# VARIATION IN VERIFICATION: UNDERSTANDING VERI-FICATION DYNAMICS IN LARGE LANGUAGE MODELS

**Anonymous authors**Paper under double-blind review

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

033

034

035

037

040

041

042

043

044

046

047

051

052

#### **ABSTRACT**

Recent advances have shown that scaling test-time computation enables large language models (LLMs) to solve increasingly complex problems across diverse domains. One effective paradigm for test-time scaling (TTS) involves LLM generators producing multiple solution candidates, with LLM verifiers assessing the correctness of these candidates without reference answers. In this paper, we study generative verifiers, which perform verification by generating chain-of-thought (CoT) reasoning followed by a binary verdict. We systematically analyze verification dynamics across three dimensions – problem difficulty, generator capability, and verifier generation capability – through empirical studies on 12 benchmarks across mathematical reasoning, knowledge, and natural language reasoning tasks using 14 open-source models (2B to 72B parameter range) and GPT-4o. Our experiments reveal three key findings about verification effectiveness: (1) Easy problems allow verifiers to more reliably certify correct responses; (2) Weak generators produce errors that are easier to detect than strong generators; (3) Verification ability is generally correlated with the verifier's own problem-solving capability, but this relationship varies with problem difficulty. These findings reveal opportunities for optimizing basic verification strategies in TTS applications. First, given the same verifier, some weak generators can nearly match stronger ones in post-verification TTS performance (e.g., the Gemma2-9B to Gemma2-27B performance gap shrinks by 75.7%). Second, we identify cases where strong verifiers offer limited advantages over weak ones, as both fail to provide meaningful verification gains, suggesting that verifier scaling alone cannot overcome fundamental verification challenges.

# 1 Introduction

Large language models (LLMs) have advanced rapidly in solving reasoning tasks such as mathematics and code generation, yet their outputs remain unreliable, often containing subtle or obvious mistakes (Ke et al., 2025; Lightman et al., 2023). LLM based verification (Angelopoulos et al., 2025; Cemri et al., 2025; Huang et al., 2023b; Mao et al., 2024) has emerged as a central mechanism to identify such errors in a scalable manner. Recent work has increasingly focused on *generative verifiers* (Liu et al., 2025d; Mahan et al., 2024; Zhang et al., 2025), which frame verification as next-token prediction: the model typically generates a chain-of-thought (CoT) reasoning trace and then outputs a binary verdict token. This approach has been shown to outperform earlier discriminative verifiers or scalar reward models (RMs, Lightman et al., 2023), as it better leverages the inherent text-generation capabilities of LLMs. One valuable downstream application of automatic verification is test-time scaling (TTS), where additional inference-time compute is allocated to improve generation performance. A popular paradigm of TTS is the use of a verifier model to evaluate candidate responses, filter errors, and identify correct solutions. This approach underlies techniques such as rejection sampling (Brown et al., 2024), re-ranking (Zhou et al., 2025), weighted majority voting (Wang et al., 2024a; 2022), and step-level generation (Snell et al., 2024).

Current practice in LLM verification often deploys strong, typically closed-source frontier models as verifiers. This practice rests on the assumption that verification quality scales with a verifier's capability to solve the same problem (i.e., its generation capability), a correlation demonstrated in recent work (Chen et al., 2025c; Krumdick et al., 2025; Tan et al., 2024). However, this practice may be suboptimal given that verifying a solution is often easier than generating one from scratch, a phenomenon referred to as "verification asymmetry" (Wei, 2025). This asymmetry appears in several fields. In convex optimization, dual certificates enable efficient validation of optimality of a proposed solution,

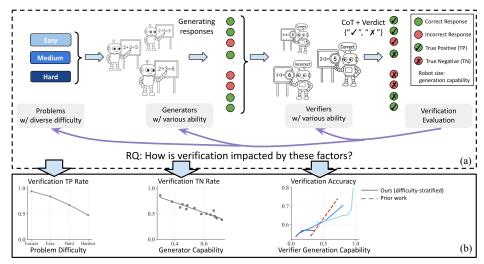


Figure 1: **Overview of our study on verification dynamics.** (a) We consider generative verification: an LLM generator produces a solution to a problem, and an LLM verifier conditions on the problem and solution to generate a verification CoT followed by a binary verdict ("Correct"/"Incorrect"). We design controlled experiments that vary problem difficulty, generator generation capability, and verifier generation capability, investigating how each of these factors influences verification performance. (b) Our analysis reveals three patterns: problem difficulty governs recognition of correct responses (true positives); generator generation capability determines error detectability (true negatives); and verifier generation capability correlates with performance in a difficulty-dependent manner, revealing non-linear regimes left uncovered in prior work. The three plots were generated by aggregating benchmark data across three domains and averaging performance metrics over 15 models.

while in factorization, verifying correctness is trivial compared to finding the prime factors. Thus, it is worth investigating verification as a distinct capability rather than merely a byproduct of generation.

Despite extensive research on generation dynamics and the factors influencing generation quality (Allen-Zhu & Li, 2025; Chen et al., 2024; Ye et al., 2025), the dynamics of verification remain largely unexplored. In particular, little is known about how problem characteristics, properties of generated responses, and model capabilities interact to determine verification effectiveness. Without understanding verification dynamics, one can risk misallocating computational resources by defaulting to expensive frontier models when simpler alternatives might suffice. This gap in understanding motivates our central research question: what factors influence verification success?

In this paper, we present a systematic study of generative verification across three dimensions–problem difficulty, generator capability, and verifier generation capability–shown in Figure 1. We quantify verification performance by measuring the probability of the verifier recognizing both correct and incorrect generated solutions in controlled experimental settings. We focus on verifiable problems with objective ground-truth answers in mathematical reasoning, knowledge question-answering (QA), and natural language (NL) reasoning domains. This allows us to objectively measure verifier and generator performance, while simulating the reference-free evaluation settings where verifiers are typically deployed in practice, e.g., in TTS. While our experiments use these domains as a testbed, we believe the insights should extend to any domain where correctness can be reliably defined and checked.

**Main Findings.** While prior work showed that verifier generation capability correlates with verification performance (Chen et al., 2025; Krumdick et al., 2025; Tan et al., 2024), we reveal that two additional factors, *problem difficulty* and *generator capability*, also critically influence verification success, as illustrated in Figure 1. Our analysis reveals:

- Problem difficulty primarily governs the recognition of correct solutions: verifiers are more likely to recognize correct solutions on easy problems than on difficult ones.
- Generator capability influences error detection: errors made by weak generators are easier to detect than those made by strong generators.
- Verifier generation capability correlates with verification performance in a manner dependent on problem difficulty: saturated (or uncorrelated) for easy problems, linear for medium problems, and threshold-limited for hard problems.

Our empirical analysis includes 2,347 math problems from eight datasets, 1,196 knowledge QA problems, and 901 NL reasoning problems, evaluated across 14 open-source models and GPT-40.

**Application to TTS.** We demonstrate the practical implications of our findings for TTS. First, given the same verifier, the TTS performance of a weak generator can nearly match the performance of a strong generator. For instance, Gemma2-9B achieves comparable performance to Gemma2-27B when both use the same verifier, GPT-4o. Second, we identify regimes where a strong verifier (e.g., GPT-4o) offers no additional benefit and can be replaced by a weaker verifier (e.g., Qwen2.5-7B), with both providing limited gains. This occurs with strong generators or with problems at either extreme of the difficulty spectrum.

# 2 RELATED WORK

Automatic Evaluation. The deployment of LLMs as evaluators has emerged as a central mechanism for scalable assessment, with efforts focusing on training specialized small evaluators through fine-tuning (Wang et al., 2024b; Whitehouse et al., 2025; Xu et al., 2025; Zhang et al., 2025). Beyond reference-based verifiers (Chen et al., 2025a; Liu et al., 2025b), verification approaches include self-verification (Chen et al., 2023; Huang et al., 2023a;b; Kumar et al., 2024; Shinn et al., 2023), where models reflect on or critique their own outputs, and multi-agent verification (Li et al., 2023; Lifshitz et al., 2025; Zhuge et al., 2024), where multiple agents collaborate in debate-style or hierarchical setups to improve reliability. Prior work identifies several factors influencing evaluation performance. One important factor is evaluator generation capability. Krumdick et al. (2025) find that evaluator performance changes significantly based on whether the evaluator is capable of answering the question or not. Tan et al. (2024) demonstrate the correlation between pairwise judging ability and generation ability on the same set of problems. Chen et al. (2025b) observe linear relationships between evaluation improvements and reasoning-required sample proportions in fine-tuned evaluators. Chen et al. (2025c) show a strong positive correlation between generation capability and evaluation accuracy. Our work extends these findings by identifying unexplored factors that influence evaluation performance. We also demonstrate that the relationship between the evaluator generation capability and evaluation quality is more nuanced than previously understood.

Verification for Test-Time Scaling. Early studies explore how to effectively apply verification methods to improve TTS performance. Snell et al. (2024) show RMs improve various TTS approaches, including Best-of-N and beam search, while Liu et al. (2025a) find that compute-optimal strategies vary with policy models and problem difficulty. Recent work explores alternatives to discriminative RMs: Zhang et al. (2025) show trained generative verifiers outperform RMs in Best-of-N, and Zhou et al. (2025)'s JETTS benchmark demonstrates generative evaluators match outcome RMs in reranking. While verification benefits from increased model size and test-time compute, recent work addresses how to reduce these computational costs. Saad-Falcon et al. (2025) propose a framework to aggregate weak verifiers to approach strong ones; Angelopoulos et al. (2025) balance weak/strong evaluators for efficiency; Stroebl et al. (2024) analyze fundamental limits of resampling with imperfect verifiers; and Singhi et al. (2025) propose strategies to balance solving-verification trade-offs. Our work studies the factors driving verification and explores their implications for TTS.

# 3 EXPERIMENTAL SETUP

## 3.1 Preliminaries

**Problem and Response Space.** Let x denote a problem with ground-truth answer  $y^*(x)$ . A model response r to x consists of a CoT solution and a final answer a(r), and we consider the response correct if  $a(r) = y^*(x)$ . As discussed in Section 1, our study uses verifiable problems with objective answers, allowing us to rigorously evaluate verifier outputs against ground-truth while simulating reference-free evaluation settings.

**Generator and Verifier.** A generator G maps a problem x to a distribution over responses, denoted  $r \sim G(\cdot|x)$ . A verifier V takes a problem–response pair (x,r) and outputs a judgment of correctness. In the binary case,  $V(x,r) \in \{0,1\}$ , where 1 indicates acceptance and 0 indicates rejection. More generally, a generative verifier produces a verification CoT explaining its reasoning, followed by an explicit verdict such as "Correct" or "Incorrect." The prompt templates are provided in Appendix A.

