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

Anonymous authors

Paper under double-blind review

ABSTRACT

Reinforcement learning with verifiable rewards (RLVR) enhances the reasoning of large language models (LLMs), but its scalability is hampered by the high cost of human-annotated labels. Label-free alternatives, such as majority voting or LLM-as-a-judge, are susceptible to false positives that lead to reward hacking and training collapse. We introduce JURY-RL, a label-free RLVR framework that separates answer proposal from reward disposal: votes from model rollouts propose a consensus answer, while a formal theorem prover disposes the final reward. Specifically, a rollout is rewarded only if the majority-voted answer is formally verified by a Lean prover. When verification is inconclusive, we activate our proposed ResZero (Residual-Zero) reward: it drops the unverifiable majority proposal and assigns a zero-mean, variance-preserving reward to the remaining (residual) answers. This unique design maintains a stable optimization gradient for RL algorithms without reinforcing spurious consensus. Experiments across mathematical reasoning, code generation, and multi-task benchmarks show that JURY-RL not only achieves more stable training but also consistently outperforms label-free baselines and even matches or surpasses supervised training with ground-truth rewards across pass@1 and pass@k.

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 & Jadon, 2025; Huang et al., 2025b). Recent advances 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 path: rather than aligning to what looks plausible, the objective aligns to what is provably correct by using verifiable signals such as from program execution or mathematical equivalence (Cobbe et al., 2021; Yu et al., 2025; AI & Agentica, 2025). However, scaling RLVR is constrained by its reliance on human-annotated answers or carefully curated specifications, which is costly and limited in coverage (Ouyang et al., 2022; Bai et al., 2022; Shao et al., 2024).

To reduce labeling cost, recent work explores 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). These approaches can learn reasoning but are prone to false positives and training collapse, often via reward hacking: models learn to satisfy the surrogate while drifting from correctness, (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, they are vulnerable to instruction and format manipulation when reasoning steps are explicitly produced, and suffer high false positives when they are hidden. They are also sensitive to prompting and temperature, and introduce nontrivial compute overhead and shared-bias risks (Pang et al., 2023; Chen et al., 2024; Thakur et al., 2025; Shi et al., 2025; Zhao et al., 2025c; a).

The shortcomings of the 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 simultaneously 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, in one degree or another, failed to meet all 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

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 learning signal to correctness. Conversely, when verification is **Inconclusive** ($\delta = 0$), all rollouts receive the proposed ResZero (Residual-Zero) Reward.

majority vote. A formal theorem prover (de Moura & 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. A positive reward is granted only if the majority-voted answer is formally verified by a Lean verifier (de Moura & Ullrich, 2021), suppressing the false positives common in other self-supervised or judge-based methods (Shafayat et al., 2025; Gao et al., 2022). This raises a critical question: what happens when verification is inconclusive? A naive zero reward would stall learning, while rewarding the majority vote would reintroduce the risk of reinforcing errors. To solve this, we introduce the **ResZero** (**Residual-Zero**) reward, a novel fallback mechanism. ResZero discards the unverified majority proposal and assigns a carefully constructed, zero-mean reward to the remaining (residual) answers. This design maintains a stable optimization gradient for learning to proceed, without amplifying a potentially spurious consensus, ensuring both training stability and truth alignment.

Contributions. (1) We introduce JURY-RL, a novel label-free RLVR framework that operationalizes a "votes propose, proofs dispose" paradigm. By strategically verifying only the majority-voted candidate, it aligns rewards with provable correctness using a formal verifier, eliminating the need for human-annotated answers. (2) We design the ResZero (Residual-Zero) reward, a principled fallback mechanism for when verification is inconclusive. By discarding the unverifiable majority and assigning a zero-mean, variance-preserving reward to residual candidates, it ensures stable optimization and prevents collapse from spurious consensus. (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 or surpassing a strong supervised ground-truth baseline across pass@1 and pass@k.

2 RELATED WORK

Large Language Model Reasoning. The general capabilities of Large language models 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 such as 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 its 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 models 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 & Ullrich, 2021; Cobbe et al., 2021; AI & 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: votes propose a candidate, and a verifier disposes the reward. Verified cases pay positive reward only to supporting trajectories, while unverifiable cases return a centered, variance-preserving group signal that avoids amplifying spurious consensus. This proof-gated design preserves label-free scalability while maintaining optimization alignment with provable correctness.

3 PRELIMINARIES

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

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

which returns 1 if a is formally certified correct under standard soundness assumptions.

We optimize a KL-regularized RLVR objective with a reference policy π_{ref} and coefficient β :

$$\max_{\pi_{\theta}} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\theta}(\cdot \mid x)} \Big[r(x, y; \mathcal{G}_x) - \beta \operatorname{D}_{\mathrm{KL}} \big(\pi_{\theta}(\cdot \mid x) \parallel \pi_{\mathrm{ref}}(\cdot \mid x) \big) \Big], \tag{1}$$

where $r(\cdot)$ is a grouped reward computed from G rollouts $\mathcal{G}_x = \{y_i\}_{i=1}^G$ for the same input.

We adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024) to estimate group-normalized advantages. Concretely, sample $y_i \sim \pi_{\text{old}}(\cdot \mid x)$ for $i = 1, \dots, G$ and compute $r_i := r(x, y_i; \mathcal{G}_x)$. The scalar group advantage is

$$\hat{A}_i = \frac{r_i - \overline{r}}{\operatorname{std}(\{r_j\}_{j=1}^G) + \varepsilon}, \qquad \overline{r} = \frac{1}{G} \sum_{j=1}^G r_j.$$
 (2)

Let the per-token ratio be $\rho_{i,t}(\theta) = \frac{\pi_{\theta}(y_{i,t}|x,y_{i,< t})}{\pi_{\text{old}}(y_{i,t}|x,y_{i,< t})}$. GRPO maximizes

$$J_{\text{GRPO}}(\theta) = \mathbb{E}\left[\frac{1}{G}\sum_{i=1}^{G}\frac{1}{|y_i|}\sum_{t=1}^{|y_i|}\min(\rho_{i,t}(\theta)\hat{A}_{i,t}, \operatorname{clip}(\rho_{i,t}(\theta), 1-\epsilon, 1+\epsilon)\hat{A}_{i,t}) - \beta D_{\text{KL}}(\pi_{\theta}||\pi_{\text{ref}})\right], \quad (3)$$

where $\hat{A}_{i,t}$ is a broadcast of \hat{A}_i (or any per-token variant compatible with GRPO), and the clipping avoids excessive policy drift. This presentation mirrors established RLVR practice for fair comparison and faithful reproduction.

Label-free Self-Reward in Reinforcement Learning. We categorize existing methods for label-free self-reward based on the origin of the reward signal: model-internal signals versus those from an external judge.

(A) Model-Internal Self-Reward Signals. These methods derive rewards from the model's own outputs without external evaluators. (i) **Output-side proxies**, such as entropy minimization or confidence-based scores (Agarwal et al., 2025; Prabhudesai et al., 2025), reward hypotheses that exhibit high certainty. However, this approach is fragile, as it can amplify errors when the model becomes confidently wrong. (ii) **Single-view agreement** rewards consistency among multiple outputs generated from the same input x (Shafayat et al., 2025). Specifically, for G responses $y_{i_{i=1}}^G$, it identifies the majority-voted answer $a_v = \arg\max_a \sum_{i=1}^G \mathbb{I}[ans(y_i) = a]$ and assigns a positive reward $r_{sv}(x,y) = a$

 $\mathbb{I}[\operatorname{ans}(y) = a_v]$ to responses that match a_v . The primary risk is reinforcing an erroneous consensus, where the model converges on a popular but incorrect answer, often by exploiting superficial heuristics such as formatting conventions. (iii) **Multi-view agreement** attempts to improve robustness by enforcing consistency across multiple, semantically equivalent prompts. For instance, the majority answer from prompt variant x' is used as a pseudo-label to supervise responses from the original prompt x (Zhang et al., 2025). This often improves training stability. However, it usually only delays rather than eliminates hacking, since spurious shortcuts can eventually propagate across multiple views.

(B) External-Judge Signals. This paradigm uses a powerful, external LLM as an automated judge to score the model's outputs. **LLM-as-a-Judge** (Lee et al., 2024; Su et al., 2025; Huang et al., 2025a; Zhao et al., 2025a) remains labelfree (no human annotation) but comes with distinct trade-offs. On the one hand, it mitigates the self-confirmation bias of internal methods. On the other hand, it introduces sensitivity to the judge's prompt design and decoding strategy, incurs significant computational cost, and risks transferring the judge's intrinsic biases into the training signal. Another approach is **LLM-based Knowledge Distillation (LLM-KD)** (Gu et al., 2024b), where a teacher model generates a reference answer to guide the student model. While potentially offering a more granular signal, it is similarly constrained by the teacher's capabilities and biases. Thus, while external judges reduce some weaknesses of the internal proxies, they bring a different set of reliability and scalability challenges.

Positioning Our Work. JURY-RL is a label-free RLVR method that *proof-gates* reward: a majority-voted answer is rewarded only if a formal verifier certifies it. When verification is inconclusive, we drop the majority and apply **ResZero**—a zero-mean, variance-preserving residual reward, so that GRPO maintains a stable gradient without reinforcing spurious consensus. This contrasts with the self-reward approaches that reward popularity or confidence, and with LLM-as-a-Judge/KD approaches that are prompt-sensitive and prone to false positives; JURY-RL aligns learning to verifiable correctness while remaining label-free.

4 JURY-RL

 JURY-RL is designed to satisfy three core principles: (i) **scalability** without costly human supervision, (ii) **truth alignment** by grounding rewards in verifiable evidence, and (iii) **optimization stability**, ensuring continuous learning even when verification is inconclusive. We achieve this by decoupling the process of proposing a candidate answer from the final disposal of its reward. Votes from the policy's own rollouts serve as a scalable proposal mechanism, while a formal theorem prover acts as a reliable judge for reward disposal.

This design choice—verifying only the single, majority-voted candidate—is crucial for maintaining computational tractability and scalability. Performing formal verification on every unique answer generated across all rollouts would be prohibitively expensive, undermining the efficiency of the label-free approach. The majority vote thus acts as an effective heuristic to identify the most promising candidate for the costly-but-reliable verification process.

