of different fallback reward mechanisms within the GRPO framework for the scenario where formal verification of the majority-voted answer is inconclusive ($\delta = 0$). The analysis focuses on the dynamics of the group-normalized advantage, $\hat{A}_i$, which serves as the learning signal for the policy update. The advantage is defined as:

$$\hat{A}_i = \frac{r_i - \bar{r}}{\text{std}(\{r_j\}_{j=1}^G) + \varepsilon}, \quad \text{where} \quad \bar{r} = \frac{1}{G} \sum_{j=1}^G r_j.$$

We analyze three cases: rewarding the majority vote (MV), assigning a zero reward, and using our proposed ResZero reward.

Case 1: Majority Voting (MV) Reward

As established in Appendix A.1, if we naively reward the majority consensus when verification fails, the reward is $r_i^{\text{MV}} = \mathbb{I}[a_i = \hat{a}]$. Let $\hat{v}$ be the vote share of the majority answer $\hat{a}$. The key statistics are $\bar{r} = \hat{v}$ and $\text{std}(\{r_j\}) = \sqrt{\hat{v}(1 - \hat{v})}$. This yields the following advantage:

$$\hat{A}_i = \begin{cases} \sqrt{\frac{1-\hat{v}}{\hat{v}}}, & \text{if } a_i = \hat{a} \text{ (Supporter)} \\ -\sqrt{\frac{\hat{v}}{1-\hat{v}}}, & \text{if } a_i \neq \hat{a} \text{ (Dissenter)} \end{cases}$$

Theoretical Implication. As a spurious consensus strengthens ($\hat{v} \rightarrow 1$), the advantage for supporters vanishes ($\hat{A}_i^+ \rightarrow 0$) while the penalty for dissenters diverges ($\hat{A}_i^- \rightarrow -\infty$). This dynamic punishes any exploration and provides no positive signal for adhering to the consensus, leading to **entropy collapse** and policy degradation.

Case 2: Zero Reward

A seemingly safe alternative is to assign a zero reward to all rollouts when verification is inconclusive. In this case, $r_i^{\text{Zero}} = 0$ for all $i \in \{1, \dots, G\}$.

The resulting statistics are trivial:

- **Mean Reward:** $\bar{r} = \frac{1}{G} \sum_{i=1}^G 0 = 0$.
- **Standard Deviation:** $\text{std}(\{r_j\}) = \sqrt{\frac{1}{G} \sum_{i=1}^G (0 - 0)^2} = 0$.

Substituting these into the advantage formula gives:

$$\hat{A}_i = \frac{0 - 0}{0 + \varepsilon} = 0$$

Theoretical Implication. The advantage signal is nullified for all rollouts in the group. This leads to a **vanishing gradient** for the policy update, effectively stalling the learning process for that entire batch. While it avoids the destructive collapse of MV, it does so at the cost of learning efficiency, rendering the update step ineffective.

Case 3: ResZero Reward

The ResZero reward is designed specifically to address the shortcomings of the above methods. It has two crucial properties by construction:

- **Zero-Mean Property:** The total reward across the group sums to zero, i.e., $\sum_{i=1}^G r_i^{\text{ResZero}} = 0$. This immediately implies the mean reward is $\bar{r} = 0$.
- **Non-Zero Variance:** By assigning a negative reward to the majority group and a structured, zero-mean reward to the residual group, r_i^{ResZero} is non-zero for most rollouts (unless all residual answers are identical). Therefore, $\text{std}(\{r_j\}) > 0$ as long as there is diversity among residual answers.

The advantage function thus becomes:

$$\hat{A}_i = \frac{r_i^{\text{ResZero}} - 0}{\text{std}(\{r_j\}) + \varepsilon} = \frac{r_i^{\text{ResZero}}}{\text{std}(\{r_j\}) + \varepsilon}$$

Theoretical Implication. This formulation establishes a stable, self-regulating update dynamic that avoids the extremes of vanishing gradients (Case 2) or diverging penalties (Case 1). Intuitively, ResZero operates as a **damped correction mechanism**:

- **Adaptive Penalization:** If the model is *uncertain* (small majority share α), the corrective signal is mild. However, if the model is *confidently wrong* (large α but unverified), both the majority penalty and residual amplification scale up. This specifically targets spurious consensus without destabilizing early training.
- **Stability against Oscillations:** Unlike a binary flip, ResZero redistributes probability mass toward residuals proportional to their internal support z_i (as defined in §A.2). This creates a zero-mean, variance-preserving gradient that gently pushes the policy toward a more balanced mixture rather than inducing unbounded oscillations.

This mechanism explains the stability observed in our experiments (§5.3): ResZero provides a directional signal to undo unverified confidence while GRPO ensures the updates remain bounded.

In summary, our theoretical analysis reveals that each fallback strategy results in a fundamentally different learning dynamic:

- **Majority Voting (MV)** leads to a **destructive update**. As a spurious consensus strengthens, it creates a diverging penalty for dissent while the positive signal for supporters vanishes. This dynamic ultimately causes **entropy collapse**, suppressing exploration and degrading the policy.
- **Zero Reward** results in an **ineffective update**. By nullifying the reward for all rollouts, the advantage signal becomes zero for the entire group. This causes a **vanishing gradient** that stalls the learning process for that step, wasting computational resources.
- **ResZero Reward** provides a **constructive update**. It maintains a stable, non-zero learning signal that is both **corrective**, by penalizing the unverified majority consensus, and **exploratory**, by rewarding promising alternatives among the residual answers.

B Implementation Details

B.1 The JURY-RL Algorithm

The full algorithm is summarized in Algorithm 1.

Algorithm 1 JURY-RL (one grouped update for a prompt x)

- 1: Sample G rollouts $y_i \sim \pi_{\text{old}}(\cdot \mid x)$; parse $a_i = \text{ans}(y_i)$.
 - 2: Compute vote shares $v(a)$ and majority $\hat{a} \in \arg \max_a v(a)$.
 - 3: Query verifier once: $\delta = \text{verify}(x, \hat{a})$.
 - 4: **if** $\delta=1$ **then**
 - 5: Set $r_i = \mathbb{I}[a_i = \hat{a}]$.
 - 6: **else**
 - 7: Form M, R ; compute $u^{(-i)}$ and $\bar{u}$; set r_i via Eq. (4).
 - 8: **end if**
 - 9: Compute group-normalized advantages $\hat{A}_i$ from $\{r_i\}_{i=1}^G$; broadcast token-wise.
 - 10: Update π_θ with GRPO (clipped ratios, KL to reference).
-

B.2 An Intuitive Example of the ResZero Reward

Consider a scenario with $G = 8$ rollouts for a given problem. A majority vote reveals that $|M| = 4$ rollouts support a proposal $\hat{a}$, so the majority share is $\alpha = |M|/G = 4/8 = 0.5$. The remaining $|R| = 4$ residual rollouts are split: three support answer b , and one supports answer c . If the verifier is inconclusive ($\delta = 0$), the ResZero reward is activated.