**Generation Capability.** We measure the generation capability of a model using its *pass rate*. For a generator G and problem x, we define  $p_G(x) = \Pr[a(r) = y^*(x) \mid r \sim G(\cdot | x)]$  as the pass rate on a single problem, i.e., the probability that G solves x correctly on one sampled attempt. We define  $p_G(\mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} p_G(x)$  as the pass rate aggregated over a dataset  $\mathcal{D}$ , which we use as the overall

measure of a model's generation capability. Since the verifiers we study are generic LLMs (e.g., GPT-40), we measure a verifier's generation capability using the same metric by prompting it as a generator.

We estimate these pass rates empirically as  $\hat{p}_G(x;K)$  and  $\hat{p}_G(\mathcal{D};K)$  by sampling K responses per model-problem pair. Since K is fixed at 64 throughout, we often omit K for simplicity, writing  $\hat{p}_G(x)$  and  $\hat{p}_G(\mathcal{D})$ . We use temperature 0.7 and top-p 1.0 as default sampling hyperparameters, and adopt recommended settings when available (e.g., temperature 0.7 and top-p 0.8 for the non-thinking mode of Qwen3). Ground-truth correctness is established with Math-Verify (Kydlíček, 2025), supplemented by LLM-as-a-judge grading to reduce false negatives (details in Appendix B.1). These 64 responses per problem-model pair are used to estimate generation capability and problem difficulty.

**Problem Difficulty.** We define the difficulty of a problem as the average pass rate across a set of diverse generators  $\mathcal{G}$ ,  $d(x) = \frac{1}{|\mathcal{G}|} \sum_{G \in \mathcal{G}} \hat{p}_G(x)$ . This score reflects how broadly solvable a problem is: if most generators succeed, d(x) is high (easy problem), while if few succeed, d(x) is low (hard problem). It provides a model-agnostic way to partition problems by difficulty, extending prior work (Snell et al., 2024), which measured difficulty relative to a single generator.

**Verification Metrics and Evaluation.** We evaluate verifiers using true positive rate (TPR), the probability of the verifier accepting a correct response:  $\text{TPR} = \mathbb{E}[V(x,r) \mid a(r) = y^*(x)]$ , and true negative rate (TNR), the probability of rejecting an incorrect response:  $\text{TNR} = \mathbb{E}[1 - V(x,r) \mid a(r) \neq y^*(x)]$ . We also report balanced accuracy,  $\text{Acc}_{\text{bal}} = \frac{1}{2}(\text{TPR} + \text{TNR})$ , which accounts for class imbalance. For verification evaluation, we subsample 8 responses from each 64-sample pool, balanced with 4 correct and 4 incorrect when possible. For very hard problems with fewer than 4 correct responses, we keep all correct ones and sample incorrect ones to reach 8 total (and vice versa for easy problems). Each verifier evaluates responses from all 15 models over the full test set using greedy decoding, unless a controlled subset is specified.

**Verification-Augmented Test-time Scaling.** We consider the TTS setting of sampling multiple responses from the generator and filtering with a verifier before evaluation. For each problem  $x \in \mathcal{D}$ , we sample K responses from the generator using a fixed temperature, with K=64 in our experiments. Without verification, TTS performance is measured as  $\hat{p}_G(\mathcal{D};K)$  (or  $\hat{p}_G(\mathcal{D})$ ), the empirical pass rate defined above. With verification, the verifier V evaluates each candidate, and only responses deemed "Correct" are retained for evaluation. The performance of verification-augmented TTS is measured as

$$\hat{p}_{G,V}(\mathcal{D};K) = \frac{1}{|\mathcal{D}|} \sum_{x \in \mathcal{D}} \left( \frac{1}{K'} \sum_{i=1}^{K} \mathbb{1} \left( a(r_i) = y^*(x) \right) \cdot V(x, r_i) \right)$$
(1)

where  $K' = \sum_{i=1}^K V(x,r_i)$ . This metric represents the conditional pass rate, i.e., the fraction of correct responses among those retained by the verifier. A corner case arises when the verifier rejects all responses (K' = 0); in this case, we set the metric to the generator's pass rate  $\hat{p}_G(\mathcal{D})$ , so evaluation reverts to selecting from the original K responses in the non-verified setting. We define the *verification gain* from verifier V as the difference relative to the performance without verification,  $\Delta \hat{p}_V = \hat{p}_{G,V}(\mathcal{D}) - \hat{p}_G(\mathcal{D})$ , which quantifies how much gain can be attributed to verification. Note that our formulation of TTS differs from the common setting where a single "best" response (e.g., by majority vote) is selected and then evaluated. Instead, we report the empirical pass rate of the verifier-retained pool, which can be interpreted as the expected accuracy of uniformly sampling one response from that pool. This expectation-based view captures the average quality of verifier-retained responses without tying performance to a specific selection strategy.

## 3.2 TASKS AND MODELS

**Mathematical Reasoning.** We collect a total of 2,347 problems from the test sets of eight mathematical reasoning benchmarks: GSM8K (Cobbe et al., 2021), MATH500 (Hendrycks et al., 2021), OlympiadBench (He et al., 2024), AIME24/25 (Li et al., 2024), AMC23 (Li et al., 2024), Minerva-Math (Lewkowycz et al., 2022), and BBEH Multi-step Arithmetic (Kazemi et al., 2025). We use the entire test sets of these benchmarks, except for GSM8K, from which we subsample 600 of 1,319 problems to balance difficulty distribution and reduce the proportion of easy problems.

<sup>&</sup>lt;sup>1</sup>If the verifier generates an invalid output (e.g., due to the CoT running out of max generation length), we treat it as an uninformative verdict of "Correct" and "Incorrect" each with probability of 50%. Computationally, we set V(x,r)=0.5 in this case, and also in Equation 1.

**Knowledge.** We use a subset of MMLU-Pro (Wang et al., 2024c) as our knowledge category. We randomly subsample 10% from each of its 14 disciplines, yielding 1,196 problems. MMLU-Pro consists of college-level multiple-choice questions spanning STEM, humanities, and social sciences.

**Natural Language Reasoning.** We collect 901 multiple-choice problems from three benchmarks. (1) ReClor (validation set, Yu et al., 2020), a multiple-choice benchmark requiring logical analysis of short passages. (2) FOLIO (Han et al., 2022), a first-order logic reasoning benchmark in natural language. (3) GPQA Diamond (Rein et al., 2024), a dataset that consists of graduate-level multiple-choice science questions, requiring multi-step reasoning.

**Models.** We use 14 open-source models from four families: (1) Qwen2.5 at 3B, 7B, and 72B (Team, 2024); Qwen3 at 4B, 8B, and 32B (Yang et al., 2025); (2) Llama-3.2 at 3B, Llama-3.1 at 8B, and Llama-3.3 at 70B (Grattafiori et al., 2024); (3) Gemma-2 at 2B, 9B, and 27B (Team et al., 2024); (4) Ministral 8B and Mistral-Small-24B; and one closed-source model GPT-40 (Hurst et al., 2024). All models are instruction-tuned versions by default. Each model is used both as a generator and a verifier. We use abbreviated model names in figures for space efficiency; see Appendix B.2 for mappings.

#### 4 EXPERIMENTAL RESULTS

Our experiments focus on how problem difficulty and generator and verifier generation capability influence verification performance. We present the three research questions and main findings below.

- **RQ1:** How does problem difficulty affect verification? (Section 4.1) TPR increases steadily with decreasing problem difficulty, meaning verifiers better recognize correct responses on easier problems. However, TNR shows no predictable relationship with problem difficulty. This indicates that problem difficulty primarily influences correctness recognition.
- RQ2: How does the generator's generation capability influence verification? (Section 4.2) As generators become stronger, TNR decreases substantially while TPR increases only slightly. This reveals that generator capability primarily determines error detectability: stronger generators produce errors that are harder for verifiers to identify.
- RQ3: How does verifier generation capability impact verification? (Section 4.3) Verifier generation capability and verification performance are generally positively correlated. However, the form of correlation depends heavily on problem difficulty: linear correlation occurs in medium-difficulty problems, while nonlinear patterns appear in other difficulty levels.

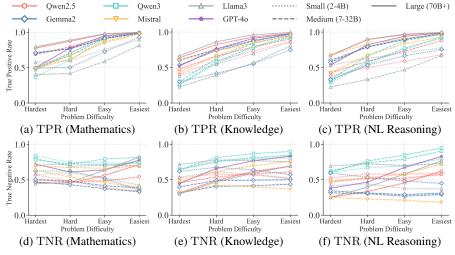


Figure 2: **Problem difficulty primarily affects** TPR **of verification.** Each curve shows verifier performance across four difficulty groups, with the x-axis indicating problem difficulty and the y-axis reporting TPR (a-c) and TNR (d-f). Colors denote model families, and line styles indicate model size.

# 4.1 How Does Problem Difficulty Affect Verification?

To examine how problem difficulty influences verification, we partition problems into four equal-sized quartiles by their difficulty score d(x), termed "hardest", "hard", "easy", and "easiest".

**Problem difficulty primarily influences the verifier's ability to recognize correct responses.** Our analysis is conducted at two levels of granularity: response level and problem level. Both analyses

reveal that problem difficulty mainly shapes the verifier's sensitivity to correct responses, while not consistently affecting its ability to identify incorrect responses.

At the response level, we compute the TPR and TNR of all responses within each difficulty quartile. As shown in Figure 2, TPR increases steadily as problems become easier, while TNR shows no clear trend. This pattern is consistent across model families and domains. At the problem level, we pool responses from all generators for each problem and compute a single TPR and TNR per problem. The distribution of these metrics within each quartile is reported in Figures 7 and 8 of Appendix C.1. We observe that easier problems yield higher and more stable TPR, while harder problems exhibit lower and more variable TPR. In contrast, TNR distributions show no consistent correlation with problem difficulty. To understand this pattern, case studies in Figure 20 show that verifiers tend to generate their own reference solutions for comparison during verification. As the problem difficulty increases, these verifier-generated answers become increasingly incorrect, producing false negatives that reduce TPR. In Appendix C.2, we show that the main verification dynamics about TPR we identified generalize to reasoning models, while extended reasoning provides benefits and alters TNR behavior.