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

4.1 The Proof-Gated Reward Framework

The JURY-RL workflow begins with a **proposal stage**. For a given problem x, we generate G trajectories $\{y_i\}_{i=1}^G \sim \pi_{\theta}(\cdot|x)$ and parse their corresponding answers $a_i = \operatorname{ans}(y_i)$. A majority vote determines the most frequent answer, which becomes our candidate proposal:

$$\hat{a} = \arg\max_{a} \sum_{i=1}^{G} \mathbb{I}[a_i = a]$$

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

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

$$r_{i} = \underbrace{\delta \cdot \mathbb{I}[a_{i} = \hat{a}]}_{\text{Verified Correctness}} + \underbrace{(1 - \delta) \cdot r_{i}^{\text{ResZero}}}_{\text{Inconclusive Fallback}}, \tag{4}$$

where r_i^{ResZero} is the Residual-Zero Reward detailed in Section 4.2.

¹Details of our Lean verifier can be found in Appendix C.

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 directly 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 our ResZero fallback rather than naively rewarding the majority consensus. This carefully centered substitute maintains a stable optimization gradient by preserving group-wise variance, which prevents learning from stalling or oscillating, a common issue when verification fails for benign reasons like search limits or incomplete libraries. Finally, by paying only for verifiable correctness, this framework inherently narrows the attack surface for prompt and format hacking, a critical 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 directly 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 spurious consensus strengthens. A simple zero-reward fallback is also suboptimal, as it would lead to a zero group-wise advantage, effectively stalling 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:

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

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:

$$r_{i}^{\text{ResZero}} = \underbrace{\alpha \cdot \mathbb{I}[i \in R] \cdot \left(z_{i} - \bar{u}\right)}_{\text{Amplify residual signals}} - \underbrace{c\alpha \cdot \mathbb{I}[i \in M]}_{\text{Penalize majority}} + \underbrace{\gamma}_{\text{Global re-centering}},$$

$$\text{where} \quad \bar{u} = \frac{1}{|R|} \sum_{j \in R} z_{j}, \quad \text{and} \quad \gamma = c\alpha^{2}.$$
(5)

Here, c is a positive hyperparameter controlling the penalty strength. γ is designed as 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 $(\sum_i r_i^{\rm ResZero} = 0)$, preserving the zero-mean property crucial for GRPO stability. ²

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. Collectively, 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 (δ =1), the policy update is guided by verifiable correctness. Conversely, if inconclusive (δ =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.

²See Appendix A.2 for a formal proof and Appendix B.2 for a worked example. We set c = 0.01 in our experiments.

5 EXPERIMENTS

5.1 **SETTING**

Backbone Models. Our experiments are conducted on a diverse range of open-source large language models to ensure broad applicability. This includes models from the Qwen2.5 (Yang et al., 2025b), Qwen3 series (Yang et al., 2025a), and the Llama3 series (Dubey et al., 2024).

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), and **Entropy** minimization (Prabhudesai et al., 2025). To further contextualize the comparison, we include a ground-truth supervised oracle baseline with GRPO (Shao et al., 2024). Additionally, we benchmark against judge-based methods, **LLM-as-a-Judge** (Pang et al., 2023; Zhao et al., 2025a) and **LLM-KD**, to provide a comprehensive evaluation. Additional details are provided in Appendix D.

Implementation Details. All methods are implemented using the VeRL framework (Sheng et al., 2025) and trained on $8 \times$ NVIDIA A100 GPUs. We train on 7,500 problems from the MATH dataset's training split (Hendrycks et al., 2021), and evaluate on the 5,000-problem validation split (referred to as MATH5000). For each reinforcement learning update step, we sample a batch of 128 problems and generate G=8 rollouts per problem. We use a learning rate of 3×10^{-6} and a KL penalty coefficient of $\beta=0.005$. To ensure fair and reproducible comparisons, we utilize the officially released chat-based prompting formats for all models. More details 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 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. The specific metrics, frameworks, and implementation details are provided in Appendix F.

5.2 Main Results

We evaluate **JURY-RL** on both in- and out-domain tasks, as shown in Tables 1³. A key finding is that JURY-RL not only outperforms all label-free baselines but also consistently surpasses the supervised GRPO with ground-truth rewards (GT), suggesting that proof-gated rewards can offer a superior learning signal to direct supervision.

Mathematical Reasoning and In-Domain Generalization. The results on mathematical reasoning benchmarks in Table 1 reveal a crucial insight into JURY-RL's learning mechanism. We observe that while its advantage over baselines is modest on the **in-distribution** MATH500 test set (which shares the same origin as our MATH training data), its superiority becomes substantially more pronounced on **out-of-distribution** math benchmarks such as GSM8K and AMC. We posit that this pattern arises because competing methods, including the supervised GT baseline, tend to overfit to the stylistic patterns and problem-solving shortcuts of the MATH dataset. In contrast, JURY-RL, by relying on formal verification, is incentivized to learn the underlying mathematical principles that are robust to such distributional shifts. This superior generalization within **the math domain** culminates in a significant overall performance leap over the GT baseline on all three backbones, with average score improvements of **+2.32** pts (+5.91% rel.) on Qwen3-1.7B, **+1.91** pts (+5.93% rel.) on Llama-3.2-3B-Instruct and **+1.53** pts (+3.28% rel.) on Qwen2.5-7B. For instance, on Qwen2.5-7B, JURY-RL's average of **48.13**% notably exceeds both the GT oracle (46.60%) and the strongest label-free competitor, LLM-KD (46.54%). This body of evidence indicates that the signal from formal verification is not merely a proxy for ground truth but can be a more potent objective for learning generalizable reasoning.

Out-of-Domain Generalization. This strong in-domain performance translates to robust out-of-domain generalization across code generation, instruction following, and multi-task knowledge tests. On Qwen2.5-7B, JURY-RL again surpasses the GT baseline, achieving an average of **40.45**% (+2.76 pts, +7.32% rel.). On the other two backbone models, its performance is statistically on par with GT while consistently ranking as the best-performing label-free method. These results demonstrate that optimizing for verifiable correctness encourages the model to learn fundamental, transferable skills that generalize beyond the mathematical domain used for training.

 $^{^{3}}$ Values are reported as mean \pm standard error of the mean (SEM).