First, we analyze the residual group R . For a trajectory y_i with answer $a_i = b$, its leave-one-out support within the residual group is $z_i = u^{(-i)}(b) = \frac{2}{|R|-1} = 2/3$. For the singleton answer $a_i = c$, the support is $z_i = u^{(-i)}(c) = 0/3 = 0$. The average support across the residual group is therefore $\bar{u} = \frac{1}{|R|}(3 \cdot \frac{2}{3} + 1 \cdot 0) = 0.5$.

Let the penalty hyperparameter be $c = 0.1$. The global re-centering term is $\gamma = c\alpha^2 = 0.1 \cdot (0.5)^2 = 0.025$. The final reward for each trajectory is assigned according to Eq. (4):

$$r_i^{\text{ResZero}} = \begin{cases} -c\alpha + \gamma = -0.025, & i \in M \\ \alpha(z_i - \bar{u}) + \gamma \approx 0.1083, & i \in R, a_i = b \\ \alpha(z_i - \bar{u}) + \gamma = -0.225, & i \in R, a_i = c \end{cases}$$

By design, the total reward sums exactly to zero: $4(-0.025) + 3(\frac{1}{12} + 0.025) + 1(-0.225) = -0.1 + 0.25 + 0.075 - 0.225 = 0$. This demonstrates how ResZero penalizes the unverifiable majority while redistributing a zero-mean signal among diverse residual answers. This maintains a useful, variance-driven learning gradient for GRPO without reinforcing a potentially spurious consensus.

C Lean Verifier Details

C.1 Lean Verifier

Lean (de Moura and Ullrich, 2021) has emerged as a transformative framework in the formal verification of mathematical proofs, grounded in a rigorous type-theoretic foundation that guarantees unprecedented levels of logical soundness and mechanized reliability. Its intrinsic dependency on computer-assisted compilation environments not only ensures formal correctness but also serves as a high-fidelity feedback mechanism for refining and validating mathematical reasoning within LLMs.

In response to this potential, we have designed and implemented a comprehensive mathematical

verification system centered on Lean. This system bridges the gap between informal natural language mathematics and machine-verifiable formalism by automatically translating problem statements and proposed solutions into syntactically and semantically well-formed Lean expressions, followed by formal proof verification via Lean’s trusted kernel.

The architecture of this verification system (illustrated in Figure 4) is a cascaded, modular pipeline comprising three specialized components:

- **Autoformalizer** Translates natural language mathematical content (question—answer pairs in our setting) into precise, executable Lean formal specifications, preserving both syntactic structure and semantic intent.
- **Consistency-Checker** Performs bidirectional semantic alignment between the original input of natural language and its formalized Lean counterpart, ensuring fidelity of meaning and detecting potential misinterpretations or translation artifacts.
- **Prover** Synthesizes formal proof scripts in Lean and submits them to the theorem prover for mechanized validation, thus certifying the logical correctness of the solution under the foundational logic of the system.

This framework not only advances the automation of mathematical reasoning verification, but also establishes a scalable, feedback-driven paradigm for integrating formal methods into natural language-based mathematical systems.

C.2 Setting

To ensure reliable and precise reward signals, our Lean Verifier employs a three-stage inference pipeline composed of the following core components:

- **Autoformalizer:** Utilizes the Goedel-Formalizer-V2-32B model (Lin et al., 2025) to translate informal mathematical statements into formal Lean specifications.
- **Consistency-Checker:** Implemented using QwQ-32B (QwenTeam, 2025), this component validates the outputs of the Autoformalizer.
- **Prover:** Employs the Goedel-Prover-V2-32B model (Lin et al., 2025) to synthesize formal

proofs for the specifications that have been validated in the previous stage.

The pipeline operates as follows: first, the Autoformalizer generates up to 8 candidate formalizations. The Consistency-Checker then evaluates effective candidates and selects one checked candidate. Finally, the Prover conducts up to 16 independent sampling trials to find a valid proof.

All formal verification outputs were generated under the same hyperparameter configuration below:

- **Temperature:** 0.7
- **Maximum Tokens:** 32,768

We make this configuration choice to balance exploration, fidelity, and resource efficiency.

C.3 Component-wise Accuracy Analysis

To better understand the reliability of our reward signal, we evaluated the independent performance of the Autoformalizer and Consistency Checker on the MATH500 dataset.

Autoformalizer Reliability. We evaluated the autoformalization success rate. The Autoformalizer successfully generated syntactically valid and type-checkable Lean 4 statements for **477 out of 500 problems**, yielding a success rate of **95.4%**. This high coverage ensures that the vast majority of mathematical problems can be converted into a verifiable format.

Consistency Checker Accuracy. The Consistency Checker is responsible for filtering out formalizations that are syntactically valid but semantically inconsistent with the original problem. As shown in Table 6, the model achieves an F1-score of **91.1%** and a Recall of **94.9%**, demonstrating its effectiveness in maintaining semantic fidelity without aggressively discarding valid candidates.

Held-out Verifier Signal Quality. To address the concern that the MATH500 evaluation above overlaps with the training distribution, we additionally evaluate both components on external held-out benchmarks that were not used to train or tune any part of the pipeline. Table 7 reports the Consistency Checker on two out-of-domain benchmarks and Table 8 reports the Autoformalizer success rate on four external datasets spanning textbook, olympiad, and competition mathematics.

Two observations follow. First, the Consistency Checker maintains strong performance on out-of-domain data, with F1 above 84% on both external




Figure 4: The pipeline of our Lean Verifier: a collaborative system that combines autoformalization, formal verification, and consistency checking to validate mathematical reasoning with high reliability.

Metric	Accuracy	Precision	Recall	F1-score
Consistency Checker	85.3%	87.7%	94.9%	91.1%

Table 6: Performance metrics of our Consistency Checker.

Benchmark	Accuracy	Precision	Recall	F1-score
ReformBench-859(Chen et al., 2025)	80.4%	80.1%	94.2%	86.6%
CriticleanBench(Peng et al., 2025)	84.4%	82.1%	88.0%	84.9%

Table 7: Held-out evaluation of the Consistency Checker on two external benchmarks.