# 4.2 How Does Generator Capability Influence Verification?

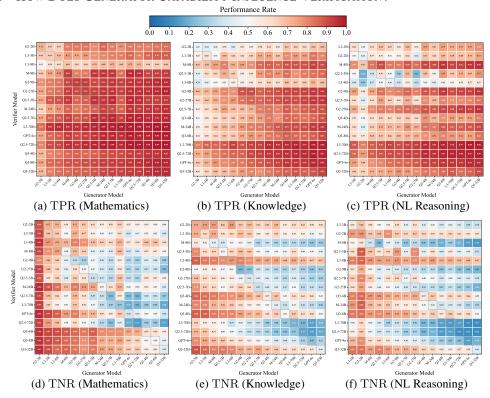


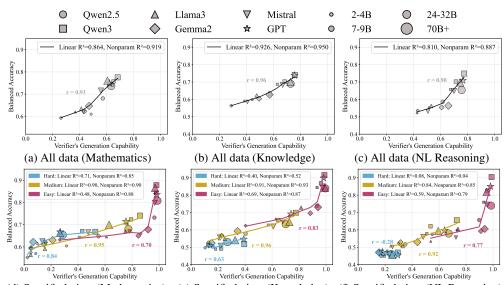
Figure 3: Generator capability influences verifier performance of identifying incorrect responses. Heatmaps show (a-c) TPR and (d-f) TNR when pairing 15 verifier models (rows) with 15 generator models (columns). Rows and columns are ordered by models' generation capability computed on all problems of each domain. Values indicate mean performance over the evaluation subset.

We study how generator capability affects verifier performance by having each verifier evaluate responses from each generator. Generators of different capabilities may produce extreme response distributions, e.g., weak generators may produce no correct response on hard problems within 64 samples. To ensure fair comparison, we compute TPR on problem subsets where all generators produce at least one correct response. Analogously, TNR is computed on problems where all generators produce at least one incorrect response. Details are provided in Appendix B.3.

As shown in Figures 3a to 3c, TPR remains uniformly high across nearly all settings and increases further with stronger generators. The heatmap is dominated by red colors, with values mostly above 0.7, indicating that most verifiers are already reliable at recognizing correct responses. As generator capability improves, TPR approaches 1.0. This suggests that generator strength influences recognition of correct responses in a relatively mild way.

Generator capability correlates with error detection in verification. In Figures 3d to 3f, moving from weaker generators on the left to stronger ones on the right, the heatmap shifts generally from red to blue, indicating a substantial decrease in TNR. For example, in the Mathematics domain, for the Qwen2.5-72B verifier, TNR drops from 0.68 on solutions generated by Llama-3.1-8B to 0.17 on those by Qwen3-32B. The overall pattern is consistent across three domains and nearly all verifiers. These results show that generator capability strongly modulates the detection of incorrect responses. Figure 21 presents case studies exploring this phenomenon. Strong generators produce internally consistent reasoning chains where early mistakes (e.g., missed cases) propagate coherently, yielding well-structured but incorrect solutions that cause verifier false positives. Weak generators produce self-contradictory solutions with apparent inconsistencies, facilitating verifier rejection. In Appendix C.2, we show that this finding generalizes to reasoning models.

# 4.3 How does Verifier Generation Capability Impact Verification?



(d) Stratified view (Mathematics) (e) Stratified view (Knowledge) (f) Stratified view (NL Reasoning)

Figure 4: **Distinct correlation forms between verification performance and generation capability.** Solid lines represent nonparametric fits to the data; r indicates the Pearson correlation coefficient. (a-c) Averaged across all problems, verifier generation capability exhibits a strong linear correlation with balanced accuracy. (d-f) When stratified by problem difficulty, distinct correlation patterns emerge.

We measure verifier generation capability and evaluate verification performance using balanced accuracy ( $Acc_{bal}$ ) on the entire test set. Each verifier is evaluated on responses from all generators, and we report results both averaged across all problems and stratified by problem difficulty. To characterize the relationship between generation capability and verification performance, we employ locally weighted regression (Cleveland, 1979) with a bandwidth of 0.6 to fit nonparametric curves. We compare  $R^2$  values between nonparametric and linear fits to assess linearity. We also report the Pearson correlation coefficient (Benesty et al., 2009) as another measure of linear correlation.

Figures 4a to 4c show a strong overall correlation between verifier generation capability and verification accuracy, with NL reasoning showing less linearity than other domains. This result is consistent with prior work showing that evaluator accuracy tends to track the evaluator's task performance, with the relationship appearing nearly linear. While this global trend validates findings in prior work (Chen et al., 2025c; Tan et al., 2024), a closer inspection of the trend reveals highly non-linear regimes.

Verifier generation capability influences verification accuracy differently based on problem difficulty. Stratified analysis reveals regime-dependent correlation with phase-transition behavior. We partition problems into 10 equal-width bins by difficulty d(x) and analyze three representative intervals: hard [0.1, 0.3), medium [0.4, 0.5), and easy [0.8, 0.9) in Figures 4d to 4f. For hard problems (blue), verification accuracy shows minimal improvement with increasing capability. Mathematics plateaus around 0.65 accuracy after initial gains, while other domains remain flat throughout. Notably, verifiers achieve below-random accuracy on hard NL Reasoning problems, which we analyze in Appendix C.3. Medium problems (yellow) exhibit steady accuracy increases with capability, indicating strong linear

relationships. This is confirmed by linear and nonparametric fits that yield nearly identical  $R^2$  values, with r>0.9. Easy problems exhibit a threshold effect at the x-axis around 0.9: below this threshold, the relationship is linear; above it, small capability improvements yield large verification gains. Hard and easy regimes show nonlinearity with nonparametric  $R^2$  exceeding linear  $R^2$  by 0.1–0.2 and r<0.85. The exception is NL Reasoning on hard problems, where both fits yield near-zero  $R^2$ , indicating no meaningful capability-accuracy relationship. Appendix C.4 provides additional results, including analysis across all difficulty intervals (Figure 12) and discusses implications of these findings.

# 5 APPLICATION TO TEST-TIME SCALING (TTS)

Our analysis in Section 4 is conducted with verification itself as the end goal. However, our findings have direct implications for TTS. We analyze two research questions in TTS settings that naturally arise out of our previous findings, and present our results below:

- RQ4: Given a fixed verifier, can a weak generator match a stronger generator in TTS? (Section 5.1) Weak generators can nearly match stronger generators' post-verification performance. Verification gains peak at weak-medium generators by achieving a high error detection rate (TNR) while maintaining a moderately high correctness recognition rate (TPR).
- RQ5: Can weak verifiers match the gains of strong verifiers in TTS? (Section 5.2) The verification gain gap between weak and strong verifiers narrows at both low and high problem difficulty extremes, and when using strong generators.

The following sections present results on the Mathematics domain, with complete results across all three domains in Appendices C.5 and C.6.

### 5.1 CAN WEAK GENERATORS MATCH STRONGER GENERATORS IN TTS?

We evaluate TTS with a fixed verifier (GPT-40) by varying generator capability and reporting pass rates before and after verification, along with the verification gain  $\Delta \hat{p}_V$ .

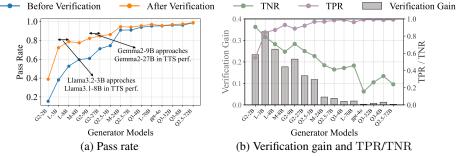


Figure 5: TTS performance before and after verification when varying generator strength. Results are reported on problems with difficulty in the range [0.7, 0.8) from the Mathematics domain, including 181 problems. (a-c) Pass rate before (blue) and after (orange) adding a fixed verifier (GPT-40). The generators in the x-axis are ordered from weaker (left) to stronger (right) by generation capability measured on the problem subset. (d-f) Bar chart shows the verification gain  $\Delta \hat{p}_V$  (left y-axis) for each generator. Lines show the verifier's TNR and TPR (right y-axis).

Verification gain peaks for weak-medium generators, enabling them to approach stronger **models post-verification.** As shown in Figure 5a, weak generators start with much lower pass rates but improve dramatically after verification, reaching levels comparable to larger models. For example, Gemma2-9B starts from a significantly lower baseline but, after verification, achieves a pass rate nearly matching Gemma2-27B. The performance gap shrinks from 10.3% to 2.5%, closing 75.7% of the original difference. Figure 5b explains this phenomenon: as generator strength increases (left to right), TNR decreases sharply while TPR rises only modestly, consistent with **RQ2** findings. Consequently, verification gain (gray bars) peaks at weak-medium generators. These generators achieve high TNR for effective error filtering while maintaining moderate TPR to preserve correct responses. For the strongest generators, errors become harder to identify, causing TNR decline and limiting gains. In Appendix C.5, we show the findings derived from the Mathematics domain generalize well to two other domains in Figure 13. We also provide additional evidence confirming the generalizability of these findings. First, verification gains peak for weak-medium generators across a broad range of problem difficulties ( $d(x) \ge 0.3$ ) in all domains (Figures 14 to 16). Second, performance gaps of most weak and strong model pairs can be reduced by verification when evaluated on the entire domain datasets (Figure 17), mostly achieving 30-50% reduction.

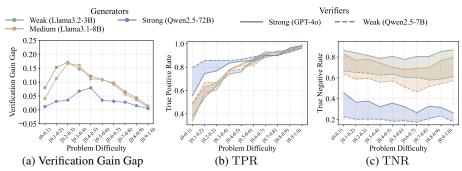


Figure 6: Weak vs. strong verifiers under varying problem difficulty and generator strength. The x-axis shows problem difficulty, ordered from hardest to easiest, measured relative to each generator. (a) Verification gain gap between strong and weak verifiers when applied to weak, medium, and strong generators. (b) TPR increases as problems become easier; shadow band indicates TPR gap between two verifiers. (c) As generators strengthen, TNR decreases overall and the TNR gap (shadow band) narrows. Results are from the Mathematics domain.

# 5.2 CAN WEAK VERIFIERS MATCH THE GAINS OF STRONG VERIFIERS IN TTS?

We analyze the verification gain gap between a strong verifier (GPT-4o) and a weaker one (Qwen2.5-7B) across problem difficulty ranges and generator strengths. The verification gain  $\Delta \hat{p}_V$  is defined in Section 3 and the gap between verifier is  $\Delta \hat{p}_{V_{\text{strong}}} - \Delta \hat{p}_{V_{\text{weak}}}$ . Our goal is to identify when this gap narrows, as such regimes suggest weak verifiers can substitute for strong ones.

The gap narrows on the extremes of problem difficulty As shown in Figure 6a, the verification gain gap shrinks as problems become easier, which corresponds to the rising TPR for both weak and strong verifiers seen in Figure 6b. This aligns with our **RQ1** findings that easier problems improve TPR for all verifiers. Even weak verifiers reliably recognize correct responses on easy problems, leaving little room for strong ones to provide additional benefit. At the opposite extreme, the gap also narrows on the hardest problems. As discussed in **RQ3** and shown in Figures 19d to 19f, increasing verifier generation capability (or scaling up to larger models) fails to improve verification accuracy on hard problems, resulting in only marginal performance differences between weak and strong verifiers.