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

Methods			Mathematics			Code		Instruction	Multi-Task	Average	
Withous	AIME24	AIME25	MATH500	GSM8K	AMC	LiveCode	CRUX	IFEval	MMLU-Pro	Average	
Qwen3-1.7B-Base											
Before RL	$6.67_{\pm 4.6}$	$6.67_{\pm 4.6}$	$45.8_{\pm 2.2}$	$63.31_{\pm 1.3}$	$26.51_{\pm 4.8}$	$4.59_{\pm0.4}$	$7.12_{\pm 0.9}$	$33.66_{\pm 1.0}$	$33.60_{\pm0.4}$	$25.33_{\pm 2.3}$	
GT-Reward	$10.00_{\pm 5.6}^{-}$	$0.00_{\pm 0.0}$	$68.2_{\pm 2.1}^{-}$	$83.09_{\pm 1.0}$	$34.94_{\pm 5.2}$	$14.84_{\pm 0.9}$	$33.75_{\pm1.7}$	$38.70_{\pm 1.0}$	$39.43_{\pm0.4}$	$35.88_{\pm 2.0}$	
Self-Certainty	$6.67_{\pm 4.6}$	$3.33_{\pm 3.3}$	$57.2_{\pm2.2}$	$74.68_{\pm1.2}$	$26.51_{\pm4.8}$	$10.37_{\pm0.7}$	$19.00_{\pm1.4}$	$38.58_{\pm 1.0}$	$35.28_{\pm0.4}$	$30.18_{\pm2.2}$	
Entropy	$10.00_{\pm 5.6}$	$3.33_{\pm 3.3}$	$65.2_{\pm 2.1}$	$79.98_{\pm 1.1}$	$32.53_{\pm 5.1}$	$12.64_{\pm0.8}$	$31.00_{\pm 1.6}$	$35.00_{\pm 1.0}$	$37.01_{\pm0.4}$	$34.08_{\pm 2.4}$	
Majority-Voting	$3.33_{\pm 3.3}$	$0.00_{\pm 0.0}$	$59.8_{\pm 2.2}$	$81.88_{\pm 1.1}$	$33.73_{\pm 5.2}$	$14.29_{\pm 0.9}$	$32.75_{\pm 1.7}$	$37.16_{\pm 1.0}$	$35.68_{\pm0.4}$	$33.18_{\pm 1.8}$	
CoReward	$10.00_{\pm 5.6}$	$3.33_{\pm 3.3}$	$66.0_{\pm 2.1}$	$82.03_{\pm 1.1}$	$33.73_{\pm 5.2}$	$14.27_{\pm 0.9}$	$32.12_{\pm 1.6}$	$37.28_{\pm 1.0}$	$37.39_{\pm0.4}$	$35.13_{\pm 2.4}$	
LLM-KD	$3.33_{\pm 3.3}$	$3.33_{\pm 3.3}$	$68.4_{\pm 2.1}$	$82.41_{\pm 1.1}$	$44.58_{\pm 5.5}$	$14.05_{\pm0.9}$	$34.62_{\pm 1.7}$	$34.65_{\pm 1.0}$	$37.97_{\pm0.4}$	$35.93_{\pm 2.1}$	
LLM-as-a-Judge	$10.00_{\pm 5.6}$	$3.33_{\pm 3.3}$	$62.8_{\pm2.2}$	$80.52_{\pm 1.1}$	$31.33_{\pm 5.1}$	$14.33_{\pm0.9}$	$36.12_{\pm 1.7}$	$36.28_{\pm 1.0}$	$33.83_{\pm0.4}$	$34.28_{\pm 2.4}$	
JURY-RL (Ours)	$13.33_{\pm 6.3}$	$6.67_{\pm 4.6}$	$68.4_{\pm 2.1}$	$83.32_{\pm 1.0}$	$36.14_{\pm 5.3}$	$14.54_{\pm 0.9}$	$35.00_{\pm 1.7}$	$36.55_{\pm 1.0}$	$38.11_{\pm 0.4}$	$36.90_{\pm 2.6}$	
	Llama-3.2-3B-Instruct										
Before RL	$3.33_{\pm 3.3}$	$0.00_{\pm 0.0}$	$42.8_{\pm 2.2}$	$69.75_{\pm 1.3}$	$13.25_{\pm 3.7}$	$3.00_{\pm 0.4}$	$25.50_{\pm 1.5}$	$54.41_{\pm 1.1}$	$32.01_{\pm0.4}$	$27.12_{\pm 1.6}$	
GT-Reward	$13.33_{\pm 6.3}$	$0.00_{\pm 0.0}$	$48.0_{\pm2.2}$	$76.72_{\pm 1.2}^{-}$	$22.89_{\pm 4.6}$	$7.05_{\pm 0.6}$	$32.62_{\pm1.7}$	$50.16_{\pm 1.1}^{-}$	$34.26_{\pm0.4}$	$31.67_{\pm 2.0}$	
Entropy	$6.67_{\pm4.6}$	$0.00_{\pm0.0}$	$40.4_{\pm 2.2}$	$68.99_{\pm 1.3}$	$13.25_{\pm 3.7}$	$5.38_{\pm 0.5}$	26.62 ± 1.6	$54.24_{\pm 1.1}$	$33.54_{\pm0.4}$	27.68 ± 1.7	
Self-Certainty	$3.33_{\pm 3.3}$	$0.00_{\pm 0.0}$	$40.2_{\pm2.2}$	$74.07_{\pm 1.2}$	$16.87_{\pm 4.1}$	$5.57_{\pm 0.6}$	$22.75_{\pm 1.5}$	$54.11_{\pm 1.1}$	$34.42_{\pm 0.4}$	$27.92_{\pm 1.6}$	
Majority-Voting	$10.00_{\pm 5.6}$	$0.00_{\pm 0.0}$	47.0 ± 2.2	79.08 ± 1.1	19.28 ± 4.3	$8.11_{\pm 0.7}$	31.87 ± 1.6	48.36 ± 1.1	33.94 ± 0.4	$30.85_{\pm 1.9}$	
CoReward	$10.00_{\pm 5.6}$	$0.00_{\pm 0.0}$	$49.2_{\pm 2.2}$	$79.76_{\pm 1.1}$	$20.48_{\pm 4.4}$	$5.57_{\pm 0.5}$	$30.38_{\pm1.6}$	$50.46_{\pm 1.1}$	$32.95_{\pm0.4}$	$30.98_{\pm 1.9}$	
LLM-KD	$10.00_{\pm 5.6}$	$0.00_{\pm 0.0}$	$49.4_{\pm 2.2}$	$78.47_{\pm 1.1}$	$20.48_{\pm 4.4}^{-}$	$7.55_{\pm 0.6}$	$32.38_{\pm 1.6}$	$49.06_{\pm 1.1}$	$34.04_{\pm0.4}$	$31.26_{\pm 1.9}$	
LLM-as-a-Judge	$6.67_{\pm 4.6}$	$0.00_{\pm0.0}$	$47.8_{\pm 2.2}$	$77.10_{\pm 1.2}$	$21.69_{\pm 4.5}$	$3.96_{\pm0.4}$	$33.88_{\pm 1.7}$	$51.46_{\pm 1.1}$	$34.22_{\pm 0.4}$	$30.75_{\pm 1.8}$	
JURY-RL (Ours)	$16.67_{\pm 6.9}$	$0.00_{\pm 0.0}$	$49.0_{\pm2.2}$	$80.74_{\pm 1.1}$	$24.10_{\pm 4.7}$	$6.16_{\pm 0.6}$	$32.62_{\pm 1.7}$	$50.09_{\pm 1.1}$	$34.54_{\pm0.4}$	$32.66_{\pm 2.1}$	
				Q	wen2.5-7B						
Before RL	$3.33_{\pm 3.3}$	$3.33_{\pm 3.3}$	$49.2_{\pm 2.2}$	$71.65_{\pm 1.2}$	$20.48_{\pm 4.4}$	$4.57_{\pm0.4}$	$28.00_{\pm 1.6}$	$40.61_{\pm 1.0}$	$44.03_{\pm0.4}$	$29.47_{\pm 2.0}$	
GT-Reward	$13.33_{\pm 6.3}$	$6.67_{\ \pm 4.6}$	76.4 ± 1.9	$89.61{\scriptstyle \pm 0.8}$	$46.99{\pm}5.5$	$12.78_{\pm 0.7}$	51.38 ± 1.8	$41.50_{\pm 1.0}$	45.09 ± 0.4	$42.64{\pm}2.6$	
Self-Certainty	$13.33_{\pm 6.3}$	$6.67_{\pm 4.6}$	$75.0_{\pm 1.9}$	$88.93_{\pm 0.9}$	$44.58_{\pm 5.5}$	$11.87_{\pm 0.7}$	$53.87_{\pm 1.8}$	$39.66_{\pm 1.0}$	$43.89_{\pm0.4}$	$41.98_{\pm 2.6}$	
Entropy	$16.67_{\pm 6.9}$	$10.00_{\pm 5.6}$	$73.4_{\pm 2.0}$	$85.52_{\pm 1.0}$	$39.76_{\pm 5.4}$	$15.68_{\pm0.8}$	$51.50_{\pm 1.8}$	$40.25_{\pm 1.0}$	$42.61_{\pm 0.4}$	$41.71_{\pm 2.8}$	
Majority-Voting	$10.00_{\pm 5.6}$	$3.33_{\pm 3.3}$	$71.4_{\pm 2.0}$	$90.22_{\pm 0.8}$	$38.55_{\pm 5.3}$	$18.37_{\pm 0.9}$	$52.38_{\pm 1.8}$	$42.72_{\pm 1.0}$	43.83 ± 0.4	$41.20_{\pm 2.4}$	
CoReward	$10.00_{\pm 5.6}$	$3.33_{\pm 3.3}$	$74.0_{\pm 2.0}$	$90.52_{\pm 0.8}$	$38.55_{\pm 5.3}$	$10.37_{\pm 0.6}$	$57.25_{\pm 1.8}$	$43.75_{\pm 1.0}$	$42.08_{\pm0.4}$	$41.09_{\pm 2.3}$	
LLM-KD	$13.33_{\pm 6.3}$	$10.00_{\pm 5.6}$	$73.6_{\pm 2.0}$	$89.99_{\pm0.8}$	$45.78_{\pm 5.5}$	$8.09_{\pm 0.5}$	$51.38_{\pm 1.8}$	$43.09_{\pm 1.0}$	43.73 ± 0.4	$42.11_{\pm 2.7}$	
LLM-as-a-Judge	$10.00_{\pm 5.6}$	$6.67_{\pm 4.6}$	$72.2_{\pm 2.0}$	$89.61_{\pm 0.8}$	$42.17_{\pm 5.4}$	$10.58_{\pm 0.6}$	$53.00_{\pm 1.8}$	$44.17_{\pm 1.0}$	$48.67_{\pm0.4}$	$41.90_{\pm 2.5}$	
JURY-RL (Ours)	$13.33_{\pm 6.3}$	$13.33_{\pm 6.3}$	$74.6_{\pm 1.9}$	$91.21_{\pm 0.8}$	$48.19_{\pm 5.5}$	$14.69_{\pm 0.8}$	$55.62_{\pm 1.8}$	$41.51_{\pm 1.0}$	$50.00_{\pm 0.4}$	$44.72_{\pm 2.8}$	

Overall Gains. JURY-RL demonstrates comprehensive performance improvements over the ground-truth (GT) baseline across both in-domain and out-of-domain tasks. Specifically, it outperforms the GT baseline by +1.02 pts (+2.84% rel.) on Qwen-1.7B-Base and +2.08 pts (+4.88% rel.) on Qwen-2.5-7B, while maintaining a robust gain of +0.99 pts (+3.13% rel.) on Llama-3.2-3B-Instruct. This indicates that the advantage of JURY-RL lies not only in addressing the training collapse problem common in label-free methods but also in delivering stable and scalable performance gains across different model scales. Crucially, the benefits of this verifiability-driven learning paradigm extend beyond single-answer accuracy (pass@1). As we will show in subsequent analyses, it also learns, at a more intrinsic level, to generate diverse and effective solution paths, leading to more substantial improvements in multi-attempt success rates (pass@k) and effectively averting the mode collapse common in self-supervised methods.

5.3 ANALYSIS

Ablation Studies of ResZero. We compare three fallback designs under the same proof-gated framework when δ =0: Zero Reward (no signal), MV Reward (rewarding the majority vote), and our proposed ResZero, with GT and Majority-Voting as references. As shown in Table 2, ResZero consistently achieves the highest *Average* across all backbones. On average, it outperforms Zero Reward by **+1.8** pts, MV Reward by **+6.9** pts, and even surpasses GT by **+1.4** pts. This stark performance gap reveals the critical nature of the fallback signal. Rewarding an unverified majority (MV Reward) is deceptive, as it causes the model to reinforce its own errors and collapse. Conversely, 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. ResZero provides a robust solution by navigating this trade-off between dangerous reinforcement and inefficient stagnation.

JURY-RL Achieves Stable Training and Avoids Collapse. To evaluate training stability, we tracked the validation accuracy of JURY-RL against key label-free baselines on MATH5000 validation set throughout the training process. As illustrated in Figure 2, the trajectory for the Entropy and Self-Certainty shows collapse after an initial performance

Table 2: Ablation results for the proposed ResZero reward ($\delta = 0$) on reasoning benchmark.