Benchmark	miniF2F	ProofNet	Putnam	AIME 2025	AVG
Success rate	97.1%	71.5%	74.2%	66.7%	77.4%

Table 8: Held-out success rate of the Autoformalizer on four external benchmarks.

benchmarks and recall remaining high (94.2% / 88.0%); the checker does not aggressively discard valid formalizations on unseen problems. Second, the Autoformalizer shows domain-dependent variation: near-perfect on miniF2F (97.1%) but lower on harder competition problems (Putnam 74.2%, AIME 2025 66.7%). Crucially, autoformalization failures are *conservative* errors: they produce $\delta = 0$ (inconclusive verification) and fall back to ResZero, rather than emitting false-positive rewards. Lower autoformalization coverage therefore reduces supervision density on the hardest domains without compromising reward precision, which is consistent with our precision-first design.

Combined with the formal soundness of the Lean Prover (which has a false positive rate of 0%), these components ensure that JURY-RL receives a reward signal that is both high-coverage and high-precision.

C.4 Performance and Cost Analysis

A crucial factor for the practical implementation of JURY-RL is the computational overhead introduced by the formal Lean Verifier, especially when compared to baselines like an LLM-as-a-Judge. To provide a clear and quantitative assessment of this cost, we conducted a wall-clock time analysis using the Qwen-2.5-7B model. The evaluation was performed under the synchronous RL training configuration specified in our main paper.

Initial (Cold-Start) Overhead. To accurately assess the computational cost, we first establish the baseline training time. The oracle setting with Ground-Truth (GT) rewards, which involves no external verification, requires approximately **100 seconds** per step. This duration represents the core training workload.

We then measured the additional verification overhead for the other methods under a worst-case

scenario, defined as processing a batch composed **entirely of unseen question-answer (QA) pairs**. The specific overheads are as follows:

- **JURY-RL (Lean Verifier):** Under the shared setting where each training step processes 128 questions and the verifier is called once per question on the majority-vote answer, the Lean-based verification adds an overhead of approximately **200 seconds per step**.
- **LLM-as-a-Judge (Qwen-2.5-72B):** Using the judge to score only the majority-vote answer for each of the 128 questions adds an overhead of approximately **80 seconds per step**.

This means that in the initial training stages, the total step time for JURY-RL is approximately 300 seconds (100s for training + 200s for verification), compared to 180 seconds for the LLM-as-a-Judge. It is important to note that the efficiency of JURY-RL is significantly bolstered by an **early-stopping mechanism**. Although our framework attempts to autoformalize and prove the statement up to k times (in the *pass@k* setting), the process terminates immediately upon the first success. Since we verify the consensus answer derived from majority voting—which has a high probability of correctness—the prover often succeeds in the very first few attempts, avoiding the computational cost of executing all k trials.

Cost Amortization and Convergence. The initial cold-start overhead is not representative of the average cost over the entire training process. Our framework uses a caching mechanism for all verification results (shown in Fig. 4). As training proceeds and the policy begins to converge, the diversity of generated answers for any given problem stabilizes. Consequently, an increasing fraction of QA pairs in subsequent batches will have been previously encountered and verified. These cached results can be retrieved almost instantly, bypassing the expensive formal verification process. In the steady-state phase of training, this caching effect becomes dominant. The per-step time cost for JURY-RL progressively converges toward that of the GT baseline, as most verification lookups are resolved via the cache. Therefore, while the initial cost of the Lean Verifier is higher, this cost is effectively amortized throughout the training run. We contend that this represents an acceptable and practical trade-off for the substantial gains in reward

fidelity, training stability, and final model performance documented in our main results.

Comparison with LLM-as-a-Judge. Beyond the runtime overhead, another critical aspect of the verifier is the trade-off between its performance and consumption, which directly impacts both the quality of the reward signal and the operational cost. To analyze this, we evaluated different prover configurations on the MATH500 dataset and compared them with an LLM-as-a-judge baseline.

Table 9 summarizes the performance, model size, and runtime characteristics of these configurations. Among the Lean-based verifiers, the Prover@64 configuration achieves the highest accuracy with the largest end-to-end evaluation time on MATH500 (655.80 s). In contrast, the Prover@16 configuration, while exhibiting only a slight and acceptable decrease in performance to 87.6% ACC, reduces the total runtime to 324.31 s. Given this highly favorable trade-off between a marginal performance drop and substantial cost savings, we adopt Prover@16 as the default Lean-based configuration for our main experiments. For comparison, Table 9 also reports an LLM-as-a-judge variant based on Qwen2.5-72B-instruct with four-way majority voting (mv@4). This judge uses a substantially larger model (72B versus 32B) and runs at a much lower normalized throughput (6.39 versus 102.82 tokens/s in A100-equivalent units). Despite requiring far fewer output tokens per response, its end-to-end verification time on MATH500 (290.30 s) is therefore comparable to that of Prover@16 (324.31 s), rather than dramatically faster.

Moreover, unlike the Lean-based verifier, the LLM-as-a-judge approach provides no formal guarantees and yields a purely heuristic natural-language scoring signal, which is more susceptible to inconsistency and reward hacking. Our empirical results show that the formally verified Prover@16 offers substantially higher reward fidelity and more stable training dynamics for JURY-RL. Taking both wall-clock runtime and supervision quality into account, we view its additional token cost as a practical and acceptable investment in verifier quality.

C.5 Prompt

We list the prompts for all three models below for reference.

Setting	Model Size	MATH500			Token Costs	Runtime	
		ACC	TPR	F1-Score	Avg Tokens per Response	Eqv. Speed (tokens/s)	Time Cost (s)
Prover@16	32B	87.6	87.6	93.4	5,858	102.82	324.31
Prover@32	32B	89.1	89.1	94.2	6,607	102.82	506.36
Prover@64	32B	91.1	91.1	95.3	6,435	102.82	655.80
LLM-as-a-judge (MV@4)	72B	87.2	87.2	93.2	464	6.39	290.30

Table 9: Performance, token cost, and runtime analysis of Lean Verifier and LLM-as-a-judge on MATH500.

Autoformalizer Prompt

Please autoformalize the following natural language problem statement in Lean 4.
 Use the following theorem name: test_problem
 The natural language statement is:
 {informal_statement_content}.
 Think before you provide the lean statement.

Consistency-Checker System Prompt

Your role is a Lean4 expert, please help me check consistency between natural language expression and its Lean4 proof statement.
 Guidelines for Consistency Checking:
 Goal:
 Determine if the Lean theorem statement is an exact and faithful formalization of the mathematical problem.
 Do not evaluate or consider the answer or the proof. Your sole task is to verify the correctness of the formalization.
 Evaluation Stages (All required):
 1. Math Assertion Analysis
 Identify all structurally and semantically relevant components of the mathematical problem, including variables, types, quantifiers, constraints, logic structure, conclusion, and so on. The analysis should be based on the actual content of the text.
 2. Lean Statement Analysis (ignore proof part)
 Extract all structurally and semantically relevant components from the Lean statement, including variables, types, conditions, quantifiers, constraints, the final claim, and so on. The analysis should reflect the actual content present in the Lean code.
 3. Comparative Verification
 Check for exact correspondence between the math and Lean statements; you may refer to aspects like:
 - Semantic alignment, logic structure, and quantifier correctness.