The gap narrows as generators become stronger. Figure 6 shows that increasing generator capability reduces the difference between weak and strong verifiers. This is consistent with **RQ2** (Section 4.2), where we observed that the verifier's TNR decreases as the generator capability increases. As both weak and strong verifiers experience lower TNR, the gap between them also shrinks, shown as the narrowing shaded band between solid and dashed curves in Figure 6c).

The results here are obtained from the Mathematics domain, and we show the findings generalize to two other domains in Appendix C.6. In regimes of very easy/hard problems or when evaluating strong generators' responses, weak verifiers provide gains to TTS performance comparable to strong verifiers. However, these convergence regimes coincide with minimal verification benefit overall. Figures 19a to 19c shows verification gains drop to 0.1 or below for both verifiers on easy and hard problems, verification on strong generators yields peak gains of only 0.1, precisely where the gap narrows. Thus, while weak and strong verifiers converge in these regimes, this convergence occurs where both provide minimal practical value. This reveals that scaling verifiers from 7B models to GPT-4o fails to overcome fundamental verification challenges, with GPT-4o providing limited improvement over small open-source models in the identified regimes.

#### 6 Conclusion

We study LLM verification across problem difficulty, generator capability, and verifier generation capability, revealing that verification success depends on their interactions. We find that problem difficulty primarily shapes correct solution recognition, generator capability influences error detectability, and verifier generation capability correlates with verification in problem difficulty-dependent patterns. We examine the implications of these findings for verification deployment in TTS, identifying both opportunities and limitations. Stronger generators may not be necessary, as weaker generators can approach the post-verification performance of stronger ones when paired with a fixed verifier. This suggests potential for strategic model pairing that could reduce computational costs in verifier-based TTS methods. Our results also identify regimes where investing in larger verifiers yields no benefit, such as when evaluating responses from strong generators or problems at difficulty extremes.

# ETHICS STATEMENT

This work exclusively evaluates large language models on publicly available academic benchmarks containing mathematical reasoning, knowledge, and natural language problems with objective ground-truth answers. All experiments involve automated evaluation of model outputs without human subject participation. The datasets used are established research benchmarks designed for educational problem-solving tasks. Our study aims to understand verification dynamics to improve the computational efficiency of LLM systems, posing no ethical concerns regarding privacy, harmful content generation, or potential misuse.

# REPRODUCIBILITY STATEMENT

We provide comprehensive details to ensure reproducibility of our findings. All experiments use publicly available datasets and open-source/commercial LLMs. We specify the model names, versions, dataset sources, and inference hyperparameters in Section 3.2. Complete prompt templates for both generation and verification tasks are provided in Appendix A. The mathematical formulations of all metrics, along with estimation procedures and aggregation methods, are formally defined and clearly described in Section 3.1 and Appendix B.

# REFERENCES

- Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 3.2, knowledge manipulation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Anastasios N Angelopoulos, Jacob Eisenstein, Jonathan Berant, Alekh Agarwal, and Adam Fisch. Cost-optimal active ai model evaluation. *arXiv preprint arXiv:2506.07949*, 2025.
- Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. Pearson correlation coefficient. In *Noise reduction in speech processing*, pp. 1–4. Springer, Berlin, Heidelberg, 2009.
- Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv* preprint arXiv:2407.21787, 2024.
- Mert Cemri, Melissa Z Pan, Shuyi Yang, Lakshya A Agrawal, Bhavya Chopra, Rishabh Tiwari, Kurt Keutzer, Aditya Parameswaran, Dan Klein, Kannan Ramchandran, et al. Why do multi-agent llm systems fail? *arXiv preprint arXiv:2503.13657*, 2025.
- Ding Chen, Qingchen Yu, Pengyuan Wang, Wentao Zhang, Bo Tang, Feiyu Xiong, Xinchi Li, Minchuan Yang, and Zhiyu Li. xverify: Efficient answer verifier for reasoning model evaluations. *arXiv* preprint arXiv:2504.10481, 2025a.
- Nuo Chen, Zhiyuan Hu, Qingyun Zou, Jiaying Wu, Qian Wang, Bryan Hooi, and Bingsheng He. Judgelrm: Large reasoning models as a judge. *arXiv preprint arXiv:2504.00050*, 2025b.
- Qiguang Chen, Libo Qin, Jiaqi WANG, Jingxuan Zhou, and Wanxiang Che. Unlocking the capabilities of thought: A reasoning boundary framework to quantify and optimize chain-of-thought. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Wei-Lin Chen, Zhepei Wei, Xinyu Zhu, Shi Feng, and Yu Meng. Do llm evaluators prefer themselves for a reason? *arXiv preprint arXiv:2504.03846*, 2025c.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023.
- Xiusi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang,
   Denghui Zhang, Tong Zhang, et al. Rm-r1: Reward modeling as reasoning. arXiv preprint
   arXiv:2505.02387, 2025d.
  - W.S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(368):829–836, 1979. doi: 10.1080/01621459.1979.10481038.

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. arXiv preprint arXiv:2110.14168,
 2021.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. The language model evaluation harness, 07 2024. URL https://zenodo.org/records/12608602.

- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. arXiv preprint arXiv:2407.21783, 2024.
- Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, et al. Folio: Natural language reasoning with first-order logic. *arXiv preprint arXiv:2209.00840*, 2022.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*, 2024.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv* preprint arXiv:2103.03874, 2021.
- Jiaxin Huang, Shixiang Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 1051–1068, Singapore, December 2023a. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.67. URL https://aclanthology.org/2023.emnlp-main.67/.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. *arXiv preprint arXiv:2310.01798*, 2023b.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. arXiv preprint arXiv:2410.21276, 2024.
- Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, Sanket Vaibhav Mehta, Lalit K Jain, Virginia Aglietti, Disha Jindal, Peter Chen, et al. Big-bench extra hard. *arXiv preprint arXiv:2502.19187*, 2025.
- Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li, Chengwei Qin, Peifeng Wang, Silvio Savarese, et al. A survey of frontiers in llm reasoning: Inference scaling, learning to reason, and agentic systems. *arXiv preprint arXiv:2504.09037*, 2025.
- Michael Krumdick, Charles Lovering, Varshini Reddy, Seth Ebner, and Chris Tanner. No free labels: Limitations of llm-as-a-judge without human grounding. *arXiv preprint arXiv:2503.05061*, 2025.
- Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.
- H. Kydlíček. Math-verify: Math verification library, 2025. URL https://github.com/huggingface/math-verify.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857, 2022.

- Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q. Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. Hugging Face repository, 2024. Available at https://huggingface.co/datasets/AI-MO/NuminaMath-CoT.
  - Ruosen Li, Teerth Patel, and Xinya Du. Prd: Peer rank and discussion improve large language model based evaluations. *arXiv preprint arXiv:2307.02762*, 2023.
  - Shalev Lifshitz, Sheila A McIlraith, and Yilun Du. Multi-agent verification: Scaling test-time compute with multiple verifiers. *arXiv preprint arXiv:2502.20379*, 2025.
  - Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
  - Runze Liu, Junqi Gao, Jian Zhao, Kaiyan Zhang, Xiu Li, Biqing Qi, Wanli Ouyang, and Bowen Zhou. Can 1b llm surpass 405b llm? rethinking compute-optimal test-time scaling. *arXiv preprint arXiv*:2502.06703, 2025a.
  - Shudong Liu, Hongwei Liu, Junnan Liu, Linchen Xiao, Songyang Gao, Chengqi Lyu, Yuzhe Gu, Wenwei Zhang, Derek F Wong, Songyang Zhang, et al. Compassverifier: A unified and robust verifier for llms evaluation and outcome reward. *arXiv* preprint arXiv:2508.03686, 2025b.
  - Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. In *Conference on Language Modeling (COLM)*, 2025c.
  - Zijun Liu, Peiyi Wang, Runxin Xu, Shirong Ma, Chong Ruan, Peng Li, Yang Liu, and Yu Wu. Inference-time scaling for generalist reward modeling. *arXiv preprint arXiv:2504.02495*, 2025d.
  - Dakota Mahan, Duy Van Phung, Rafael Rafailov, Chase Blagden, Nathan Lile, Louis Castricato, Jan-Philipp Fränken, Chelsea Finn, and Alon Albalak. Generative reward models. *arXiv preprint arXiv:2410.12832*, 2024.
  - Yujun Mao, Yoon Kim, and Yilun Zhou. Champ: A competition-level dataset for fine-grained analyses of llms' mathematical reasoning capabilities. In *Findings of the Association for Computational Linguistics ACL* 2024, pp. 13256–13274, 2024.
  - David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.
  - Jon Saad-Falcon, E Kelly Buchanan, Mayee F Chen, Tzu-Heng Huang, Brendan McLaughlin, Tanvir Bhathal, Shang Zhu, Ben Athiwaratkun, Frederic Sala, Scott Linderman, et al. Shrinking the generation-verification gap with weak verifiers. *arXiv* preprint arXiv:2506.18203, 2025.
  - Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652, 2023.
  - Nishad Singhi, Hritik Bansal, Arian Hosseini, Aditya Grover, Kai-Wei Chang, Marcus Rohrbach, and Anna Rohrbach. When to solve, when to verify: Compute-optimal problem solving and generative verification for llm reasoning. *arXiv preprint arXiv:2504.01005*, 2025.
  - Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
  - Benedikt Stroebl, Sayash Kapoor, and Arvind Narayanan. Inference scaling flaws: The limits of llm resampling with imperfect verifiers. *arXiv preprint arXiv:2411.17501*, 2024.
  - Sijun Tan, Siyuan Zhuang, Kyle Montgomery, William Y Tang, Alejandro Cuadron, Chenguang Wang, Raluca Ada Popa, and Ion Stoica. Judgebench: A benchmark for evaluating llm-based judges. *arXiv preprint arXiv:2410.12784*, 2024.

- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya
   Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al.
   Gemma 2: Improving open language models at a practical size. arXiv preprint arXiv:2408.00118,
   2024.
- Qwen Team. Qwen2 technical report. arXiv preprint arXiv:2407.10671, 2024.
  - Han Wang, Archiki Prasad, Elias Stengel-Eskin, and Mohit Bansal. Soft self-consistency improves language model agents. *arXiv preprint arXiv:2402.13212*, 2024a.
  - Peifeng Wang, Austin Xu, Yilun Zhou, Caiming Xiong, and Shafiq Joty. Direct judgement preference optimization. *arXiv preprint arXiv:2409.14664*, 2024b.
  - Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
  - Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multitask language understanding benchmark. *Advances in Neural Information Processing Systems*, 37: 95266–95290, 2024c.
  - Jason Wei. The asymmetry of verification and verifier's law. https://www.jasonwei.net/blog/asymmetry-of-verification-and-verifiers-law, 2025.
  - Chenxi Whitehouse, Tianlu Wang, Ping Yu, Xian Li, Jason Weston, Ilia Kulikov, and Swarnadeep Saha. J1: Incentivizing thinking in llm-as-a-judge via reinforcement learning. *arXiv* preprint *arXiv*:2505.10320, 2025.
  - Austin Xu, Yilun Zhou, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. J4r: Learning to judge with equivalent initial state group relative policy optimization. *arXiv preprint arXiv:2505.13346*, 2025.
  - An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
  - An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.
  - Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. Physics of language models: Part 2.1, grade-school math and the hidden reasoning process. In *The Thirteenth International Conference on Learning Representations*, 2025.
  - Weihao Yu, Zihang Jiang, Yanfei Dong, and Jiashi Feng. Reclor: A reading comprehension dataset requiring logical reasoning. *arXiv preprint arXiv:2002.04326*, 2020.
  - Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative verifiers: Reward modeling as next-token prediction. In *The Thirteenth International Conference on Learning Representations*, 2025.
  - Yilun Zhou, Austin Xu, Peifeng Wang, Caiming Xiong, and Shafiq Joty. Evaluating judges as evaluators: The jetts benchmark of llm-as-judges as test-time scaling evaluators. *arXiv preprint arXiv:2504.15253*, 2025.
  - Mingchen Zhuge, Changsheng Zhao, Dylan Ashley, Wenyi Wang, Dmitrii Khizbullin, Yunyang Xiong, Zechun Liu, Ernie Chang, Raghuraman Krishnamoorthi, Yuandong Tian, et al. Agent-as-a-judge: Evaluate agents with agents. *arXiv preprint arXiv:2410.10934*, 2024.

# **APPENDIX**

# THE USE OF LARGE LANGUAGE MODELS (LLMS)

LLMs were used exclusively as writing assistance tools in preparing this manuscript. Specifically, we employed LLMs for grammar checking. An LLM-based image generation tool was used to create the robot caricature in Figure 1. All research ideation, experimental design, analysis, and scientific conclusions are entirely the work of the authors. The LLMs played no role in the conception of research questions, methodology development, or interpretation of results. Authors take full responsibility for all content in this paper, including any text refined with LLM assistance.

# A PROMPT TEMPLATES

**Response Generation Prompt.** Here, we provide the prompts to generate model responses to questions from three domains. For each model, we use its default system prompt as specified in the model documentation.

# ### User Prompt {problem} Please reason step by step, and put your final answer within \\boxed{{}}.

```
### User Prompt
The following are multiple choice questions (with answers) about {category}. Think step
by step and then output the answer in the format of \"The answer is (X)\" where X is the
correct letter choice.

Question: {question}
Options:
{options}
Answer: Let's think step by step.
```

# Natural Language Reasoning: ReClor ### User Prompt Answer the following multiple choice question. The last line of your response should be of the following format: 'Answer: \$LETTER' (without quotes) where \$LETTER is one of ABCD. Think step by step before answering. {context} {question} {choices} Output your answer strictly following this format: Reasoning: <step-by-step reasoning> Answer: <your choice>

```
Natural Language Reasoning: FOLIO

### User Prompt
{question}

Based on the above, can the following be inferred?
{conclusion}
```

```
Think step-by-step before outputting your answer, which must be "Yes", "No", or "Uncertain"

Output your answer strictly following this format:
Reasoning: Your reasoning here
Answer: Yes or No or Uncertain
```

```
Natural Language Reasoning: GPQA Diamond

### User Prompt
Answer the following multiple choice question. The last line of your response should be of the following format: 'Answer: $LETTER' (without quotes) where $LETTER is one of ABCD. Think step by step before answering.

Question: {question}
Options: {choices}
Output your answer strictly following this format:
Reasoning: <step-by-step reasoning>
Answer: <your choice>
```

**Verification Evaluation Prompt.** Below we present the prompt template used to evaluate verification performance.

```
### System Prompt
Please act as an impartial judge and evaluate the correctness of the response provided
by an AI assistant to the user prompt displayed below. You will be given the assistant's
response.
When evaluating the assistant's response, identify any mistakes or inaccurate
information. Be as objective as possible. Avoid any biases, such as order of responses,
length, or stylistic elements like formatting.
Before providing an your final verdict, think through the judging process and output
your thoughts as an explanation
After providing your explanation, you must output only one of the following choices as
your final verdict with a label:
1. The response is correct: [[Correct]]
2. The response is incorrect: [[Incorrect]]
Use the following template:  \\
Explanation: Your detailed thought process as an explanation.
Verdict: [[Correct]] or [[Incorrect]].
### User Prompt
<|User Prompt|>
{question}
<|The Start of Assistant's Answer|>
{response}
<|The End of Assistant's Answer|>
```

**Fallback Correction Check Prompt.** When Math-Verify returns unparsable or incorrect results, we employ LLM-as-judge as a fallback mechanism for correctness verification. Below, we provide the prompt template used for this secondary verification step:

```
### User Prompt
Given a math problem, its correct answer, and the model's generated answer, determine if
the model's generated answer is correct.

VALIDATION CRITERIA:
1. Identify the final answer, which is usually put inside \\boxed{{answer}} or
**answer**.
2. The answer must be mathematically equivalent to the correct answer
```

```
810
           3. The answer must be complete with a clear final result
811
           4. The answer must not just contain similar numbers - it must reach the correct
812
           5. If the generated answer contains multiple different final answers or is ambiguous
813
           about which is the final answer, mark it as 'False'
814
           IMPORTANT: Just having the same numbers as the ground truth is NOT sufficient - the
815
           model must actually solve the problem correctly and provide the correct final answer in
816
           the designated format.
817
           Respond with 'True' if the answer is correct and complete, and 'False' if it is
818
           incorrect or incomplete.
           Directly provide your judgement 'True' or 'False' without any other description.
819
820
           Problem: {problem}
           Correct Answer: {ground_truth_answer}
821
           Model's Generated Answer: {model_response}
822
           Your judgement:
823
```

# B ADDITIONAL PRELIMINARIES AND SETUP

# B.1 Details of Math Problems Correctness Check

Here we detail the evaluation procedure for establishing response correctness, including fallback methods. Ground-truth correctness is determined using Math-Verify (Kydlíček, 2025). If Math-Verify fails to parse an answer or returns incorrect, we recheck with other string-matching verifiers from open-source repositories lm-eval (Gao et al., 2024), Dr. GRPO (Liu et al., 2025c), and Qwen2.5-Math (Yang et al., 2024). We further apply GPT-4.1-mini and Qwen2.5-72B to conduct reference-based evaluation and check the equivalence of the model prediction and ground-truth answers. The prompt template for LLM-based verification is provided in Appendix A.

#### **B.2** Model Naming Conventions

Throughout this paper, we use abbreviated model names in figures and tables to improve readability and space efficiency. Table 1 provides the complete mapping between abbreviations and full model names. All models referenced are instruction-tuned versions unless otherwise specified.

Table 1: Mapping between abbreviated model names used in figures and their full names. All models are instruction-tuned versions.

Abbreviation	Full Model Name	Abbreviation	Full Model Name
G2-2B	Gemma2-2B	Q3-4B	Qwen3-4B
G2-9B	Gemma2-9B	Q3-8B	Qwen3-8B
G2-27B	Gemma2-27B	Q3-32B	Qwen3-32B
L3-3B (L-3B)	Llama3.2-3B	M-8B	Ministral-8B
L3-8B (L-8B)	Llama3.1-8B	M-24B	Mistral-Small-24B
L3-70B (L-70B)	Llama3.3-70B	gpt-4o	GPT-4o
Q2.5-3B	Qwen2.5-3B		
Q2.5-7B	Qwen2.5-7B		
Q2.5-72B	Qwen2.5-72B		

# B.3 DETAILED EXPERIMENTAL SETUP OF RQ2

Due to significant differences in generator capability, when measuring TPR, for some very difficult problems, none of the 64 responses sampled from a weak model are correct. To ensure fair evaluation unaffected by intrinsic problem difficulty, we exclude these problems and keep only those where every generator produces at least one correct response. We apply analogous filtering for TNR, keeping only problems where each generator produces at least one incorrect response. Beyond filtering problems, we also carefully balance how many responses we evaluate from each generator.

As described in Section 3, we subsample 8 responses from each generator's 64-sample pool for verification evaluation, aiming for 4 correct and 4 incorrect when possible. However, across these 8-response subsets, stronger generators may have produced more correct responses than weaker ones. This would bias our metrics by creating different denominators per generator. To address this, we randomly select one correct response per problem from each generator's 8-response pool when computing TPR (and analogously for TNR). We repeat this evaluation with random selections eight times and report the mean.

# .

# C ADDITIONAL RESULTS

#### C.1 More Details in Problem Difficulty Analysis

In Section 4.1, we show that problem difficulty primarily influences the verifier's ability to recognize correct responses. As discussed in the main paper, our analysis is conducted at two levels of granularity: response level and problem level. Figure 2 shows results at the response level. Figures 7 and 8 show results at the problem level, summarizing the distribution of TPR and TNR across difficulty quartiles. Together, these results confirm our main finding that problem difficulty strongly correlates with TPR but has no systematic effect on TNR.

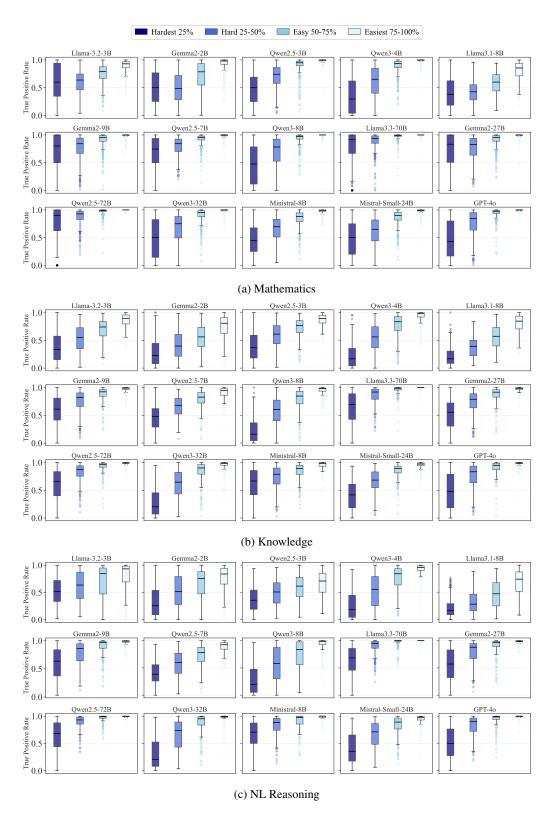


Figure 7: **Problem difficulty correlates with verification TPR on per-problem level across three domains.** Each boxplot shows the distribution of per-problem TPR for 15 verifier models, grouped by difficulty quartiles. TPR exhibits a strong positive correlation with problem easiness: easier problems consistently yield higher and less variable TPR.