Methods			Mathematics			Code		Instruction	Multi-Task	Average	
Methods	AIME24	AIME25	MATH500	GSM8K	AMC	LiveCode	CRUX	IFEval	MMLU-Pro	Average	
Qwen3-1.7B-Base											
GT-Reward $\begin{vmatrix} 10.00_{\pm 5.6} & 0.00_{\pm 0.0} & 68.2_{\pm 2.1} & 83.09_{\pm 1.0} & 34.94_{\pm 5.2} & 14.84_{\pm 0.9} & 33.75_{\pm 1.7} & 38.70_{\pm 1.0} & 39.43_{\pm 0.4} & 35.88_{\pm 0.0} & 38.84_{\pm 0.0} & 39.43_{\pm 0.0} & 39$											
Majority-Voting	$3.33_{\pm 3.3}$	$0.00_{\pm 0.0}$	$59.8_{\pm 2.2}$	$81.88_{\pm 1.1}$	$33.73_{\pm 5.2}$	$14.29_{\pm 0.9}$	$32.75_{\pm 1.7}$	$37.16_{\pm 1.0}$	$35.68_{\pm0.4}$	$33.18_{\pm 1.8}$	
Proof-Gate + Zero Reward	$3.33_{\pm 3.3}$	$6.67_{\pm 4.6}$	$69.0_{\pm 2.1}$	$83.70_{\pm 1.0}$	$36.14_{\pm 5.3}$	$14.82_{\pm 0.9}$	$30.63_{\pm 1.6}$	$37.07_{\pm 1.0}$	$37.40_{\pm0.4}$	$35.42_{\pm 2.3}$	
Proof-Gate + MV Reward	$0.00_{\pm 0.0}$	$0.00_{\pm 0.0}$	$41.2_{\pm 2.2}$	$82.87_{\pm 1.0}$	$8.43_{\pm 3.0}$	14.48 ± 0.9	$34.00_{\pm 1.7}$	36.78 ± 1.0	35.33 ± 0.4	$28.12_{\pm 1.1}$	
JURY-RL (Proof-Gate + ResZero)	$13.33_{\pm 6.3}$	$6.67_{\pm4.6}$	$68.4_{\pm2.1}$	$83.32_{\pm 1.0}$	$36.14_{\pm 5.3}$	$14.54_{\pm0.9}$	$35.00_{\pm1.7}$	$36.55_{\pm 1.0}$	$38.11_{\pm 0.4}$	$36.90_{\pm 2.6}$	
Llama-3.2-3B-Instruct											
GT-Reward	$13.33_{\pm 6.3}$	$0.00_{\pm 0.0}$	$48.0_{\pm2.2}$	$76.72_{\pm 1.2}$	$22.89 {\scriptstyle \pm 4.6}$	$7.05_{\ \pm0.6}$	$32.62_{\pm 1.7}$	$50.16_{\pm 1.1}$	$34.26 {\scriptstyle \pm 0.4}$	$31.67_{\pm 2.0}$	
Majority-Voting	$10.00_{\pm 5.6}$	$0.00_{\pm 0.0}$	$47.0_{\pm 2.2}$	$79.08_{\pm 1.1}$	$19.28_{\pm 4.3}$	$8.11_{\pm 0.7}$	$31.87_{\pm 1.6}$	$48.36_{\pm 1.1}$	33.94 ± 0.4	$30.85_{\pm 1.9}$	
Proof-Gate + Zero Reward	$3.33_{\pm 3.3}$	$0.00_{\pm 0.0}$	$51.4_{\pm 2.2}$	$78.85_{\pm 1.1}$	$21.69_{\pm 4.5}$	$6.81_{\pm 0.6}$	$30.12_{\pm 1.6}$	$50.12_{\pm 1.1}$	$33.27_{\pm0.4}$	$30.62_{\pm 1.6}$	
Proof-Gate + MV Reward	$10.00_{\pm 5.6}$	$0.00_{\pm 0.0}$	$47.4_{\pm 2.2}$	$79.23_{\pm 1.1}$	$25.30_{\pm 4.8}$	$7.07_{\pm 0.6}$	$31.37_{\pm 1.6}$	$48.49_{\pm 1.1}$	$33.68_{\pm0.4}$	$31.39_{\pm 1.9}$	
JURY-RL (Proof-Gate + ResZero)	$16.67_{\pm 6.9}$	$0.00_{\pm 0.0}$	$49.0_{\pm 2.2}$	$80.74_{\pm 1.1}$	$24.10_{\pm 4.7}$	$6.16_{\pm 0.6}$	$32.62_{\pm 1.7}$	$50.09_{\pm 1.1}$	$34.54_{\pm0.4}$	$32.66_{\pm 2.1}$	
				Qwen2.5-7	В						
GT-Reward	$13.33_{\pm 6.3}$	$6.67{\scriptstyle \pm 4.6}$	76.4 ± 1.9	$89.61{\scriptstyle \pm 0.8}$	$46.99{\scriptstyle \pm 5.5}$	12.78 ± 0.7	$51.38{\scriptstyle\pm1.8}$	$41.50{\scriptstyle \pm 1.0}$	$45.09{\scriptstyle \pm 0.4}$	$42.64{\scriptstyle \pm 2.6}$	
Majority-Voting	10.00 ±5.6	$3.33_{\pm 3.3}$	$71.4_{\pm 2.0}$	$90.22_{\pm 0.8}$	$38.55_{\pm 5.3}$	$18.37_{\pm 0.9}$	$52.38_{\pm 1.8}$	$42.72_{\pm 1.0}$	$43.83_{\pm0.4}$	$41.20_{\pm 2.4}$	
Proof-Gate + Zero Reward	$13.33_{\pm 6.3}$	$6.67_{\pm 4.6}$	$75.8_{\pm 1.9}$	$90.67_{\pm 0.8}$	$39.76_{\pm 5.4}$	$18.14_{\pm 0.9}$	$52.25_{\pm 1.8}$	$41.67_{\pm 1.0}$	47.88 ± 0.4	$42.91_{\pm 2.6}$	
Proof-Gate + MV Reward	$0.00_{\pm 0.0}$	$0.00_{\pm 0.0}$	$51.6_{\pm 2.2}$	$89.69_{\pm0.8}$	$16.87_{\pm 4.1}$	$17.29_{\pm 0.9}$	$53.25_{\pm 1.8}$	$43.64_{\pm 1.0}$	$35.01_{\pm0.4}$	$34.15_{\pm 1.2}$	
JURY-RL (Proof-Gate + ResZero)	$13.33_{\pm 6.3}$	$13.33_{\pm 6.3}$	$74.6_{\pm 1.9}$	$91.21_{\pm 0.8}$	$48.19_{\pm 5.5}$	$14.69_{\pm0.8}$	$55.62_{\pm 1.8}$	$41.51_{\pm 1.0}$	$50.00_{\pm0.4}$	$44.72_{\pm 2.8}$	

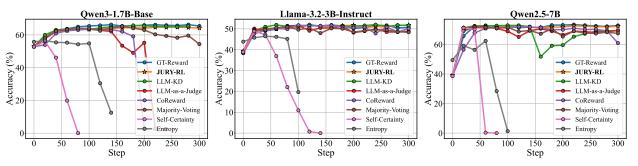


Figure 2: Accuracy on MATH5000 Validation set over training steps.

gain, as the model begins to reinforce spurious consensus. The LLM-as-a-Judge/LLM-KD/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.

Does JURL-RL Enhance Diversity? Yes. Table 3 shows that JURY-RL achieves substantially larger improvements in pass@k than in pass@1 (e.g., +4.05 pp vs. +2.32 pp for Qwen3-1.7B-Base), indicating an increase in solution diversity. This stems from our ResZero reward mechanism, which penalizes flawed consensus and redistributes rewards to explore alternative reasoning paths. This process actively counters mode collapse and incentivizes exploration, directly boosting multi-attempt success rates. The effect is visually confirmed in Figure 3, which tracks the average number of unique answers generated per problem during training. While baselines like Majority-Voting quickly suffer from mode collapse—converging to a single answer—JURY-RL sustains a high level of diversity, an indication of its robustness. Per-benchmark breakdowns are available in Appendix G.1.

Table 3: Performance of JURY-RL vs. GT-Reward on math reasoning tasks (pass@k and pass@1). k=16 for AIME, k=4 for MATH500 and GSM8K, and k=8 for AMC.

Model	Avera	age (pass@k)		Average (pass@1)			
1,10401	GT-Reward	JURY-RL	Δ (pp)	GT-Reward	JURY-RL	Δ (pp)	
Qwen3-1.7B-Base	55.36	59.41	+4.05	39.25	41.57	+2.32	
Llama-3.2-3B-Instruct	45.46	48.48	+3.02	32.19	34.10	+1.91	
Qwen2.5-7B	62.48	64.04	+1.56	46.60	48.13	+1.53	

433

434

435

436

437

438

439

440

441

442

443

444

445

446

447

448

449 450

451

452

453

455

456

457

458

459

460

461 462 463

464

465

466

467

468

469

470

472

473 474 475

476 477

478

479

480

481

482

483

484

485

Analysis of Verifier Signal Quality. Is Lean a better judge than LLM? While a formal verifier like Lean theoretically offers near-zero false positives, our practical pipeline involves upstream processes like auto-formalization and consistency checks, which can introduce errors. It is therefore crucial to compare the signal quality of our verifier against an LLM-as-a-Judge. As shown in Table 4, our Lean verifier provides a superior reward signal compared to the LLM-as-a-Judge. It achieves substantially higher **precision** (84.5% vs. 75.9%) at the cost of moderately lower recall (88.0% vs. 96.1%). This trade-off is paramount: high precision drastically reduces false positives, preventing reward hacking and tightly aligning the training objective with verifiable correctness. Conversely, the LLM-judge's noisy signal, stemming from low precision, risks reinforcing errors despite its higher recall. The Lean verifier's higher F1-score (86.2%) confirms its better overall balance, validating our "Proofs Dispose" principle of prioritizing signal fidelity for stable learning. While our verifier is not perfect, its imperfections stem from upstream components rather than the prover's core logic. We provide a deeper analysis of these nuances and their effect on training dynamics in Appendix G.2. **Impact of** c. We analyze the impact of the hyperparameter c in Eq. 5, convergence on the primary task objectives.

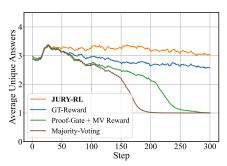


Figure 3: Average unique answers per sample over training steps on Qwen3-1.7B-Base.

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

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

which controls the penalty strength on the unverified majority proposal in our ResZero reward. As illustrated in Figure 4, 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 tradeoff: 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

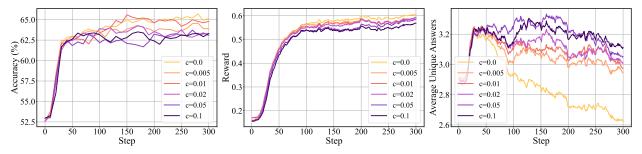


Figure 4: Training dynamics under different values of the hyperparameter c.

CONCLUSION

We introduced JURY-RL, a label-free RLVR framework that decouples proposal from disposal: majority voting across rollouts merely proposes a candidate answer, while a formal Lean verifier disposes the final reward. If the proposal is verified, only the supporting trajectories are rewarded; when verification is inconclusive, ResZero discards the unverifiable majority and assigns a zero-mean, variance-preserving residual reward that keeps optimization wellconditioned. This design jointly achieves three goals: scalability without human labels, truth alignment via verifiable correctness, and optimization stability in the absence of proof. Across mathematical reasoning, code generation, and multi-task evaluations, JURY-RL trains more stably than label-free baselines and achieves performance comparable to or better than supervised training with ground-truth rewards. Our work demonstrates that grounding RL in sparse but formally verified signals is a promising strategy for building robust and generalizable reasoning models without human labels.

