

ON THE ROLE OF TEMPERATURE SAMPLING IN TEST-TIME SCALING

Anonymous authors

Paper under double-blind review

ABSTRACT

Large language models (LLMs) can improve reasoning at inference time through test-time scaling (TTS), where multiple reasoning traces are generated and the best one is selected. Prior work shows that increasing the number of samples K steadily improves accuracy. In this paper, we demonstrate that this trend does not hold indefinitely: at large K , further scaling yields no gains, and certain *hard* questions remain unsolved regardless of the number of traces. Interestingly, we find that different sampling temperatures solve different subsets of problems, meaning single-temperature scaling explores only part of a model’s potential. We therefore propose scaling along the temperature dimension, which enlarges the reasoning boundary of LLMs. Temperature scaling enables base models to reach performance comparable to reinforcement learning (RL)-trained counterparts, without additional post-training. We further provide a comprehensive analysis of this phenomenon and design a multi-temperature voting method that reduces the overhead of temperature scaling. Overall, our findings suggest that TTS is more powerful than previously thought, and that temperature scaling offers a simple and effective way to unlock the latent potential of base models.

1 INTRODUCTION

Large language models (LLMs) have demonstrated strong reasoning capabilities for complex problems at test time (Wei et al., 2022). As illustrated in Figure 1a, two main approaches have emerged to achieve such reasoning. The first trains models to produce long reasoning traces with self-reflection and correction, often implemented through reinforcement learning (RL) (Guo et al., 2025a; Yang et al., 2025c). While effective, this approach requires costly and time-consuming training (Liu et al., 2025a). The second, known as test-time scaling (TTS) (Brown et al., 2024; Snell et al., 2025; Zhao et al., 2025), shifts the burden to inference: the model generates multiple reasoning traces in parallel and a verifier selects the most reliable one (Saad-Falcon et al., 2025a). Unlike RL, TTS requires no large-scale post-training and generates relatively short, prefix-sharing traces. This enables efficient reuse of the KV cache (Hooper et al., 2025) and offers speed advantages in modern serving systems (Kwon et al., 2023; Zheng et al., 2024).

Recent studies on TTS have shown that increasing the number of samples K can enhance reasoning performance (Schaeffer et al., 2025). As illustrated in Figure 1b, scaling K up to 1,024 shows a clear trend of steadily improving accuracy. However, when we push K further to 13,312, the improvement stops: some questions remain unsolved no matter how many samples are drawn. *If further scaling K brings no improvement, have we reached the ceiling of TTS performance?*

The answer is no. As shown in Figure 1c, we observe an interesting phenomenon: when scaling K up to 1,024 under different sampling temperatures T , the sets of solvable questions differ. For example, a question unsolved at $T = 0.5$ may become solvable at $T = 0.7$. The model’s overall solvable set is therefore larger than the solvable set under any single temperature. This indicates that scaling K at a fixed temperature only explores part of the model’s potential. To unlock the full boundary, we scale along the temperature dimension: given any budget, we divide samples evenly across multiple temperatures. As shown in Figure 1d, this multi-temperature strategy achieves a much higher performance upper bound. These results highlight the importance of temperature sampling in TTS.

In this paper, we show that temperature provides a new dimension for scaling at test time. Through extensive experiments across model sizes and datasets, we demonstrate that temperature scaling ex-

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107




Figure 1: Observations and motivation for temperature scaling in TTS. (a) RL vs. TTS: RL produces long single traces, while TTS generates multiple shorter ones. (b) $\text{Pass}@K$ and $-\log(\text{Pass}@K)$ curves at $T = 0.7$ on Qwen3-4B (AIME 2025); no gain beyond $K = 1,024$. (c) Question solvability on AIME 2025 for Qwen3-4B: different temperatures solve different subsets of questions. (d) Single-temperature vs. multi-temperature scaling: the latter expands the reasoning boundary.

pands the reasoning boundary of LLMs. This effect arises because, while all temperatures can solve *easy* questions, some *hard* questions are only solvable at specific temperatures. With temperature scaling, a base model can reach performance comparable to RL-trained models. We further conduct entropy-based analyses and case studies to understand this phenomenon. Finally, we design a multi-temperature voting method that identifies and exits *easy* questions early, making temperature scaling more efficient. Overall, our contributions are threefold:

- **Scaling test-time compute to a new dimension.** We show that scaling the sampling temperature expands the model’s reasoning boundary, revealing a new axis of test-time compute.
- **Analyzing the dynamics of temperature scaling.** Through comprehensive analysis, we show that the enlarged reasoning boundary arises because different temperatures solve different *hard* questions, while *easy* questions can be solved by all temperatures.
- **Designing efficient methods for temperature scaling.** We propose a multi-temperature voting strategy that exits *easy* questions early, reducing overhead while preserving the benefits of temperature scaling across models and datasets.

2 SCALING TEMPERATURE AT TEST TIME

In this section, we first introduce the concept of temperature sampling and describe the experimental setup for scaling temperature. We then present the performance improvements achieved through temperature scaling, and compare this approach against further scaling K and RL-based methods.