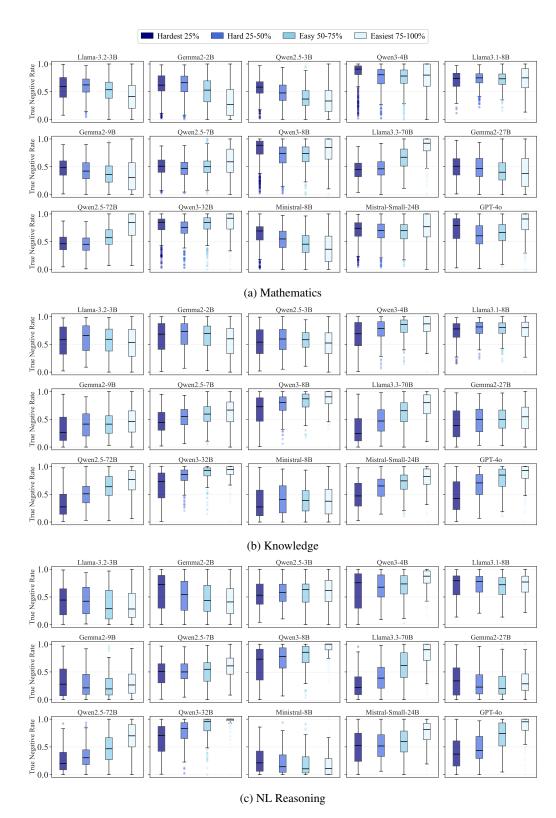


Figure 8: Problem difficulty shows no systematic correlation with verification TNR on perproblem level across three domains. Each boxplot shows the distribution of per-problem metrics for 15 verifier models, grouped by difficulty quartiles. TNR doesn't show obvious correlation with problem difficulty, exhibiting inconsistent trends across models.

#### C.2 ANALYSIS ON REASONING MODEL

 Our main analysis focuses on instruction-tuned models, which represent the typical setting for verification systems in current practice, including recent judge models (Tan et al., 2024; Wang et al., 2024b) and verifier work (Liu et al., 2025b; Zhang et al., 2025). We prioritize models without extensive CoT reasoning because verification often demands low-latency solutions, particularly for reinforcement learning training and TTS applications where rapid evaluation is critical. However, a recent trend involves training long-reasoning evaluators (Chen et al., 2025d; Whitehouse et al., 2025) that generate extended CoT before making verification decisions. To examine whether our findings generalize to this emerging paradigm, we include two reasoning models (Qwen3-8B-Thinking and Qwen3-32B-Thinking²) and analyze how they perform across our research questions. These models generate longer reasoning traces before producing binary verdicts, representing the state-of-the-art in reasoning-enhanced verification.

In Figures 9 and 10, we evaluate the conclusion of **RQ1** (Section 4.1). We observe that reasoning models exhibit the same TPR pattern as instruction-tuned models: easier problems consistently yield higher TPR across all three domains. This indicates that the fundamental relationship between problem difficulty and correctness recognition persists with extended reasoning. However, reasoning models exhibit a notable difference in TNR behavior. Unlike instruction-tuned models, where TNR showed no systematic relationship with problem difficulty, both reasoning models demonstrate improved TNR as problems become easier across all three domains. This pattern suggests that, with extended reasoning, error detection becomes easier when problems become easier.

In Figure 11, we evaluate the findings of **RQ2** (Section 4.2) on reasoning models and find that they maintain the core patterns observed in instruction-tuned models. TPR remains consistently high with mild increases as generator strength increases, while TNR decreases more significantly (goes from red to white) with stronger generators. This indicates that the fundamental challenge of detecting errors from capable generators persists despite enhanced reasoning capabilities.

These findings demonstrate that reasoning models offer some advantages for error detection on easier problems while preserving the core verification dynamics we identified. Problem difficulty continues to govern correctness recognition, and generator capability primarily influences error detectability across different verification paradigms.

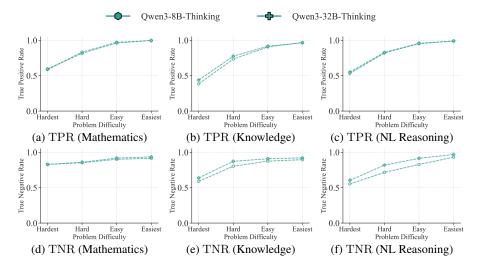


Figure 9: Verification performance of reasoning models across problem difficulty at the perresponse level. TPR (a-c) and TNR (d-f) for Qwen3-8B-Thinking and Qwen3-32B-Thinking across difficulty quartiles in three domains. Both reasoning models show increasing TPR and TNR as problem difficulty decreases.

<sup>&</sup>lt;sup>2</sup>We use the suggested sampling hyperparameter (temperature 0.6, top-p 0.95).

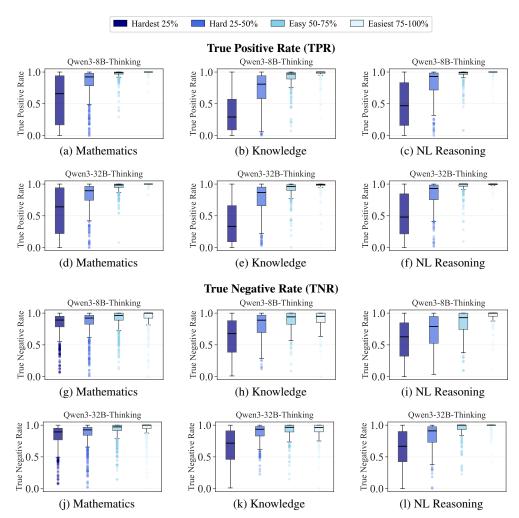


Figure 10: Verification metrics for reasoning models across difficulty quartiles at the perproblem level. Each boxplot shows the distribution of per-problem TPR and TNR for Qwen3-8B-Thinking and Qwen3-32B-Thinking across difficulty quartiles in three domains. Both TPR and TNR distributions shift higher and become less variable as problems become easier.

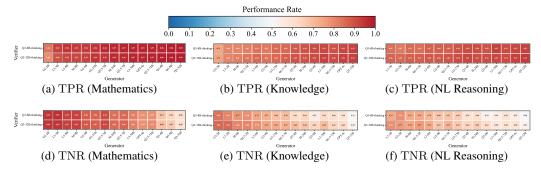


Figure 11: **Reasoning models as verifiers paired with generators of varying capability.** TPR (a-c) and TNR (d-f) for Qwen3-8B-Thinking and Qwen3-32B-Thinking verifiers when evaluating responses from 15 generator models across three domains. Generators are ordered left-to-right by increasing generation capability, measured separately for each domain. Red indicates higher performance, blue indicates lower performance.

# C.3 EXPLANATION OF BELOW-RANDOM VERIFICATION PERFORMANCE ON HARD PROBLEMS

In Figure 4f, we observe that verifiers achieve balanced accuracy below the random baseline of 0.5 on hard problems from the NL Reasoning domain, a result that needs explanation. This phenomenon can occur in reference-free evaluation when verifiers employ a "solve-and-match" verification strategy, where they attempt to solve the problem independently and then compare their answer with the generator's response.

For NL Reasoning tasks with 3-way or 4-way multiple choice formats, this mechanism can produce below-random performance when verifiers consistently fail to solve hard problems correctly. In such cases, the verifier never correctly identifies true positive responses (TPR=0) because it always produces wrong answers that don't match correct generator responses. However, it can still identify some true negatives when both the generator and verifier happen to select the same wrong answer. For three-way choices, the TNR=0.5. With TPR near zero and TNR remaining positive, the balanced accuracy falls below 0.5.

This phenomenon is specific to tasks with limited answer spaces. The affected problems are those in the hard set with d(x) < 0.3, where even strong models achieve very low pass rates. It occurs in NL Reasoning because this domain includes three-way multiple-choice questions from datasets like FOLIO. It does not occur in Mathematics, where responses are open-ended strings, or in Knowledge domains with 10-way multiple choice, where the large answer space dilutes the effect.

#### C.4 ADDITIONAL RESULTS OF VERIFIER GENERATION CAPABILITY

Here we present additional results for **RQ3** from Section 4.3, providing correlation analysis between verifier generation capability and verification accuracy across the entire problem difficulty range in Figure 12. The results confirm our finding from the main paper that the correlation form varies with problem difficulty: medium problems show strong positive linear relationships, while hard and easy problems exhibit non-linear trends.

These findings highlight the need for regime-aware verifier strategies. On hard problems, strong verifiers are unnecessary as performance plateaus regardless of capability. On medium problems, selecting models with better generation capability consistently yields better verification. On easy problems, selecting higher-capability models works well among weak-to-medium verifiers, but strong models with similar capabilities show vastly different verification performance. Thus, optimal selection of strong verifiers requires supplementary benchmarking or alternative evaluation metrics.

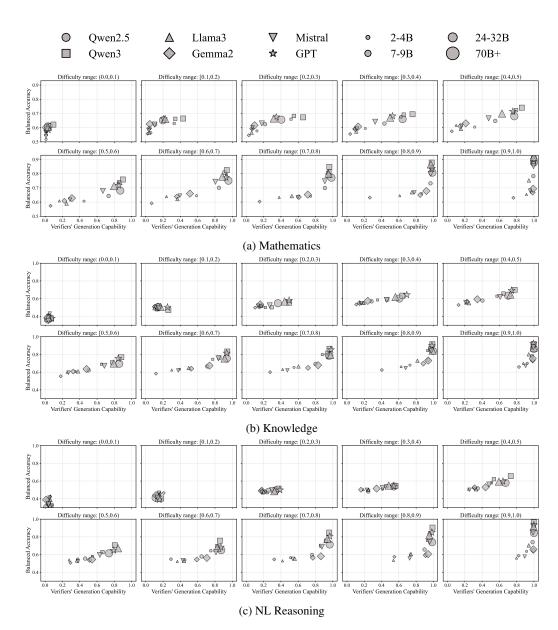


Figure 12: Correlation between verification performance and generation capability across problem difficulty ranges on three domains. Balanced accuracy as a function of verifier generation capability for difficulty ranges from (0.0,0.1) to [0.9,1.0). Performance exhibits three distinct regimes: plateaus on hard problems, strong positive correlation on medium problems, and high variance with saturated capability on easy problems. Marker shapes indicate model family; sizes represent model scale.