ETHICS STATEMENT

This work studies label-free reinforcement learning with verifiable rewards (RLVR) for mathematical and programmatic reasoning. It does not involve human subjects, crowd workers, user studies, or the collection of personally identifiable information. All datasets used (e.g., MATH, GSM8K, AMC/AIME, LiveCodeBench, CRUX, MMLU-Pro, IFEval) are publicly available and used under their respective licenses; we redistribute nothing and provide only references and scripts to download from the original sources.

Potential risks include (i) *misuse*: stronger automated reasoning could be used to complete graded assignments or to generate convincing but incorrect solutions; (ii) *bias/coverage*: public benchmarks may contain stylistic biases or limited topical coverage; (iii) *safety*: LLM judges can transfer prompt or format biases into training signals. Our design explicitly mitigates these issues by proof-gating rewards with a formal verifier (Lean) to reduce false positives and by using the Residual-Zero (ResZero) fallback to avoid reinforcing unverifiable consensus. We report precisely which datasets, prompts, and evaluation harnesses are used, and we provide ablations and diagnostics to surface failure modes (e.g., collapse under majority-voting). We release only code, configuration files, and download scripts; no private or proprietary data are included.

REPRODUCIBILITY STATEMENT

We are committed to ensuring the full reproducibility of our research. All necessary components, including code, models, datasets, and experimental settings, are detailed below.

Theory. Full proofs for our theoretical analysis (including the zero-mean property of RESZERO) are provided in Appendix A.

Code. The implementation of JURY-RL and all baseline methods is based on the publicly available VeRL framework (Sheng et al., 2025). The complete source code for our experiments, including scripts for training and evaluation, will be made publicly available upon publication. All evaluation frameworks used are standard and open-source, with specific details and links provided in Appendix F.

Models. The backbone models used in our experiments are all publicly available open-source models from the Qwen and Llama3 series. Specifically, we used **Qwen3-1.7B-Base**, **Llama-3.2-3B-Instruct**, and **Qwen2.5-7B**. These models can be accessed through official repositories such as Hugging Face. The judge model used in our LLM-as-a-Judge baseline is qwen-2.5-72b-instruct, which is also publicly accessible. We will release our trained model upon acceptance.

Datasets. All datasets used for training and evaluation are standard, publicly available benchmarks. We train our models on a 7,500-problem subset of the official **MATH** training split (Hendrycks et al., 2021). We evaluate on the following benchmarks: **AIME24/25**, **MATH500**, **GSM8K**, **AMC**, **LiveCodeBench**, **CRUX**, **IFEval**, and **MMLU-Pro**. References and links for each are provided in Section 5.1.

Experimental Setup and Hyperparameters. All crucial hyperparameters for training, optimization, and generation are provided in Table 6 in Appendix E. The prompt formats used for all models are their officially released chat-based formats to ensure faithful reproduction. Details of the baseline implementations are described in Appendix D. The design of our Lean verifier is detailed in Appendix C.

REFERENCES

Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, and Hao Peng. The unreasonable effectiveness of entropy minimization in llm reasoning. *CoRR*, abs/2505.15134, 2025. doi: 10.48550/ARXIV.2505.15134. URL https://doi.org/10.48550/arXiv.2505.15134.

Together AI and Agentica. Deepcoder: A fully open-source 14b coder at o3-mini level. Blog post, 2025. URL https://www.together.ai/blog/deepcoder. 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, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah,

Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback. *CoRR*, abs/2204.05862, 2022. doi: 10.48550/ARXIV.2204.05862. URL https://doi.org/10.48550/arXiv.2204.05862.

- Jasmin C.Blanchette, Sascha Böhme, and Lawrence C.Paulson. Extending Sledgehammer with SMT Solvers. In *Proceedings of the 23rd International Conference on Automated Deduction (CADE23)*, volume 6803 of *Lecture Notes in Artificial Intelligence*, pp. 116–130. Springer, 2011. doi: 10.1007/978-3-642-22438-6_9.
- Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. Humans or LLMs as the judge? a study on judgement bias. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 8301–8327, Miami, Florida, USA, 2024. Association for Computational Linguistics. doi: 10. 18653/v1/2024.emnlp-main.474. URL https://aclanthology.org/2024.emnlp-main.474/.
- 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. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021. doi: 10.48550/ARXIV.2110.14168. URL https://doi.org/10.48550/arXiv.2110.14168.
- Leonardo de Moura and Sebastian Ullrich. The Lean4 Theorem Prover and Programming Language. In *Automated Deduction— CADE28*, volume 12699 of *Lecture Notes in Computer Science*, pp. 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37. URL https://lean-lang.org/papers/lean4.pdf.
- 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, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, and Wangding Zeng. Deepseek-v3 technical report. *CoRR*, abs/2412.19437, 2024. doi: 10.48550/ARXIV.2412.19437. URL https://doi.org/10.48550/arxiv.2412.19437.
- 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, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, and S. S. Li. Deepseek-r1: Incentivizing reasoning capability in Ilms via reinforcement learning. *CoRR*, abs/2501.12948, 2025. doi: 10.48550/ARXIV.2501.12948. URL https://doi.org/10.48550/arxiv.2501.12948.
- 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. The llama 3 herd of models. *CoRR*, abs/2407.21783, 2024. doi: 10.48550/ARXIV.2407.21783. URL https://doi.org/10.48550/arXiv.2407.21783.
- Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. *CoRR*, abs/2210.10760, 2022. doi: 10.48550/ARXIV.2210.10760. URL https://doi.org/10.48550/arXiv.2210.10760.
- Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I. Wang. Cruxeval: A benchmark for code reasoning, understanding and execution. *CoRR*, abs/2401.03065, 2024a. doi: 10.48550/ARXIV.2401.03065. URL https://doi.org/10.48550/arXiv.2401.03065.

Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 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, 2024b. URL https://openreview.net/forum?id=5h0qf7IBZZ.

- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In Joaquin Vanschoren and Sai-Kit Yeung (eds.), Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/be83ab3ecd0db773eb2dc1b0a17836a1-Abstract-round2.html.
- Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum. Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base model. *CoRR*, abs/2503.24290, 2025. doi: 10.48550/ARXIV.2503.24290. URL https://doi.org/10.48550/arXiv.2503.24290.
- Hui Huang, Xingyuan Bu, Hongli Zhou, Yingqi Qu, Jing Liu, Muyun Yang, Bing Xu, and Tiejun Zhao. An empirical study of llm-as-a-judge for LLM evaluation: Fine-tuned judge model is not a general substitute for GPT-4. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics*, *ACL 2025*, *Vienna*, *Austria*, *July 27 August 1*, 2025, pp. 5880–5895. Association for Computational Linguistics, 2025a. URL https://aclanthology.org/2025.findings-acl.306/.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 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, 2025b. doi: 10.1145/3703155. URL https://arxiv.org/abs/2311.05232.
- Hugging Face H4. AIME 2024 Benchmark. https://huggingface.co/datasets/HuggingFaceH4/ aime_2024, 2024.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *CoRR*, abs/2403.07974, 2024. doi: 10.48550/ARXIV.2403.07974. URL https://doi.org/10.48550/arXiv.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, Haiqing Guo, Han Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li, Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su, Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye, Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu, Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong, Weiran He, Weixiao Huang, Wenhao Wu, Wenyang He, Xianghui Wei, Xianqing Jia, Xingzhe Wu, Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du, Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, and Zonghan Yang. Kimi k1.5: Scaling reinforcement learning with llms. *CoRR*, abs/2501.12599, 2025. doi: 10.48550/ARXIV.2501.12599. URL https://doi.org/10.48550/arxiv.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 Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tülu 3: Pushing frontiers in open language model post-training. *CoRR*, abs/2411.15124, 2024. doi: 10.48550/ARXIV.2411.15124. URL https://doi.org/10.48550/arXiv.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. Rlaif vs. rlhf: Scaling reinforcement learning from human feedback with ai feedback. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 26874–26901. PMLR, 2024. URL https://arxiv.org/abs/2309.00267.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Letś verify step by step. *CoRR*, abs/2305.20050, 2024. doi: 10.48550/ARXIV. 2305.20050. URL https://doi.org/10.48550/arXiv.2305.20050.

651 652

653 654

655 656 657

658

659 660 661

667 668 669

666

670 671 672

673 674 675

680

685

686

687 689 690

692 693 694

691

- 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. Goedel-prover-v2: Scaling formal theorem proving with scaffolded data synthesis and self-correction, 2025. URL https://arxiv.org/abs/2508.03613.
- math-ai Team. Amc23: American mathematics competitions 2023 test set. https://huggingface.co/ datasets/math-ai/amc23, 2024. Hugging Face dataset; 40 AMC 2023 problems; accessed 2025-09-16.
- OpenAI. GPT-4 technical report. CoRR, abs/2303.08774, 2023. doi: 10.48550/ARXIV.2303.08774. URL https: //doi.org/10.48550/arXiv.2303.08774.
- OpenCompass. AIME 2025 Benchmark. https://huggingface.co/datasets/opencompass/ AIME2025, 2025.
- 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. Training language models to follow instructions with human feedback. CoRR, abs/2203.02155, 2022. doi: 10.48550/ARXIV.2203.02155. URL https://doi.org/10.48550/arXiv.2203.02155.
- Jing-Cheng Pang, Pengyuan Wang, Kaiyuan Li, Xiong-Hui Chen, Jiacheng Xu, Zongzhang Zhang, and Yang Yu. Language model self-improvement by reinforcement learning contemplation. CoRR, abs/2305.14483, 2023. doi: 10.48550/ARXIV.2305.14483. URL https://doi.org/10.48550/arXiv.2305.14483.
- Avinash Patil and Aryan Jadon. Advancing reasoning in large language models: Promising methods and approaches. arXiv preprint arXiv:2502.03671, 2025. URL https://arxiv.org/abs/2502.03671.
- Gérard P.Huet, Gilles Kahn, and Christine Paulin-Mohring. The Coq Proof Assistant: A Tutorial. Technical report, INRIA, 1997. URL https://coq.inria.fr.
- Mihir Prabhudesai, Lili Chen, Alex Ippoliti, Katerina Fragkiadaki, Hao Liu, and Deepak Pathak. Maximizing confidence alone improves reasoning. CoRR, abs/2505.22660, 2025. doi: 10.48550/ARXIV.2505.22660. URL https://doi.org/10.48550/arXiv.2505.22660.
- QwenTeam. Qwq-32b: Embracing the power of reinforcement learning, March 2025. URL https://qwenlm. github.io/blog/qwq-32b/.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), 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, 2023. URL http://papers.nips.cc/paper_files/paper/2023/ hash/a85b405ed65c6477a4fe8302b5e06ce7-Abstract-Conference.html.
- Z. Z. Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia 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. Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. CoRR, abs/2504.21801, 2025. doi: 10.48550/ARXIV.2504.21801. URL https://doi. org/10.48550/arXiv.2504.21801.
- Sheikh Shafayat, Fahim Tajwar, Ruslan Salakhutdinov, Jeff Schneider, and Andrea Zanette. Can large reasoning models self-train? CoRR, abs/2505.21444, 2025. doi: 10.48550/ARXIV.2505.21444. URL https://doi. org/10.48550/arXiv.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. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. CoRR, abs/2402.03300, 2024. doi: 10.48550/ARXIV.2402.03300. URL https://doi.org/10. 48550/arXiv.2402.03300.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 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, pp. 1279-1297. ACM, 2025. doi: 10.1145/3689031.3696075. URL https://doi.org/10.1145/3689031. 3696075.