- Preservation of constraints and boundary assumptions.
- Accurate typing and use of variables.
- Syntactic validity and proper Lean usage (free from errors).
- Use of symbols and constructs without semantic drift.
- No missing elements, no unjustified additions, and no automatic corrections or completions.

4. Final Judgement

Based solely on the above analysis, judge whether the Lean statement is a correct and exact formalization of the mathematical problem.

5. Accuracy Confirmation

If correct: clearly confirm why all elements match.

If incorrect: list all mismatches and explain how each one affects correctness.

Note: While the analysis may be broad and open to interpreting all relevant features, the final judgment must be based only on what is explicitly and formally expressed in the Lean statement. ****Do not consider or assess any part of the proof. Your judgment should be entirely about the accuracy of the statement formalization.****

You should present the results following the format:

Input:

The Natural Language Statement:

A math problem and its answer (no proof).

The Formal Statement in Lean4:

```
```lean
```

A Lean 4 theorem statement formalizing the problem. Proof is intentionally omitted (e.g., sorry).

```
```
```

Output Format:

Return exactly one xml object

```
<comments>
```

Your brief analysis:

Math Assertion Analysis: [...]
 Lean Statement Analysis (Proof Ignored): [...]
 Comparative Verification: [...]
 Conclusion: [...]
 Accuracy Confirmation: [... match confirmation or list of discrepancies...]
 </comments>
 <consistency>[Correct/Incorrect]</consistency>

Consistency-Checker User Prompt

Input Data:
 The Natural Language Statement:
 {informal_prefix}
 The Formal Statement in Lean4:
 ```lean  
 {formal\_statement}  
 ```

Prover Prompt

Complete the following Lean 4 code:
 ```lean4  
 {formal\_statement}  
 ```  
 Before producing the Lean 4 code to formally prove the given theorem, provide a detailed proof plan outlining the main proof steps and strategies.
 The plan should highlight key ideas, intermediate lemmas, and proof structures that will guide the construction of the final formal proof.

D Baseline Details

- **Majority Voting (MV).** A self-supervised consensus reward (Shafayat et al., 2025). For each problem, we generate G rollouts. A rollout y_i receives reward 1 if its extracted answer $\text{ans}(y_i)$ matches the majority answer among the G rollouts, and 0 otherwise. This directly reinforces popular answers regardless of correctness.
- **Self-Certainty.** A confidence-derived signal (Zhao et al., 2025b). The reward is computed from the log-probabilities of the tokens composing the final answer (as extracted by $\text{ans}(\cdot)$); higher cumulative log-probability indicates greater model certainty and yields a higher reward.
- **Entropy Minimization.** A low-entropy proxy (Prabhudesai et al., 2025). The reward is in-

versely related to the policy’s output entropy for the final-answer tokens, encouraging more deterministic, high-confidence predictions.

- **CoReward.** Contrastive agreement over input-level variants (Zhang et al., 2025). For each training question, CoReward generates semantically equivalent paraphrased versions, samples multiple rollouts for both the original and paraphrased questions, and aggregates their answers via majority voting to obtain pseudo-consensus labels. These majority-voted answers are then used in a cross-over manner as rewards for policy optimization on both sides.
- **Ground Truth (GT).** The ground-truth supervised oracle baseline. Using human-annotated labels, a rollout receives 1 if its extracted answer matches the ground truth and 0 otherwise. We train this baseline with the same GRPO objective (Shao et al., 2024) as our method for a fair comparison.
- **LLM-as-a-Judge.** An external-judge paradigm (Pang et al., 2023; Zhao et al., 2025a). We employ qwen-2.5-72b-instruct as the judge. It assesses the reasoning process and the final answer of each rollout; its evaluation (numeric score or a binary correct/incorrect label) is used as the reward for the RL update. The prompt for it is shown below.

- **LLM-KD (Knowledge Distillation).** Teacher–student distillation has been used in Label-Free RLVR (Zhao et al., 2025a). We use qwen-2.5-72b-instruct to produce an answer for each problem and treat it as a pseudo label. The policy model is trained to align its outputs with these machine-generated references.

LLM-as-a-Judge prompt

You are a professional math QA pair evaluator. Your sole task is to determine whether a given mathematical question and answer are correctly matched. First explain your reasoning, then end your response with your final judgment: True or False.

E Additional Experimental Details

Detailed training and testing settings are provided in Table 10.

Hyperparameter	Value
<i>Training Configuration</i>	
Train Batch size	128
Rollouts per problem (G)	8
Max prompt length	512
Max new tokens	3072
Training epoch	6
<i>Optimizer Parameters (AdamW)</i>	
Learning rate	3×10^{-6}
β_1	0.9
β_2	0.999
ϵ	10^{-8}
Warmup style	Cosine
Warmup steps ratio	0.1
<i>GRPO Algorithm Parameters</i>	
KL loss coefficient (β)	0.005
clip ratio (ϵ)	0.2
<i>Generation Parameters</i>	
Training temperature	1.0
Evaluation temperature	0.8
top-p	0.95

Table 10: Reinforcement learning training hyperparameters. This configuration is consistently applied across all experiments to ensure a fair comparison.

F Benchmark and Metric Details

This section details the specific metrics and software frameworks used for each benchmark mentioned in the main text.

- **AIME24/25, MATH500, GSM8K:** We report **avg@16** accuracy for **AIME24/25**, computed by averaging correctness over 16 independently sampled solutions per problem, and **avg@4** accuracy for **MATH500** and **GSM8K**. All benchmarks are evaluated using the lighteval framework: <https://github.com/huggingface/lighteval>
- **AMC:** We report **avg@8** accuracy. This is evaluated using the ttrl implementation: https://github.com/ruixin31/Spurious_Rewards/tree/main/code/ttrl

- **LiveCodeBench:** We report **avg@5** accuracy. This is evaluated using the official library: <https://github.com/LiveCodeBench/LiveCodeBench>

- **CRUX:** We report **avg@5** accuracy. This is evaluated using the ZeroEval framework: <https://github.com/WildEval/ZeroEval>

- **MMLU-Pro, IFEval:** We report **pass@1** accuracy. This is evaluated using the lm-evaluation-harness: <https://github.com/EleutherAI/lm-evaluation-harness>