#### C.5 ADDITIONAL RESULTS OF GENERATOR ANALYSIS IN TEST-TIME SCALING

This subsection provides complementary results for the generator analysis presented in Section 5.1, demonstrating the generalizability of our findings across domains and problem difficulties. Figure 13 extends the analysis from the main paper to Knowledge and NL Reasoning domains. Our central finding from **RQ4** holds consistently. As Figures 13a to 13c shows, verification gains peak at weak-medium generator strength, enabling these generators to substantially close performance gaps with stronger models. The underlying mechanism driving this phenomenon, identified in **RQ2**, remains consistent across domains. Figures 13d to 13f shows that, as generator strength increases, TNR decreases sharply while TPR rises only modestly. For the strongest generators, the collapsed TNR limits verification gains as errors become increasingly difficult to detect. This brings high

 verification gains at weak-medium generator levels. In the main paper, we show results on problems with difficulty range  $d(x) \in [0.7, 0.8)$  in Figure 5. Here, Figures 14 to 16 report results across the entire difficulty range for three domains, respectively. Figure 17 shows the percentage of performance gap closed by verification for all weak-to-strong generator pairs, computed on all problems within each domain.

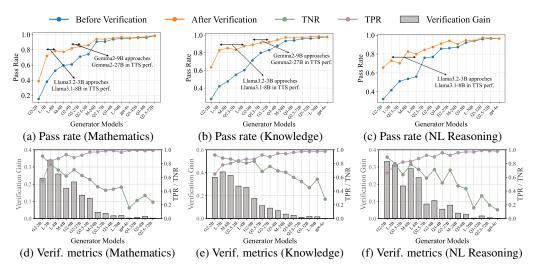


Figure 13: **TTS** performance before and after verification when sweeping generator strength. (a-c) Pass rate before (blue) and after (orange) adding a fixed verifier (GPT-40), across generators ordered from weaker (left) to stronger (right) by generation capability. (d-f) Bar chart shows the verification gain  $\Delta \hat{p}_V$  (left y-axis) for each generator. Lines show the verifier's TNR and TPR on the same datasets (right y-axis). Results are reported on problems with difficulty in the range [0.7, 0.8) for three domains. Problem counts across domains: 181 (Mathematics), 154 (Knowledge), 97 (NL Reasoning).

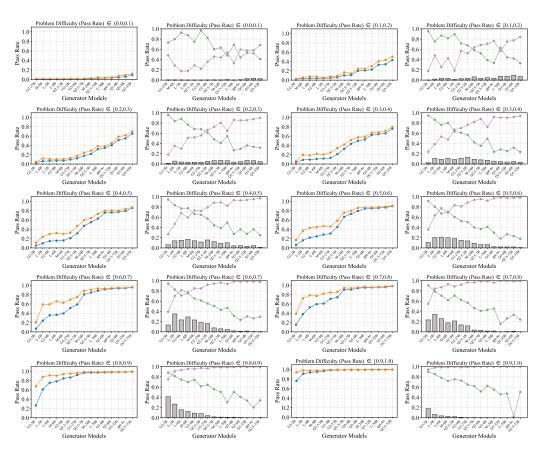


Figure 14: Verification-augmented TTS performance across the full range of problem difficulties, shown here for the Mathematics domain. Each pair of figure corresponds to a different difficulty interval (measured by pass rate d(x)), with the left panel showing pass rates before (blue) and after (orange) verification, and the right panel showing verification gain  $\Delta \hat{p}_V$  (bars) alongside the verifier's TNR (green) and TPR (purple). Compared to Figure 5, which focused only on problems with  $d(x) \in [0.7, 0.8)$ , this includes the entire difficulty range.

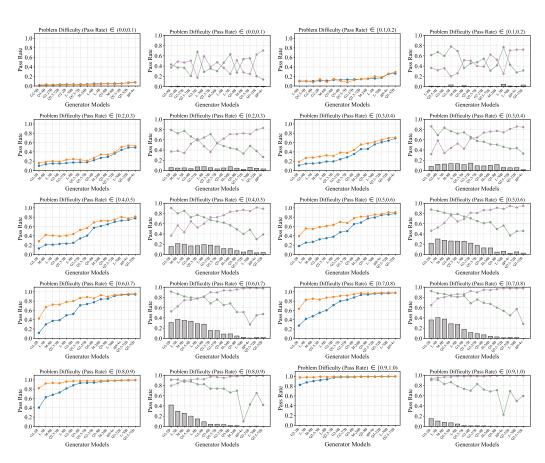


Figure 15: Verification-augmented TTS performance across the full range of problem difficulties, shown here for the Knowledge domain. Each pair of figure corresponds to a different difficulty interval (measured by pass rate d(x)), with the left panel showing pass rates before (blue) and after (orange) verification, and the right panel showing verification gain  $\Delta \hat{p}_V$  (bars) alongside the verifier's TNR (green) and TPR (purple). Compared to Figure 5, which focused only on problems with  $d(x) \in [0.7, 0.8)$ , this includes the entire difficulty range.

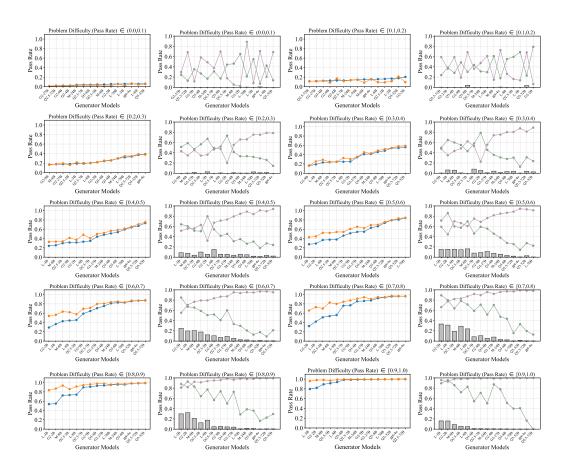


Figure 16: Verification-augmented TTS performance across the full range of problem difficulties, shown here for the NL reasoning domain. Each pair of figure corresponds to a different difficulty interval (measured by pass rate d(x)), with the left panel showing pass rates before (blue) and after (orange) verification, and the right panel showing verification gain  $\Delta \hat{p}_V$  (bars) alongside the verifier's TNR (green) and TPR (purple).

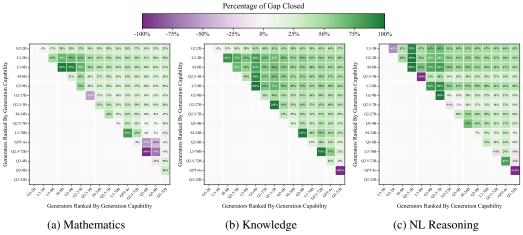


Figure 17: Percentage of TTS performance gap between weak and strong generators closed by verification. Each heatmap shows the fraction of the performance gap between a weaker generator (x-axis) and a stronger generator (y-axis) that is closed by verification with a fixed verifier GPT-40. Green cells indicate a larger gap closure, meaning the weaker model approaches the stronger one after verification. A value greater than 100% means that the originally weaker model performs better with verifier augmentation. Purple cells indicate negative values where verification increases the gap.

#### C.6 ADDITIONAL RESULTS OF VERIFIER ANALYSIS IN TEST-TIME SCALING

Here we provide complementary results to the verifier analysis in Section 5.2. Figure 18 presents two other domains' results. Figure 19 presents additional metrics in the same setup for complete analysis, including balanced accuracy and verification gains. We can see that our findings from Mathematics generalize to other domains. Figures 18a to 18c shows that weak verifiers can approximate strong verifier performance in TTS, at the extremes of problem difficulty or responses generated by strong generators. While on these regimes, we show that both verifiers provide limited verification gain in Figures 19a to 19c. Notably, in the NL Reasoning domain (Figure 18i), on the strongest generator, both verifiers' TNR fall below 0.5, and the weak verifier's TNR exceeds that of the strong verifier. Despite this TNR inversion, the overall verification performance gap (shown as balanced accuracy in Figure 19f) remains narrow, with the strong verifier maintaining a slight advantage due to its superior TPR.

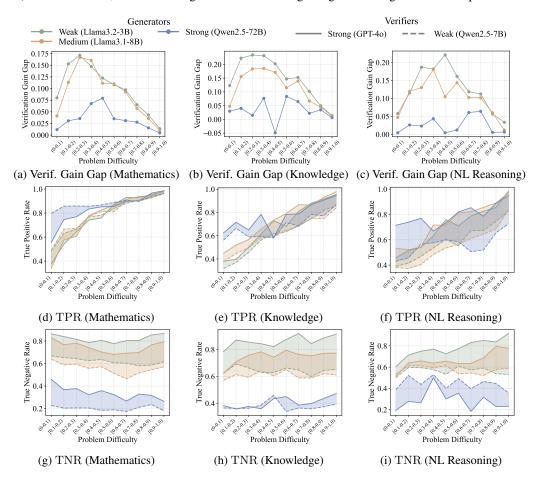


Figure 18: Analyzing verification gain gaps and TPR/TNR between verifiers under varying problem difficulty and generator strength. The x-axis shows problem difficulty measured relative to each generator. Shaded regions visualize the difference in metrics between verifiers for each generator. (a-c) Verification gain gap between strong and weak verifiers. (d-f) TPR increases as problems become easier for all generator-verifier combinations. (g-i) TNR decreases as generators become stronger, with TNR gap narrowing.

1513

1515

1516

1517

1518

1519 1520 1521

1522

1531 1532

1533

1534

1535

1536

1537 1538 1539

1540 1541

1565

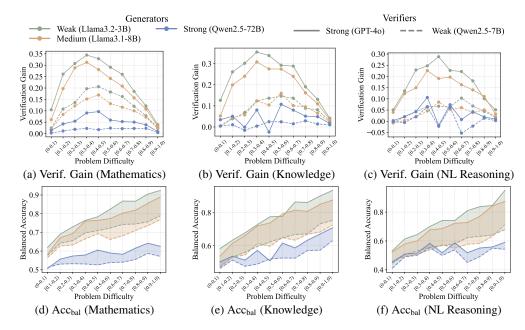


Figure 19: Analyzing verification gains and Acc<sub>bal</sub> between verifiers under varying problem difficulty and generator strength. The x-axis shows problem difficulty measured relative to each generator. Shaded regions visualize the difference in metrics between verifiers for each generator. (a-c) Verification gain for both strong (solid lines) and weak (dashed lines) verifiers across three generators: weak, medium, and strong. (d-f) Balanced accuracy for the same verifier-generator combinations. The accuracy gap (shaded regions) between verifiers is smallest on the hard problems (left side of the x-axis).