- Jiawen Shi, Zenghui Yuan, Yinuo Liu, Yue Huang, Pan Zhou, Lichao Sun, and Neil Zhenqiang Gong. Optimization-based prompt injection attack to llm-as-a-judge, 2025. URL https://arxiv.org/abs/2403.17710.
- Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu. Crossing the reward bridge: Expanding RL with verifiable rewards across diverse domains. *CoRR*, abs/2503.23829, 2025. doi: 10. 48550/ARXIV.2503.23829. URL https://doi.org/10.48550/arXiv.2503.23829.
- Aman Singh Thakur, Kartik Choudhary, Venkat Srinik Ramayapally, Sankaran Vaidyanathan, and Dieuwke Hupkes. Judging the judges: Evaluating alignment and vulnerabilities in Ilms-as-judges, 2025. URL https://arxiv.org/abs/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, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023. doi: 10.48550/ARXIV.2307.09288. URL https://doi.org/10.48550/arxiv.2307.09288.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=pZ3i2yt5DoY.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *CoRR*, abs/2406.01574, 2024. doi: 10.48550/ARXIV.2406.01574. URL https://doi.org/10.48550/arXiv.2406.01574.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. *CoRR*, abs/2201.11903, 2022. doi: 10.48550/ARXIV.2201.11903. URL https://doi.org/10.48550/arXiv.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. Qwen3 technical report. *CoRR*, abs/2505.09388, 2025a. doi: 10.48550/ARXIV.2505.09388. URL https://doi.org/10.48550/arXiv.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, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zihan Qiu, and et al. Qwen2.5 technical report. CoRR, abs/2412.15115, 2025b. doi: 10.48550/ARXIV.2412.15115. URL https://doi.org/10.48550/arXiv.2412.15115.
- Qiying 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, Yuan Cheng, Qingchuan Pu, Feilin Yang, Yuning Jiang, Xinjie Hao, Sheng Wang, Weizhong Tian, Yu Xi, Bin Wang, Cong Li, Jiang Guo, Heyi Ouyang, Kaiwen Wang, Brandon Timmons, and et al. Dapo: An open-source llm reinforcement learning system at scale. *CoRR*, abs/2503.14476, 2025. doi: 10.48550/ARXIV.2503.14476. URL https://doi.org/10.48550/arxiv.2503.14476.
- Zizhuo Zhang, Jianing Zhu, Xinmu Ge, Zihua Zhao, Zhanke Zhou, Xuan Li, Xiao Feng, Jiangchao Yao, and Bo Han. Co-reward: Self-supervised reinforcement learning for large language model reasoning via contrastive agreement. *CoRR*, abs/2508.00410, 2025. doi: 10.48550/ARXIV.2508.00410. URL https://doi.org/10.48550/arXiv.2508.00410.

Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. Absolute zero: Reinforced self-play reasoning with zero data. *CoRR*, abs/2505.03335, 2025a. doi: 10.48550/ARXIV.2505.03335. URL https://doi.org/10.48550/arXiv.2505.03335.

Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. Learning to reason without external rewards. *CoRR*, abs/2505.19590, 2025b. doi: 10.48550/ARXIV.2505.19590. URL https://doi.org/10.48550/arXiv.2505.19590.

Yulai Zhao, Haolin Liu, Dian Yu, S. Y. Kung, Haitao Mi, and Dong Yu. One token to fool llm-as-a-judge, 2025c. URL https://arxiv.org/abs/2507.08794.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. *CoRR*, abs/2311.07911, 2023. doi: 10.48550/ARXIV. 2311.07911. URL https://doi.org/10.48550/arXiv.2311.07911.

APPENDIX

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. 5).
- §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 $y_i \sim \pi_\theta(\cdot \mid x)$ and parse answers $a_i = \text{ans}(y_i)$. Let $\hat{a} = \text{arg}\max_a \sum_{j=1}^G \mathbb{I}[a_j = a]$ 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 $\overline{r} = \hat{v}$, std $= \sqrt{\hat{v}(1-\hat{v})}$, hence

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

As a spurious consensus strengthens $(\hat{v} \to 1)$, supporters receive vanishing signal $(\hat{A}_i^+ \to 0)$ while dissenters incur diverging penalties $(\hat{A}_i^- \to -\infty)$, 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

Here, we provide a formal proof that the adaptive ResZero reward, as defined in Equation (5), maintains a strict zero-mean property. This property, $\sum_{i=1}^{G} r_i^{\text{ResZero}} = 0$, 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,

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:

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

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:

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

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.

$$-c\alpha|M| + G\gamma = 0$$
$$\gamma = \frac{c\alpha|M|}{G}$$

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.

 $[\]overline{|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.

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:

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

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 $\operatorname{std}(\{r_j\}) = \sqrt{\hat{v}(1-\hat{v})}$. 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 $(\hat{v} \to 1)$, the advantage for supporters vanishes $(\hat{A}_i^+ \to 0)$ while the penalty for dissenters diverges $(\hat{A}_i^- \to -\infty)$. 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: $std(\{r_j\}) = \sqrt{\frac{1}{G} \sum_{i=1}^{G} (0-0)^2} = 0.$

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, $\operatorname{std}(\{r_i\}) > 0$ as long as there is diversity among residual answers.

The advantage function thus becomes:

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

Theoretical Implication. This formulation provides a rich, structured learning signal that is both **corrective** and **exploratory**.

- 1. **Prevents Stagnation:** Since \hat{A}_i is non-zero, it provides a meaningful gradient, ensuring that learning continues.
- 2. Corrects Spurious Consensus: For rollouts in the majority group $(i \in M)$, r_i^{ResZero} is negative. This results in a negative advantage $(\hat{A}_i < 0)$, actively discouraging the policy from converging on the unverified consensus.
- 3. **Guides Exploration:** For rollouts in the residual group $(i \in R)$, the sign of r_i^{ResZero} (and thus \hat{A}_i) depends on the relative popularity of their answers within the minority. This rewards more promising alternative paths and penalizes less popular ones, guiding exploration in a principled manner.

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)

```
    Sample G rollouts y<sub>i</sub> ~ π<sub>old</sub> (· | x); parse a<sub>i</sub> = ans(y<sub>i</sub>).
    Compute vote shares v(a) and majority â ∈ arg max<sub>a</sub> v(a).
    Query verifier once: δ = verify(x, â).
    if δ=1 then
    Set r<sub>i</sub> = I[a<sub>i</sub> = â].
    else
    Form M, R; compute u<sup>(-i)</sup> and ū; set r<sub>i</sub> via Eq. (5).
    end if
    Compute group-normalized advantages Â<sub>i</sub> from {r<sub>i</sub>}<sup>G</sup><sub>i=1</sub>; broadcast token-wise.
    Update π<sub>θ</sub> 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. (5):

$$r_i^{\text{ResZero}} = \begin{cases} -c\alpha + \gamma = -0.1(0.5) + 0.025 = -0.025, & \text{if } i \in M \text{ (proposal \hat{a})} \\ \alpha(z_i - \bar{u}) + \gamma = 0.5(\frac{2}{3} - 0.5) + 0.025 = \frac{1}{12} + 0.025 \approx 0.0667, & \text{if } i \in R, a_i = b \\ \alpha(z_i - \bar{u}) + \gamma = 0.5(0 - 0.5) + 0.025 = -0.225, & \text{if } i \in R, a_i = c \end{cases}$$

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 & 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 5) 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 8 candidate formalizations. The Consistency-Checker then evaluates all candidates and selects the top 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:

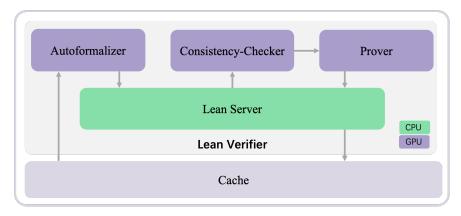


Figure 5: The pipeline of our Lean Verifier: a collaborative system that combines autoformalization, formal verification, and consistency checking to validate mathematical reasoning with high reliability.

• Temperature: 0.7

• Maximum Tokens: 32,768

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

C.3 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, including the defined batch size and the pass@8 verification setting.

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): The verification process introduces an additional overhead of approximately 200 seconds per step.
- **LLM-as-a-Judge:** The baseline using Qwen-2.5-72B for verification 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.

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.5. 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.

Beyond the runtime overhead, another critical aspect of the verifier is the trade-off between its performance and token 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.

Table 5 presents the performance and cost analysis of different prover configurations. While the Prover@64 configuration achieves the highest accuracy, it comes at a significant computational cost, averaging 67,429 tokens per verification. In contrast, the Prover@16 configuration, while exhibiting a slight and acceptable decrease in performance to 87.0% ACC, reduces the token cost by more than half to 33,346. Given this highly favorable trade-off between a marginal performance drop and substantial cost savings, we adopt Prover@16 as the default configuration for our main experiments to ensure efficiency.

Table 5: Performance and token cost analysis of Lean Verifier across different Prover configurations.

Models	MATH500			Token Costs				
111000015	ACC	TPR	F1-Score	Max Tokens per Verify	Avg Tokens per Verify	Avg Tokens per Response		
Prover@16	87.0	87.0	93.0	351,799	33,346	5,858		
Prover@32	89.0	89.0	94.0	658,061	52,064	6,607		
Prover@64	91.0	91.0	95.0	1,369,038	67,429	6,435		

C.4 PROMPT

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

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