G Further Analysis

G.1 Impact of c

We analyze the impact of the hyperparameter c in Eq. 4, which controls the penalty strength on the unverified majority proposal in our ResZero reward. As illustrated in Figure 5, c is critical in navigating the trade-off between preventing mode collapse and maximizing task performance. The right panel clearly demonstrates that a non-zero c is essential for maintaining solution diversity. When $c = 0$, the framework effectively degenerates to a zero-reward fallback, leading to a sharp decline in the average number of unique answers as training progresses—a classic symptom of the model converging to a spurious consensus. In contrast, any positive c value successfully sustains a high level of diversity. However, the left and center panels reveal a subtle trade-off: an overly aggressive penalty (e.g., $c = 0.1$) can slightly suppress the final reward and accuracy. This suggests that while the penalty is crucial for exploration, an excessive value may overly restrict the policy from exploiting a potentially correct, high-consensus answer. This ablation validates that a moderately tuned c (e.g., $c = 0.01$ in our experiments) strikes an optimal balance, ensuring robust training stability and solution diversity without compromising convergence on the primary task objectives.

G.2 Pass@k Results

Main Results across Benchmarks. This appendix complements The Diversity Analysis of JURY-RL in Section 5.2. The complete numerical results are listed in Table 11 for easy cross-reference. The pass@k results presented in Table 11 unequivocally demonstrate the effectiveness of our JURY-RL framework. Across all three




Figure 5: Training dynamics under different values of the hyperparameter c .

Methods	AIME24 pass@16	AIME25 pass@16	MATH500 pass@4	GSM8K pass@4	AMC pass@8	Average
<i>Qwen3-1.7B-Base</i>						
Before RL	16.67	13.33	72.0	89.31	56.63	49.59
GT-Reward	26.67	16.67	82.0	92.42	59.04	55.36
LLM-KD	20.00	16.67	82.2	93.25	55.42	53.51
Entropy	20.00	23.33	81.2	90.75	59.04	54.86
Self-Certainty	23.33	26.67	77.2	90.83	57.83	55.17
Majority-Voting	23.33	16.67	77.2	92.34	53.01	52.51
CoReward	26.67	16.67	77.4	92.04	54.22	53.40
LLM-as-a-Judge	16.67	13.33	74.6	91.74	55.42	50.35
JURY-RL (Ours)	30.00	30.00	82.0	92.42	62.65	59.41
<i>Llama-3.2-3B-Instruct</i>						
Before RL	23.33	6.67	65.0	86.88	49.40	46.26
GT-Reward	23.33	6.67	63.8	90.14	43.37	45.46
LLM-KD	23.33	3.33	65.0	89.99	49.40	46.21
Entropy	16.67	3.33	58.8	84.46	50.60	42.77
Self-Certainty	13.33	10.00	58.0	84.69	44.58	42.12
Majority-Voting	16.67	0.00	63.6	90.98	39.76	42.20
CoReward	23.33	0.00	60.4	90.45	43.37	43.51
LLM-as-a-Judge	26.67	10.00	62.2	89.84	40.96	45.93
JURY-RL (Ours)	30.00	3.33	64.8	90.07	54.22	48.48
<i>Qwen2.5-7B</i>						
Before RL	30.00	20.00	76.2	93.78	63.86	56.77
GT-Reward	33.33	30.00	85.2	95.22	68.67	62.48
LLM-KD	33.33	33.33	87.8	95.75	67.47	63.54
Entropy	30.00	36.67	85.2	93.78	67.47	62.62
Self-Certainty	23.33	30.00	84.6	93.86	69.88	60.33
Majority-Voting	33.33	30.00	85.4	95.00	63.86	61.52
CoReward	30.00	26.67	83.8	95.60	63.86	59.99
LLM-as-a-Judge	20.00	13.33	82.2	95.38	59.04	53.99
JURY-RL (Ours)	36.67	30.00	86.6	95.83	71.08	64.04

Table 11: Pass@k results (%) of RL performance comparison on math reasoning benchmarks.

backbone models, JURY-RL not only surpasses every label-free baseline but also consistently outperforms the strong supervised GT-Reward baseline, which is trained directly on ground-truth answers.

Ablation Study on Fallback Rewards. This appendix complements Ablation Studies of ResZero in Section 5.2. The complete numerical results are listed in Table 12 for easy cross-reference with Table 3. Table 12 demonstrates that our proposed

ResZero fallback reward consistently achieves the highest average performance.

G.3 Test-Time Training Results

While our main experiments focus on offline RLVR, recent work has emphasized *test-time* reinforcement learning (TTRL) (Zuo et al., 2025), where the model is adapted on the evaluation distribution itself. To examine whether the conclusions

Methods	AIME24 pass@16	AIME25 pass@16	MATH500 pass@4	GSM8K pass@4	AMC pass@8	Average
<i>Qwen3-1.7B-Base</i>						
GT-Reward	26.67	16.67	82.0	92.42	59.04	55.36
Majority-Voting	23.33	16.67	77.2	92.34	53.01	52.51
Proof-Gate + Zero Reward	23.33	6.67	70.2	92.65	46.99	47.97
Proof-Gate + Random Reward	20.00	6.67	78.0	92.34	54.22	50.25
Proof-Gate + MV Reward	0.00	0.00	48.8	92.19	13.25	30.85
JURY-RL (Proof-Gate + ResZero)	30.00	30.00	82.0	92.42	62.65	59.41
<i>Llama-3.2-3B-Instruct</i>						
GT-Reward	23.33	6.67	63.8	90.14	43.37	45.46
Majority-Voting	16.67	0.00	63.6	90.98	39.76	42.20
Proof-Gate + Zero Reward	20.00	6.67	64.0	89.92	40.96	44.31
Proof-Gate + Random Reward	23.33	6.67	62.6	88.32	46.99	45.58
Proof-Gate + MV Reward	20.00	0.00	62.2	89.76	44.58	43.31
JURY-RL (Proof-Gate + ResZero)	30.00	3.33	64.8	90.07	54.22	48.48
<i>Qwen2.5-7B</i>						
GT-Reward	33.33	30.00	85.2	95.22	68.67	62.48
Majority-Voting	33.33	30.00	85.4	95.00	63.86	61.52
Proof-Gate + Zero Reward	30.00	33.33	85.6	95.38	62.65	61.39
Proof-Gate + Random Reward	33.33	16.67	79.6	93.63	63.86	57.42
Proof-Gate + MV Reward	0.00	0.00	58.2	95.60	19.28	34.62
JURY-RL (Proof-Gate + ResZero)	36.67	30.00	86.6	95.83	71.08	64.04

Table 12: Ablation results (Pass@k) for the proposed ResZero reward ($\delta = 0$) on math reasoning benchmark.

in §5.3 also hold in this regime, we instantiate a TTRL variant of JURY-RL and compare it with the majority-vote (MV) reward used in the public `ttr1` implementation. For each benchmark (AIME24, AIME25, AMC23, MATH500 and GSM8K), we treat the validation split as the TTRL environment, start from the same base policy, and run GRPO for 30 epochs with the hyperparameters in Table 10. We run this setup on two backbones: Qwen2.5-7B (Figure 6) and Qwen2.5-14B (Table 13 and Figure 7); in both cases, the only difference between the two runs is the reward: MV vs. our ResZero fallback. To keep the 14B comparison aligned with the 7B figure, we contrast JURY-RL with the majority-vote baseline (MV-TTRL) only.