# D CASE STUDIES

```
1542
                Problem domain, dataset and difficulty: Mathematics, MATH-500, Hardest (d(x) = 0.075)
                Generator: Owen2.5-3B
1543
                Verifier: GPT-40
                Question: Bob and Alice each have a bag that contains one ball of each of the colors, blue, green, orange, red, and violet. Alice
                randomly selects one ball from her bag and puts it into Bob's bag. Bob then randomly selects one ball from his bag and puts it into
                Alice's bag. What is the probability that after this process the contents of the two bags are the same?
                Ground-Truth Solution: After Alice puts the ball into Bob's bag, his bag will contain six balls: two of one color and one of each of
1547
                the other colors. After Bob selects a ball and places it into Alice's bag, the two bags will have the same contents if and only if Bob
                picked one of the two balls in his bag that are the same color. Because there are six balls in the bag when Bob makes his selection, the
1548
1549
                probability of selecting one of the same colored pair is 2/6
1550
                Generator Output (Correct): To determine the probability that the contents of both Bob's and Alice's bags are the same after the
1551
                process described, we need to analyze each possible step and its impact on the contents of the bags.
                Initially, each bag contains one ball of each of the colors: blue, green, orange, red, and violet.
1552
                ### Step-by-Step Analysis [ Details omitted ]
1553
                ### Identifying the Conditions for Equality [ Details omitted ]
1554
                ### Calculating the Probability [ Details omitted ]
1555
                The probability is then calculated as follows:
                                                                    Number of favorable outcomes
1556
                                                     Probability =
                                                                       Total number of outcomes
1557
                Thus, the probability that after this process the contents of the two bags are the same is
                Verifier Output (False negative): Explanation: The problem involves two steps: Alice transferring a ball to Bob's bag and then Bob
                transferring a ball back to Alice's bag. Initially, both bags have one ball of each color. After Alice transfers a ball, Bob's bag will have
1561
                two balls of one color and one of each of the other four colors. For the bags to be identical after Bob's transfer, Bob must return the
                same color ball that Alice initially transferred. This is the only scenario where the bags will end up identical, as it restores the original
                distribution of colors in both bags.
1563
```

The assistant's response correctly identifies that there are 6 possible outcomes for Bob's selection (since he has 6 balls to choose from) and that there are 2 favorable outcomes (returning one of the two balls of the color that Alice transferred). However, the assistant incorrectly concludes that the probability is 1/3. The correct probability should be 1/5, as Bob must specifically return the ball of the

1592

1593

1566 color that Alice transferred, which is only one specific ball out of the 6 in his bag. [Verifier generates a wrong answer 1/5 and uses it to 1567 compare with generator's correct answer 1/3] 1568 Verdict: [[Incorrect]] 1569 1570 1571 **Problem domain, dataset and difficulty**: Natural Language Reasoning, FOLIO, Hardest (d(x) = 0.122) 1572 Generator: Llama3.1-8B Verifier: Qwen2.5-72B 1573 Question: Some mammals have teeth. Platypuses have no teeth. Platypuses are mammals. Humans have teeth. 1574 Based on the above, can the following be inferred? Platypuses are reptiles 1575 **Ground-Truth Solution:** Uncertain Generator Output (Correct): Reasoning: 1. The statement that some mammals have teeth and humans have teeth suggests that having 1576 teeth is a characteristic of mammals, but it's not exclusive to them. 2. Platypuses are indeed mammals, as stated in the information provided. 3. The fact that platypuses have no teeth does not contradict the fact that they are mammals, as having no teeth is also a characteristic of some mammals, as mentioned in point 1. 4. The only group of animals that are not mammals and have been ruled out as a category for platypuses are not the mammals, but rather the reptiles. 1579 However, concluding that platypuses are reptiles because they have no teeth is not a sound inference, because having no teeth does not preclude an animal from being a mammal, as shown in point 3. Moreover, it contradicts the given statement that platypuses are indeed 1581 5. Therefore, platypuses being reptiles cannot be inferred based on the given information. Answer: Uncertain Verifier Output (False negative): Explanation: The assistant's response is mostly correct in its reasoning but contains a minor 1584 inaccuracy in the final conclusion. The assistant correctly identifies that the fact that platypuses have no teeth does not make them reptiles, as some mammals also lack teeth. However, the assistant's final answer of "Uncertain" is not appropriate. Given the information 1585 provided, it is clear that platypuses are mammals, and there is no evidence to suggest they are reptiles. Therefore, the correct answer 1586 should be "No" because the statement "Platypuses are reptiles" can be definitively refuted based on the given information. [Although the 1587 verifier recognizes that the generator's reasoning process is correct, it rejects the conclusion due to a mismatch with its own (incorrect) answer.] Verdict: [[Incorrect]] 1590

Figure 20: Case study: How problem difficulty impacts TPR (RQ1). False negative verification on high-difficulty problems demonstrates why TPR decreases as problems become harder. When verifiers cannot correctly solve challenging problems, they produce incorrect reference solutions and reject generators' correct answers based on these flawed references.

```
1594
1595
                Problem domain, dataset and difficulty: Mathematics, MATH-500, Easy (d(x) = 0.571)
1596
                Generator: Llama-3.1-8B
                Verifier: Qwen2.5-72B
                 Question: Paul and Jesse each choose a number at random from the first six primes. What is the probability that the sum of the numbers
1598
                Ground-Truth Solution: [Details omitted] Therefore, the probability that the sum of the numbers Paul and Jesse choose IS even is
                             13
                             18
                Generator Output (Incorrect): To find the probability that the sum of the numbers Paul and Jesse choose is even, we need to first
                understand what numbers can make an even sum.
                The sum of two numbers is even if both numbers are either even or odd.
1604
                Since the first six primes are 2, 3, 5, 7, 11, and 13, we have only one even prime, which is 2.
                Now, let's consider the possible scenarios:
1605
                1. Both Paul and Jesse choose the number 2. In this case, the sum is even.
                2. One of Paul or Jesse chooses the number 2, and the other chooses an odd prime. In this case, the sum is odd.
                3. Both Paul and Jesse choose odd primes. In this case, the sum is even. [Generator correctly identifies two cases for even sum, one case
                for odd sum.]
1609
                [ Probability calculations omitted omitted ]
1610
                Now, let's add the probabilities of the first two scenarios (both choose 2, and one chooses 2 and the other an odd prime) and the third
                scenario (both choose an odd prime) to get the total probability of an even sum:
1611
                Probability = 1/36 + 10/36 + 25/36 Probability = 36/36 Probability = 1
1612
                The probability that the sum of the numbers Paul and Jesse choose is even is 1 or 100%. This is because it is impossible for the sum to
1613
                be odd, as they will always choose either two even numbers or two odd numbers. [Generator makes self-contradiction: Earlier stated
1614
                  cenario 2 gives odd sum, now count scenario 2 as even sum and claims odd sum is impossible]
1615
                Verifier Output (True negative): The assistant's response contains a logical error in the final calculation. The assistant correctly
                identifies the scenarios where the sum of the numbers chosen by Paul and Jesse is even. However, the assistant incorrectly adds
1616
                the probabilities of all scenarios, including the scenario where the sum is odd, which should not be included in the final probability
1617
                calculation. [Verifier catches the self-contradiction of generator.]
1618
                The correct calculation should only include the probabilities of the scenarios where the sum is even: - Both choose 2: \frac{1}{36} - Both choose
                odd primes: \frac{25}{36}
1619
```

1658

1659

1662

1663

1664 1665

1671

1673

```
1620
                The probability of the sum being even is the sum of these two probabilities:
1621
                                                                      \frac{1}{36} + \frac{25}{36} = \frac{26}{36} =
1622
1623
1624
                                                       13
                                                             [Verifier can solve the problem correctly]
                Therefore, the final answer should be
1625
                                                       18
1626
                Verdict: [[Incorrect]]
1627
1628
1629
                Same problem and verifier as above with a strong generator:
                Generator: GPT-40
                Verifier: Qwen2.5-72B
                Generator Output (Incorrect): To find the probability that the sum of the numbers chosen by Paul and Jesse from the first six prime
                numbers is even, we proceed as follows:
                1. **List the first six prime numbers**: The first six prime numbers are 2, 3, 5, 7, 11, and 13.
1633
                2. **Identify even and odd primes**:
                - Even prime: 2
1635
                - Odd primes: 3, 5, 7, 11, 13
                3. **Understand when the sum is even**:
                - A sum is even if both numbers are even or both numbers are odd.
1637
                 Since there is only one even prime (2), for the sum to be even, both numbers must be odd. [Fatal error: Ignores that both choosing 2
                (even+even=even) is also possible
1639
                4. **Calculate the total number of ways to choose numbers**: [ Calculation details omitted omitted ]
1640
                5. **Calculate the number of ways to get an even sum**: [ Calculation details omitted ]
                6. **Calculate the probability**: [ Calculation details omitted omitted ]
1641
1642
                                                                                                   [Final answer is wrong because following steps
                Thus, the probability that the sum of the numbers they choose is even is
1643
                consistently based on fatal error made above.]
1644
                Verifier Output (False positive): Explanation: The assistant's response is logically sound and mathematically accurate. Here's a
1645
                breakdown of the steps
                1. **Identifying the first six primes**: The assistant correctly identifies the first six prime numbers as 2, 3, 5, 7, 11, and 13.
1646
                2. **Condition for the sum to be even**: The assistant correctly states that for the sum to be even, both numbers must be odd (since the
1647
                only even prime is 2, and choosing 2 would make the sum odd if the other number is odd). [Verifier is convinced by the generator and
1648
                ignores the error here.]
                3. **Total number of possible outcomes **: The assistant correctly calculates the total number of outcomes as 6 * 6 = 36.
1650
                4. **Number of outcomes where both numbers are odd**: The assistant correctly identifies that there are 5 odd primes and calculates
                the number of outcomes where both numbers are odd as 5 * 5 = 25.
                5. **Probability calculation**: The assistant correctly calculates the probability as the number of favorable outcomes (25) divided by
1652
                the total number of outcomes (36), resulting in \frac{25}{36}.
                The assistant's reasoning and calculations are correct throughout the response. [Although the verifier demonstrated correct problem-
                solving when evaluating Llama3.1-8B in last example, it doesn't apply the same mechanism to GPT-40, instead accepting the GPT-40's
                incorrect reasoning.
                Verdict: [[Correct]]
1656
```

Figure 21: Case study: How generator capability influences TNR (RQ2). We compare verifier behavior on the same problem when evaluating weak (Llama3.1-8B) versus strong (GPT-4o) generators. The weak generator correctly lists all the cases but then contradicts itself in the final calculation, producing an inconsistent solution. The verifier catches this apparent mistake and rejects it. The strong generator makes a single early-stage error but maintains internal consistency throughout its reasoning, generating a well-structured and self-contained chain that appears correct. Despite the solution containing an actual error, the verifier fails to detect it and accepts the solution.