```
1134
           Based solely on the above analysis, judge whether the Lean statement is a correct and exact formalization of the
1135
          mathematical problem.
1136
           5. Accuracy Confirmation
1137
          If correct: clearly confirm why all elements match.
1138
          If incorrect: list all mismatches and explain how each one affects correctness.
1139
          Note: While the analysis may be broad and open to interpreting all relevant features, the final judgment must be
1140
          based only on what is explicitly and formally expressed in the Lean statement. **Do not consider or assess any
1141
           part of the proof. Your judgment should be entirely about the accuracy of the statement formalization.**
           You should present the results following the format:
1142
          Input:
1143
          The Natural Language Statement:
1144
          A math problem and its answer (no proof).
1145
          The Formal Statement in Lean4:
1146
           ` ` 'lean
1147
          A Lean 4 theorem statement formalizing the problem. Proof is intentionally omitted (e.g., sorry).
1148
1149
          Output Format:
1150
          Return exactly one xml object
1151
           <comments>
1152
           Your brief analysis:
          Math Assertion Analysis: [...]
1153
          Lean Statement Analysis (Proof Ignored): [...]
1154
          Comparative Verification: [...]
1155
          Conclusion: [...]
1156
           Accuracy Confirmation: [... match confirmation or list of discrepancies...]
1157
           </comments>
1158
           <consistency>[Correct/Incorrect]</consistency>
1159
```

Consistency-Checker User Prompt

Input Data:

The Natural Language Statement:

{informal_prefix}

The Formal Statement in Lean4:

```lean

{formal\_statement}

1167

1160

1161 1162

1163

1164

1165

1166

1169

11701171

1172

1173 1174

1175

1176

1177

117811791180

1181 1182

1183

1184

1185

1186

1187

# **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 $ans(y_i)$ 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 ans(·)); 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 inversely related to the policy's output entropy for the final-answer tokens, encouraging more deterministic, high-confidence predictions.
- **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 <code>qwen-2.5-72b-instruct</code> 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 6.

Table 6: Reinforcement learning training hyperparameters. This configuration is consistently applied across all experiments to ensure a fair comparison.

| Hyperparameter | Value |
|------------------------------------------------|--------------------|
| Training Configuration | |
| Train Batch size (Number of Sampled Questions) | 128 |
| Rollouts per problem (G) | 8 |
| Max prompt length | 512 |
| Max new tokens | 3072 |
| Training epoch | 6 |
| Optimizer Parameters (AdamW) | |
| Learning rate | 3×10^{-6} |
| eta_1 | 0.9 |
| eta_2 | 0.999 |
| ϵ | 10^{-8} |
| Warmup style | Cosine |
| Warmup steps ratio | 0.1 |
| GRPO Algorithm Parameters | |
| KL loss coefficient (β) | 0.005 |
| clip ratio (ϵ) | 0.2 |
| Generation Parameters | |
| Training temperature | 1.0 |
| Evaluation temperature | 0.8 |
| top-p | 0.95 |

F BENCHMARK AND METRIC DETAILS

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

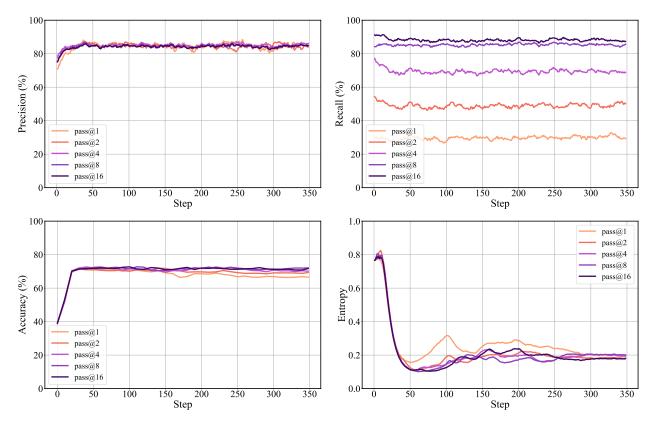


Figure 6: Training dynamics of precision, recall, validation accuracy, and training entropy under different Lean pass@k verification settings.

- AIME24/25, MATH500, GSM8K: We report pass@1 accuracy. All three are evaluated using the lighteval framework: https://github.com/huggingface/lighteval
- AMC: We report pass@1 accuracy. This is evaluated using the ttrl implementation: https://github.com/ruixin31/Spurious_Rewards/tree/main/code/ttrl
- LiveCodeBench: We report pass@1 accuracy. This is evaluated using the official library: https://github.com/LiveCodeBench/LiveCodeBench
- CRUX: We report pass@1 accuracy. This is evaluated using the ZeroEval framework: https://github.com/WildEval/ZeroEval
- MMLU-Pro: 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 PASS@K RESULTS FOR ALL COMPARISONS

This appendix complements The Diversity Analysis of JURY-RL in Section 5.3. The complete numerical results are listed in Table 7 for easy cross-reference. The pass@k results presented in Table 7 unequivocally demonstrate the effectiveness of our JURY-RL framework. Across all three 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.

Table 7: Pass@k Results (%) of RL performance comparison on math reasoning benchmarks.

| Methods | AIME24
pass@16 | AIME25
pass@16 | MATH500
pass@4 | GSM8K
pass@4 | AMC
pass@8 | Average | | | | |
|-----------------|-------------------|---------------------------------|-----------------------------|---------------------------------|----------------------------------|---------------------------------|--|--|--|--|
| Qwen3-1.7B-Base | | | | | | | | | | |
| Before RL | 16.67 ±6.9 | $13.33_{\pm 6.3}$ | $72.0_{\pm 2.0}$ | $89.31_{\pm 0.8}$ | $56.63_{\pm 5.4}$ | $49.59_{\pm 4.3}$ | | | | |
| GT-Reward | $26.67_{\pm 8.2}$ | $16.67{\scriptstyle \pm 6.9}$ | $82.0{\scriptstyle\pm1.7}$ | $92.42 {\scriptstyle \pm 0.7}$ | $59.04_{\pm5.4}$ | $55.36 {\scriptstyle\pm4.6}$ | | | | |
| Entropy | $20.00_{\pm7.4}$ | $23.33_{\pm 7.8}$ | $81.2_{\pm 1.8}$ | $90.75_{\pm0.8}$ | $59.04_{\pm 5.4}$ | $54.86_{\pm 4.7}$ | | | | |
| Self-Certainty | $23.33_{\pm7.8}$ | $26.67 {\scriptstyle \pm 8.2}$ | 77.2 ± 1.9 | $90.83_{\pm0.8}$ | $57.83_{\pm 5.4}$ | $55.17_{\pm4.8}$ | | | | |
| Majority-Voting | $23.33_{\pm 7.8}$ | $16.67_{\pm 6.9}$ | $77.2_{\pm 1.9}$ | $92.34_{\pm 0.7}$ | $53.01_{\pm 5.5}$ | $52.51_{\pm 4.6}$ | | | | |
| CoReward | $26.67_{\pm 8.2}$ | $16.67_{\pm 6.9}$ | $77.4_{\pm 1.9}$ | $92.04_{\pm0.8}$ | $54.22_{\pm 5.5}$ | $53.40_{\pm 4.6}$ | | | | |
| LLM-KD | 20.00 ± 7.4 | $16.67{\scriptstyle \pm 6.9}$ | 82.2 ± 1.7 | 93.25 ± 0.7 | $55.42_{\pm 5.5}$ | 53.51 ± 4.4 | | | | |
| LLM-as-a-Judge | 16.67 ± 6.9 | $13.33 \pm \scriptstyle 6.3$ | $74.6{\scriptstyle~\pm1.9}$ | $91.74_{\pm 0.8}$ | 55.42 ± 5.5 | 50.35 ± 4.3 | | | | |
| JURY-RL (Ours) | $30.00_{\pm 8.5}$ | $30.00_{\pm 8.5}$ | $82.0{\scriptstyle~\pm1.7}$ | $92.42_{\pm0.7}$ | $62.65{\scriptstyle\pm5.3}$ | $59.41_{\pm 5.0}$ | | | | |
| | | Llama-3 | 3.2-3B-Instruc | t | | | | | | |
| Before RL | 23.33 ±7.8 | $6.67_{\pm 4.6}$ | $65.0_{\pm2.1}$ | $86.88_{\pm 0.9}$ | $49.40_{\pm 5.5}$ | $46.26_{\pm 4.2}$ | | | | |
| GT-Reward | 23.33 ± 7.8 | 6.67 ± 4.6 | 63.8 ± 2.1 | 90.14 ± 0.8 | $43.37 \pm _{5.4}$ | $45.46{\scriptstyle \pm 4.2}$ | | | | |
| Entropy | $16.67_{\pm 6.9}$ | $3.33_{\pm 3.3}$ | $58.8_{\pm2.2}$ | $84.46_{\pm 1.0}$ | $50.60_{\pm 5.5}$ | $42.77_{\pm 3.8}$ | | | | |
| Self-Certainty | $13.33_{\pm 6.3}$ | $10.00{\scriptstyle\pm5.6}$ | $58.0_{\pm2.2}$ | $84.69{\scriptstyle \pm 1.0}$ | 44.58 ± 5.5 | 42.12 ± 4.1 | | | | |
| Majority-Voting | $16.67_{\pm 6.9}$ | 0.00 ± 0.0 | 63.6 ± 2.1 | $90.98_{\pm0.8}$ | $39.76_{\pm 5.4}$ | $42.20{\scriptstyle\pm3.0}$ | | | | |
| CoReward | 23.33 ± 7.8 | 0.00 ± 0.0 | 60.4 ± 2.2 | $90.45_{\pm0.8}$ | $43.37 \pm _{5.4}$ | 43.51 ± 3.3 | | | | |
| LLM-KD | $23.33_{\pm7.8}$ | 3.33 ± 3.3 | $65.0_{\pm2.1}$ | $89.99_{\pm0.8}$ | $49.40_{\pm 5.5}$ | $46.21 {\scriptstyle \pm 3.9}$ | | | | |
| LLM-as-a-Judge | $26.67_{\pm 8.2}$ | $10.00_{\pm 5.6}$ | $62.2_{\pm2.2}$ | $89.84_{\pm0.8}$ | $40.96{\scriptstyle \pm 5.4}$ | $45.93_{\pm4.4}$ | | | | |
| JURY-RL (Ours) | $30.00_{\pm 8.5}$ | 3.33 ± 3.3 | 64.8 ± 2.1 | $90.07_{\pm0.8}$ | $54.22_{\pm 5.5}$ | $48.48_{\pm4.0}$ | | | | |
| | | Q _H | ven2.5-7B | | | | | | | |
| Before RL | $30.00_{\pm 8.5}$ | 20.00 ± 7.4 | 76.2 ± 1.9 | $93.78_{\pm0.7}$ | 63.86 ± 5.3 | $56.77_{\pm4.8}$ | | | | |
| GT-Reward | $33.33_{\pm 8.8}$ | 30.00 ± 8.5 | 85.2 ± 1.6 | $95.22 {\scriptstyle \pm 0.6}$ | $68.67_{\pm5.1}$ | $62.48{}_{\pm 4.9}$ | | | | |
| Entropy | $30.00_{\pm 8.5}$ | 36.67 ± 8.9 | 85.2 ± 1.6 | $93.78_{\pm0.7}$ | $67.47_{\pm 5.1}$ | 62.62 ± 5.0 | | | | |
| Self-Certainty | 23.33 ± 7.8 | 30.00 ± 8.5 | 84.6 ± 1.6 | 93.86 ± 0.7 | 69.88 ± 5.0 | 60.33 ± 4.7 | | | | |
| Majority-Voting | $33.33_{\pm 8.8}$ | $30.00{\scriptstyle \pm 8.5}$ | $85.4_{\pm1.6}$ | $95.00_{\pm0.6}$ | $63.86 {\scriptstyle\pm}_{5.3}$ | $61.52_{\pm 4.9}$ | | | | |
| CoReward | $30.00_{\pm 8.5}$ | $26.67 {\scriptstyle \pm 8.2}$ | $83.8_{\pm1.6}$ | $95.60_{\pm0.6}$ | $63.86 {\scriptstyle\pm}_{5.3}$ | $59.99_{\pm 4.8}$ | | | | |
| LLM-KD | $33.33_{\pm 8.8}$ | $33.33_{\pm 8.8}$ | $87.8_{\pm 1.5}$ | $95.75_{\pm0.6}$ | $67.47_{\pm 5.1}$ | $63.54_{\pm 4.9}$ | | | | |
| LLM-as-a-Judge | $20.00_{\pm7.4}$ | $13.33_{\pm 6.3}$ | $82.2_{\pm1.7}$ | $95.38_{\pm0.6}$ | $59.04_{\pm5.4}$ | $53.99_{\pm 4.3}$ | | | | |
| JURY-RL (Ours) | $36.67_{\pm 8.9}$ | $30.00_{\pm 8.5}$ | $86.6_{\pm1.5}$ | $95.83_{\pm 0.6}$ | $71.08_{\pm 5.0}$ | $64.04_{\pm 4.9}$ | | | | |