Results on Qwen2.5-7B. On Qwen2.5-7B, the dynamics on AIME24 and AMC highlight the brittleness of MV in the test-time setting: the validation reward first increases slightly and then degrades, eventually falling below the initial performance, whereas JURY-RL continues to make steady progress with a smoother learning curve, consistent with our analysis in Appendix A.1 that spurious consensus under MV drives entropy collapse and hurts generalization. For AIME25 in

particular, the first points in Figure 6 correspond to 0/30 (JURY-RL) and 1/30 (MV) solved problems; since both runs start from the same base model, this small gap is due to stochasticity in GRPO rather than an inherent disadvantage of JURY-RL. As training proceeds, JURY-RL attains a more stable plateau on AIME25 and still ends slightly ahead of MV, while on MATH500 and GSM8K the two curves remain close but JURY-RL consistently finishes higher.

Results on Qwen2.5-14B. The same pattern holds at 14B scale. As summarized in Table 13, JURY-RL improves AVG pass@1 by **+3.53** points and AVG pass@k by **+5.16** points over MV-TTRL on Qwen2.5-14B, with the largest gains on the hardest competition-style benchmarks (AIME24 pass@1 +6.67, AIME25 pass@1 +6.67, AMC pass@k +10.84). The training curves in Figure 7 follow the same overall pattern observed at 7B: JURY-RL exhibits smoother and more sustained improvement than MV-TTRL on AIME24, AIME25, and AMC, whereas MV-TTRL again shows the characteristic initial gain followed by degradation. Taken together, these results indicate that the benefit of JURY-RL in the test-time setting is not a




Figure 6: **Test-Time Training Dynamics on Qwen2.5-7B.** We plot exponentially moving-averaged (EMA) validation accuracy for Jury-RL (purple) and majority-voting TTRL (orange) over 400 test-time training steps. Jury-RL exhibits more stable and continued improvement on challenging benchmarks.




Figure 7: **Test-Time Training Dynamics on Qwen2.5-14B.** We plot exponentially moving-averaged (EMA) validation accuracy for JURY-RL (purple) and majority-voting TTRL (orange) over 150 test-time training steps. JURY-RL exhibits more stable and continued improvement on challenging benchmarks.

Method	AVG pass@1	AVG pass@k	AIME24	AIME25	AMC	GSM8K	MATH500
MV-TTRL	46.23	57.38	10.00 / 20.00	3.33 / 23.33	44.58 / 62.65	94.62 / 95.53	78.60 / 85.40
JURY-RL	49.76	62.54	16.67 / 26.67	10.00 / 30.00	46.99 / 73.49	94.54 / 96.74	80.60 / 85.80

Table 13: Test-time training results on Qwen2.5-14B.

7B-specific phenomenon and transfers to 14B under a matched setup.

G.4 Verifier Branch Dynamics over Training

To characterize how the proof-gated and fallback branches of JURY-RL evolve over training, we tracked two diagnostics on a Qwen2.5-7B run with group size $G = 8$ under the main experimental configuration in Table 10. All measurements use benchmark ground truth as the reference rather than pipeline verifiability, to avoid circular evaluation.

Evolution of the proof-gated fraction $P(\delta)$. At checkpoints $\{0, 50, 100, 200, 400\}$, the fraction of training steps in which the plurality answer is formally verified by the Lean pipeline ($\delta = 1$) rises monotonically, while the fallback fraction $P(\delta = 0)$ falls correspondingly (Table 14). This is consistent with our design intent: ResZero acts as an early- and mid-stage stabilizer when inconclusive verification is common, and is progressively phased out as proof-gated updates become dominant.

Checkpoint	0	50	100	200	400
$P(\delta = 1)$ (%)	53.0	56.7	60.0	66.2	70.7
$P(\delta = 0)$ (%)	47.0	43.3	40.0	33.8	29.3

Table 14: Evolution of the proof-gated fraction $P(\delta = 1)$ and the fallback fraction $P(\delta = 0)$ over training checkpoints on Qwen2.5-7B ($G = 8$).

Minority-correct missed rate E_t . A concern for proof-gated RLVR is that a correct minority answer is overlooked when the plurality candidate is wrong and is selected for verification. To quantify this, we define

$$E_t = \Pr[\text{plurality answer is GT-incorrect} \\ \wedge \exists \text{ a minority rollout that is GT-correct}].$$

At the same checkpoints $\{0, 50, 100, 200, 400\}$, E_t is reported in Table 15. These missed opportunities are real but not dominant: E_t peaks at 9.3% during the mid-training phase and then declines to 5.4% by checkpoint 400, as the correct answer increasingly coincides with the plurality candidate (the




Figure 8: Training dynamics of precision, recall, validation accuracy, and training entropy under different Lean $pass@k$ verification settings.

Checkpoint	0	50	100	200	400
E_t (%)	7.6	8.7	9.3	7.4	5.4

Table 15: Evolution of the minority-correct missed rate E_t over training checkpoints on Qwen2.5-7B ($G = 8$).

majority-correct rate rises from 50.9% to 67.9% over the same training window). Among the remaining groups in which the plurality is wrong, roughly 78–85% contain no correct rollout at all, which explains why verifying the top-2 or top-3 candidates instead of top-1 yields little additional benefit (Figure 8).

G.5 Deeper Dive into Verifier Imperfections

While a formal verifier is theoretically sound, its practical application is imperfect. We analyze how the number of verification attempts, k in a $pass@k$ setting, affects the quality of the reward signal. As shown in Figure 8, the Lean verifier exhibits a highly desirable trade-off.

The **precision** (top-left panel) remains consistently high, around 85%. It is crucial to note that this precision gap from a perfect 100% is not a flaw in the Lean prover’s core logic, which is

formally sound. Instead, it primarily stems from imperfections in upstream components like **auto-formalization and consistency checks**, which translate the problem into a verifiable format. Despite this, the signal’s high fidelity is vital for preventing reward hacking and ensuring the model learns from genuinely correct examples.

In contrast, the **recall** (top-right panel) shows a clear dependency on k . With a single attempt ($pass@1$), the recall is modest (around 30%), but it steadily increases with more attempts, reaching nearly 90% at $pass@16$. This indicates that while a single verification attempt might fail to prove a correct answer (e.g., due to search limits), multiple attempts significantly increase the chance of success, thereby improving the richness of the training signal.

Crucially, despite the wide variance in recall, the final **validation accuracy** (bottom-left) and **training entropy** (bottom-right) converge to similar stable states across different values of k . This suggests that the high precision of the verifier’s signal is the dominant factor for successful and stable training. This behavior stands in stark contrast to an LLM-as-a-Judge. **Mechanistically**, an LLM-judge’s potential for error is inherent to its prob-




Figure 9: Accuracy on MATH5000 Validation set over training steps.

abilistic and opaque reasoning process, making it fundamentally prone to biases, format gaming, and confident mistakes. Its reward signal is therefore inherently noisier and less reliable. The imperfections in our pipeline, however, are largely confined to the modular pre-processing steps, leaving the core judgment by the Lean prover itself principled and trustworthy.

Effect of top- k verification on performance. To quantify how many candidates actually need to be verified, we compare JURY-RL training runs on Qwen2.5-7B that verify only the top-1, top-2, or top-3 majority-voting candidates selected by the policy, and report accuracy on the MATH5000 validation set in Figure 9. All three curves reach a similar plateau around 72% accuracy, and they remain almost indistinguishable throughout training, with final differences well within random fluctuation. This indicates that almost all useful learning signal is already contained in the highest-probability candidate, and additional verified candidates mainly increase computational cost without improving final accuracy. Consequently, we adopt the more efficient top@1 verify setting as our default choice in all main experiments.

Verifier Precision on Effective Groups. Since groups with zero reward variance produce zero advantages under GRPO, overall verifier precision may not reflect the supervision that actually drives learning. We therefore evaluate precision on *effective groups*, i.e., groups with non-zero reward variance. As shown in Table 16, JURY-RL achieves 82.24% precision on effective groups, compared with 62.31% for MV plurality labels, yielding a substantially larger advantage (+19.93 pp) than on all groups (+7.03 pp). This shows that the proof-

	JURY-RL	MV	Δ
All groups	92.30	85.27	+7.03
Effective groups	82.24	62.31	+19.93

Table 16: Precision (%) of JURY-RL verifier signals and MV plurality labels. Effective groups have non-zero reward variance and thus contribute non-zero GRPO advantages.

gated signal remains substantially more reliable than unverified consensus on the samples that actually contribute to policy updates.

G.6 Effect of Proof-Gated Disposal

To disentangle the contributions of ResZero and proof-gated disposal, we replace the Proof-Gate with an LLM judge while keeping ResZero unchanged. As shown in Table 17, adding ResZero improves LLM-as-a-Judge from 43.84% to 44.58% on average, while replacing the LLM judge with the Proof-Gate further improves performance to 46.91% (+2.33 pp). These results indicate that both ResZero and proof-gated disposal contribute to JURY-RL.

H Case Study

To illustrate our “Proofs Dispose” principle, we present a case study comparing an LLM-as-a-Judge with a Lean formal verifier. This case highlights a core problem with LLM judges: their reliability is limited by their own reasoning. The first example shows that an LLM judge can reach a wrong conclusion due to flawed reasoning, which generates an incorrect training signal. In contrast, once the autoformalization and consistency checks are passed, a Lean proof has the property of zero false positives. As shown in the second example, it provides a definitive and highly reliable signal through formal logic. This comparison explains why in JURY-RL, we use proofs, not an LLM, to dispose the final reward.

Question

Let x and y be positive real numbers
 $\rightarrow$ such that $3x + 4y < 72$. Find the
 $\rightarrow$ maximum value of xy ($72 - 3x - 4y$).

Answer

1152

Method	AIME24	AIME25	MATH-500	GSM8K	AMC	Avg.
LLM-as-a-Judge	10.42	5.62	72.30	89.73	41.11	43.84
LLM-as-a-Judge + ResZero	11.67	5.62	73.45	90.30	41.87	44.58
JURY-RL (Proof-Gate + ResZero)	13.75	7.50	75.65	90.35	47.29	46.91

Table 17: Comparison of proof-gated and LLM-judge-based disposal with ResZero on Qwen2.5-7B.

LLM-as-a-Judge Output (Judge Wrong)

To evaluate whether the given answer of
 $\hookrightarrow$ 1152 is correct for the problem, we
 $\hookrightarrow$ need to find the maximum value of
 $\hookrightarrow$ the function $\backslash(f(x, y) = xy (72 - 3x - 4y) \backslash)$ subject to the
 $\hookrightarrow$ constraint $\backslash(3x + 4y < 72 \backslash)$.

First, let's rewrite the function:
 $\backslash[f(x, y) = xy (72 - 3x - 4y). \backslash]$

We can use the method of Lagrange
 $\hookrightarrow$ multipliers or analyze the function
 $\hookrightarrow$ directly. However, a simpler
 $\hookrightarrow$ approach is to use symmetry and
 $\hookrightarrow$ calculus.

Consider the function $\backslash(g(t) = t(72 - 3t - 4t) = t(72 - 7t) = 72t - 7t^2 \backslash)$.
 $\hookrightarrow$ This is a quadratic function in
 $\hookrightarrow$ terms of $\backslash(t \backslash)$, where $\backslash(t = x + y \backslash)$. The maximum value of this
 $\hookrightarrow$ quadratic function occurs at the
 $\hookrightarrow$ vertex, which is given by:
 $\backslash[t = -\frac{b}{2a} = -\frac{72}{-14} = \frac{36}{7} \backslash]$
 $\hookrightarrow \frac{36}{7}$

Now, we need to find $\backslash(x \backslash)$ and $\backslash(y \backslash)$
 $\hookrightarrow$ such that $\backslash(x + y = \frac{36}{7} \backslash)$
 $\hookrightarrow$ and $\backslash(3x + 4y \backslash)$ is minimized while
 $\hookrightarrow$ still satisfying $\backslash(3x + 4y < 72 \backslash)$.