G.2 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 6, 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 probabilistic 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.

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 such that x + 4y < 72. Find the x + 4y < 72.
```

Answer

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

```
To evaluate whether the given answer of 1152 is correct for the problem, we need
   \hookrightarrow to find the maximum value of the function \( f(x, y) = xy (72 - 3x - 4y) \)
   \hookrightarrow subject to the constraint \( 3x + 4y < 72 \).
First, let's rewrite the function:
[f(x, y) = xy (72 - 3x - 4y).]
We can use the method of Lagrange multipliers or analyze the function directly.
   \hookrightarrow However, a simpler approach is to use symmetry and calculus.
Consider the function (g(t) = t(72 - 3t - 4t) = t(72 - 7t) = 72t - 7t^2). This
   \hookrightarrow is a quadratic function in terms of \( t \), where \( t = x + y \). The
   \hookrightarrow maximum value of this quadratic function occurs at the vertex, which is
   \hookrightarrow given by:
Now, we need to find (x \setminus) and (y \setminus) such that (x + y = \frac{36}{7}) and
   \hookrightarrow \( 3x + 4y \) is minimized while still satisfying \( 3x + 4y < 72 \).
Given the symmetry and the nature of the problem, we can assume \setminus ( x = y \setminus). Then:
[2x = \frac{36}{7} \le x = \frac{18}{7}.]
\hookrightarrow 7\left(\frac{18}{7}\right) = 18. \]
\lceil f \cdot (\frac{18}{7}, \frac{18}{7} \cdot ) =
   → \left(\frac{18}{7}\right)\left(\frac{18}{7}\right) (72 -
   \hookrightarrow 3\left(\frac{18}{7}\right) - 4\left(\frac{18}{7}\right)) \]
[ = \left( \frac{18}{7}\right)^2 (72 - 7\left( \frac{18}{7}\right) ) ]
```

1415 1416

1417

1418

1419 1420

1421 1422

1423

1424

1425

1426

1427

1428

1429 1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450 1451

1452

1453

1454

1455

1456

1457

```
1404
        [ = \left( \frac{18}{7} \right)^2 (72 - 18) ]
1405
        \  \  = \left(\frac{18}{7}\right)^2 \cdot 54 
1406
        [ = \frac{324}{49} \cdot 54 ]
1407
        [ = \frac{17496}{49} ]
        [ = 357.0612244898. ]
1408
1409
        This value does not match 1152. Therefore, the given answer of 1152 is incorrect.
1410
1411
        Final judgment: False.
1412
```

```
Lean Prover Output (Prove Success)
import Mathlib
import Aesop
set_option maxHeartbeats 0
open BigOperators Real Nat Topology Rat
theorem test_problem :
  sSup {z : \mathbb{R} | \exists x y : \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land z = x * y * (72 - 3*x
    \hookrightarrow - 4*y)} = 1152 := by
  have h_1: (1152 : \mathbb{R}) \in \{z : \mathbb{R} \mid \exists x y : \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land z =
    \rightarrow 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 h_2 : \forall (z : \mathbb{R}), z \in {z : \mathbb{R} | \exists x y : \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land z
    \hookrightarrow = x * y * (72 - 3*x - 4*y)} → z ≤ (1152 : \mathbb{R}) := by
    intro z hz
    rcases hz with \langle x, y, hx, hy, hxy, rfl \rangle
    have h_3: 0 < x := hx
    have h_4 : 0 < y := hy
    have h_5: 3 * x + 4 * y < 72 := hxy
     have h_6: 0 < 72 - 3 * x - 4 * y := by linarith
     have h_7: 0 < 3 * x := by positivity
     have h_8: 0 < 4 * y := by positivity
     have h_9: 0 < (3 * x) * (4 * y) := by positivity
     have h_{10}: 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) \le 1152
    have h_{11}: (3 * x + 4 * y + (72 - 3 * x - 4 * y)) ^ 3 <math>\geq 27 * (3 * x) * (4 * y)
    \hookrightarrow * (72 - 3 * x - 4 * y) := by
      nlinarith [sq_nonneg (3 * x - 4 * y), sq_nonneg (3 * x - (72 - 3 * x - 4 *
    \hookrightarrow 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 h_{12}: (72 : \mathbb{R}) ^ 3 \geq 27 * (3 * x) * (4 * y) * (72 - 3 * x - 4 * y) := by
       have h_{13}: (3 * x + 4 * y + (72 - 3 * x - 4 * y)) = 72 := by ring
       \texttt{rw} \text{ [$h_{13}$] at $h_{11}$}
       linarith
    have h_{13} : x * y * (72 - 3 * x - 4 * y) \leq 1152 := by have h_{14} : (72 : \mathbb{R}) ^ 3 = 373248 := by norm_num
       have h_{15}: 27 * (3 * x) * (4 * y) * (72 - 3 * x - 4 * y) = 324 * (x * y * (72
    \hookrightarrow - 3 * x - 4 * y)) := by
         ring
       rw [h_{14}] at h_{12}
       rw [h<sub>15</sub>] at h<sub>12</sub>
       nlinarith
     linarith
```

```
1458
              have h_3 : sSup {z : \mathbb{R} | \exists x y : \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land z = x * y *
1459
                \hookrightarrow (72 - 3*x - 4*y)} = 1152 := by
1460
                have h_4: IsLUB {z: \mathbb{R} | \exists x y: \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land z = x * y
1461
                \leftrightarrow * (72 - 3*x - 4*y)} 1152 := by
                   refine' \langle \text{fun z hz} \Rightarrow \text{h}_2 \text{ z hz}, \text{ fun y hy} \Rightarrow \_ \rangle
1462
                   -- We need to show that 1152 is the least upper bound.
1463
                   -- Since 1152 is in the set, any upper bound y must be at least 1152. have h_5: (1152: \mathbb{R}) \in {z : \mathbb{R} | \exists x y : \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land
1464
1465
                \hookrightarrow z = x * y * (72 - 3*x - 4*y)} := h<sub>1</sub>
1466
                   have h_6: (1152: \mathbb{R}) \leq y:= by
                      apply hy
1467
                      exact h_5
1468
                   linarith
1469
                 -- The supremum of the set is 1152 because it is the least upper bound.
1470
                have h_5: sSup {z : \mathbb{R} | \exists x y : \mathbb{R}, x > 0 \land y > 0 \land 3*x + 4*y < 72 \land z = x * y
1471
                \leftrightarrow * (72 - 3*x - 4*y)} = 1152 := by
                   apply le_antisymm
1472
                   \cdot -- Show that the supremum is less than or equal to 1152.
1473
                      apply csSup_le
1474
                      · -- The set is nonempty.
1475
                        exact \langle 1152, h_1 \rangle
1476
                      \cdot -- Every element in the set is less than or equal to 1152.
1477
                         intro z hz
                         exact h<sub>2</sub> z hz
1478
                   \cdot -- Show that 1152 is less than or equal to the supremum.
1479
                      apply le_csSup
1480
                      \cdot -- The set is bounded above.
1481
                        use 1152
1482
                         intro z hz
                        exact h2 z hz
1483
                      \cdot -- 1152 is in the set.
1484
                         exact h<sub>1</sub>
1485
                 exact h<sub>5</sub>
1486
1487
              exact h<sub>3</sub>
1488
```