Given the symmetry and the nature of the
 $\hookrightarrow$ problem, we can assume $\backslash(x = y \backslash)$.
 $\hookrightarrow$ Then:

$\backslash[2x = \frac{36}{7} \implies x = \frac{18}{7} \backslash]$
 $\hookrightarrow \frac{18}{7}$
 $\backslash[3x + 4y = 3\left(\frac{18}{7}\right) + 4\left(\frac{18}{7}\right) = 7\left(\frac{18}{7}\right) = 18. \backslash]$
 $\hookrightarrow 18$

Substituting $\backslash(x = y = \frac{18}{7} \backslash)$
 $\hookrightarrow$ into the function:
 $\backslash[f\left(\frac{18}{7}, \frac{18}{7}\right) = \left(\frac{18}{7}\right)\left(\frac{18}{7}\right)\left(72 - 3\left(\frac{18}{7}\right) - 4\left(\frac{18}{7}\right)\right) \backslash]$
 $\hookrightarrow \frac{18}{7} \cdot \frac{18}{7} \cdot (72 - \frac{54}{7} - \frac{72}{7}) \backslash]$
 $\hookrightarrow \frac{18}{7} \cdot \frac{18}{7} \cdot (72 - \frac{126}{7}) \backslash]$
 $\hookrightarrow \frac{18}{7} \cdot \frac{18}{7} \cdot (72 - 18) \backslash]$
 $\hookrightarrow \frac{18}{7} \cdot \frac{18}{7} \cdot 54 \backslash]$
 $\hookrightarrow \frac{324}{49} \cdot 54 \backslash]$
 $\hookrightarrow \frac{17496}{49} \backslash]$
 $\hookrightarrow 357.0612244898. \backslash]$

This value does not match 1152. Therefore,
 $\hookrightarrow$ the given answer of 1152 is
 $\hookrightarrow$ incorrect.

Final judgment: False.

Lean verifier output.

Lean Prover Output (Prove Success)

```
import Mathlib
import Aesop

set_option maxHeartbeats 0

open BigOperators Real Nat Topology Rat

theorem test_problem :
  sSup {z : ℝ | ∃ x y : ℝ, x > 0 ∧ y > 0 ∧
    3*x + 4*y < 72 ∧ z = x * y * (72 - 3*x - 4*y)} = 1152 := by
  have h1 : (1152 : ℝ) ∈ {z : ℝ | ∃ x y :
    ℝ, x > 0 ∧ y > 0 ∧ 3*x + 4*y < 72 ∧
    z = x * y * (72 - 3*x - 4*y)} := by
    refine' ⟨8, 6, by norm_num, by
      norm_num, by norm_num, _⟩
    <|> norm_num
    <|> ring_nf
    <|> norm_num

  have h2 : ∀ (z : ℝ), z ∈ {z : ℝ | ∃ x y
    : ℝ, x > 0 ∧ y > 0 ∧ 3*x + 4*y <
    72 ∧ z = x * y * (72 - 3*x - 4*y)}
    → z ≤ (1152 : ℝ) := by
    intro z hz
    rcases hz with ⟨x, y, hx, hy, hxy, rfl⟩
    have h3 : 0 < x := hx
    have h4 : 0 < y := hy
    have h5 : 3 * x + 4 * y < 72 := hxy
    have h6 : 0 < 72 - 3 * x - 4 * y := by
      linarith
    have h7 : 0 < 3 * x := by positivity
    have h8 : 0 < 4 * y := by positivity
    have h9 : 0 < (3 * x) * (4 * y) := by
      positivity
    have h10 : 0 < (3 * x) * (4 * y) * (72 - 3 * x - 4 * y) := by positivity
    -- Use AM-GM inequality to show that x
    -- * y * (72 - 3 * x - 4 * y) ≤ 1152
    have h11 : (3 * x + 4 * y + (72 - 3 * x - 4 * y)) ^ 3 ≥ 27 * (3 * x) * (4 * y) * (72 - 3 * x - 4 * y) := by
      nlinarith [sq_nonneg (3 * x - 4 * y),
        sq_nonneg (3 * x - (72 - 3 * x - 4 * y)),
        sq_nonneg (4 * y - (72 - 3 * x - 4 * y)),
        mul_nonneg h7.le h8.le, mul_nonneg
        h7.le h6.le, mul_nonneg h8.le h6.le]
```

```

-- Simplify the inequality to get the
↪ desired result
have h12 : (72 : ℝ) ^ 3 ≥ 27 * (3 * x)
↪ * (4 * y) * (72 - 3 * x - 4 * y) :=
↪ by
  have h13 : (3 * x + 4 * y + (72 - 3 *
↪ x - 4 * y)) = 72 := by ring
  rw [h13] at h11
  linarith
have h13 : x * y * (72 - 3 * x - 4 * y)
↪ ≤ 1152 := by
  have h14 : (72 : ℝ) ^ 3 = 373248 :=
↪ by norm_num
  have h15 : 27 * (3 * x) * (4 * y) *
↪ (72 - 3 * x - 4 * y) = 324 * (x * y
↪ * (72 - 3 * x - 4 * y)) := by
    ring
  rw [h14] at h12
  rw [h15] at h12
  nlinarith
linarith

have h3 : sSup {z : ℝ | ∃ x y : ℝ, x > 0
↪ ∧ y > 0 ∧ 3*x + 4*y < 72 ∧ z = x *
↪ y * (72 - 3*x - 4*y)} = 1152 := by
  have h4 : IsLUB {z : ℝ | ∃ x y : ℝ, x
↪ > 0 ∧ y > 0 ∧ 3*x + 4*y < 72 ∧ z =
↪ x * y * (72 - 3*x - 4*y)} 1152 := by
    refine' ⟨fun z hz => h2 z hz, fun y
↪ hy => _⟩
    -- We need to show that 1152 is the
↪ least upper bound.
    -- Since 1152 is in the set, any
↪ upper bound y must be at least 1152.
    have h5 : (1152 : ℝ) ∈ {z : ℝ | ∃ x
↪ y : ℝ, x > 0 ∧ y > 0 ∧ 3*x + 4*y <
↪ 72 ∧ z = x * y * (72 - 3*x - 4*y)}
↪ := h1
    have h6 : (1152 : ℝ) ≤ y := by
      apply hy
      exact h5
    linarith
  -- The supremum of the set is 1152
↪ because it is the least upper bound.
have h5 : sSup {z : ℝ | ∃ x y : ℝ, x >
↪ 0 ∧ y > 0 ∧ 3*x + 4*y < 72 ∧ z = x
↪ * y * (72 - 3*x - 4*y)} = 1152 := by
  apply le_antisymm
  · -- Show that the supremum is less
↪ than or equal to 1152.
    apply csSup_le
    · -- The set is nonempty.
      exact ⟨1152, h1⟩
    · -- Every element in the set is
↪ less than or equal to 1152.
      intro z hz
      exact h2 z hz
    · -- Show that 1152 is less than or
↪ equal to the supremum.
      apply le_csSup
      · -- The set is bounded above.
        use 1152
        intro z hz
        exact h2 z hz
      · -- 1152 is in the set.
        exact h1
  exact h5

exact h3

```