2.1 TEMPERATURE SAMPLING

Temperature sampling. Temperature-based sampling has long been studied in probabilistic modeling (Ackley et al., 1985), and it remains a core component of decoding in LLMs. At each step of autoregressive generation, the model conditions on the input x and previously generated tokens $y_{1:t-1}$ to yield a distribution over the next token, from which y_t is sampled:

$$y_t \sim p(\cdot | x, y_{1:t-1}) = \text{softmax}\left(\frac{f_\theta(x, y_{1:t-1})}{T}\right).$$

Here, $f_\theta(\cdot)$ denotes the model logits and T is the sampling temperature. The temperature rescales the logits before applying the softmax, thereby shaping the probability distribution used for sampling.

Effect of temperature. The sampling temperature T is non-negative. When $T = 0$, the distribution collapses to a delta on the maximum-logit token, equivalent to deterministic decoding. For $T > 0$, smaller values sharpen the distribution, making generation more deterministic and favoring high-probability tokens, while larger values flatten it, increasing randomness and promoting diversity.

2.2 EXPERIMENTAL SETUP

Sampling temperature and number of samples. We vary the sampling temperature from 0.0 to 1.2 in increments of 0.1. At temperature 0.0, we generate a single reasoning trace for each question, whereas at other temperatures we generate 1,024 traces per question.

Table 1: Results (%) of temperature scaling across models and datasets. Base reports Pass@1, 024 under $T = 0.6$. $+T$ reports Pass@All when scaling the sampling temperature from 0.0 to 1.2, with 1,024 samples per temperature. Δ denotes the difference between $+T$ and Base.

Models	AIME2025			AIME2024			MATH500			LiveCodeBench			Hi-ToM		
	Base	$+T$	Δ	Base	$+T$	Δ	Base	$+T$	Δ	Base	$+T$	Δ	Base	$+T$	Δ
Qwen3-8B	60.0	66.7	6.7	73.3	80.0	6.7	97.0	97.8	0.8	32.6	40.0	7.4	93.0	95.5	2.5
Qwen3-4B	60.0	73.3	13.3	66.7	76.7	10.0	95.5	98.5	3.0	36.0	40.0	4.0	83.0	92.0	9.0
Qwen3-1.7B	46.7	50.0	3.3	43.3	50.0	6.7	92.5	95.5	3.0	29.7	35.4	5.7	37.0	55.0	18.0
Qwen3-0.6B	20.0	36.7	16.7	23.3	30.0	6.7	88.1	94.0	5.9	25.7	31.4	5.7	71.5	82.5	11.0

Datasets and Prompts. We evaluate on reasoning benchmarks spanning multiple domains. AIME 2024, AIME 2025, and MATH500 (Hendrycks et al., 2021) are used to assess mathematical reasoning. LiveCodeBench v6 (Jain et al., 2025) evaluates code generation, and Hi-ToM (Wu et al., 2023) focuses on social and logical reasoning. All these benchmarks support automatic output verification. Further details about the datasets and prompts are provided in Appendix A.

Platform and models. All experiments are conducted using vLLM (Kwon et al., 2023) to enable large-batch rollouts. We run our evaluations on a cluster using 64 NVIDIA H100 GPUs. The models include the Qwen3 series (0.6B, 1.7B, 4B, 8B) (Yang et al., 2025a), and Polaris-4B-Preview (An et al., 2025), an RL-trained variant of Qwen3-4B.

Evaluation metrics. Our primary evaluation metric is Pass@ K , where K is the number of sampled traces. It measures the probability of obtaining at least one correct answer when sampling K times. Let N be the total number of generated samples and C the number of correct ones, then

$$\text{Pass}@K = 1 - \frac{\binom{N-C}{K}}{\binom{N}{K}}.$$

In addition, the average accuracy Avg@ $N = C/N$ approximates the model’s correctness probability when N is large. As our aim is to highlight scaling behavior, we use ground-truth verification instead of a reward model (RM) to avoid confounds from RM quality.

AIME 2024/2025 reasoning trace validation. Each AIME problem has an integer answer. It has been argued that LLMs may sometimes arrive at the correct answer by chance rather than through valid reasoning (Wen et al., 2025). To address this concern, we conduct additional validation on all problems where a model’s Avg@1, 024 is below 3%. For these problems, the model-generated reasoning traces and human-written reference solutions are jointly reviewed by gpt-5 as an automatic judge. Further details of this validation process are provided in Appendix A.

2.3 RESULTS ON TEMPERATURE SCALING

Main results. We compare scaling at the default temperature ($T = 0.6$) with multi-temperature scaling. As shown in Table 1, temperature scaling consistently enlarges the reasoning boundary across all models and datasets. For example, on AIME 2025, Qwen3-4B gains 13.3 points after temperature scaling, indicating that any single temperature covers only part of the model’s reasoning capability. However, the effect varies by dataset: on MATH500, where Qwen3-8B already performs strongly, the additional gain from temperature scaling is relatively small.

No single temperature works for all questions. Each question shows its own temperature preference. For example, in Figure 2a, one AIME 2025 problem requires $T = 1.1$ with 128 traces to reach Pass@ $K = 1$, while $T = 0.7$ needs 256 traces. Figure 2b further shows the temperature preferences of Qwen3-4B on all AIME 2024/2025 questions: each question has a different optimal temperature, and no single value performs best across the board. Interestingly, while higher temperatures increase diversity and creativity, we find no consistent link to improved reasoning capability.

Pass@ K vs. Avg@ N . Figure 2c shows Pass@ K scaling curves of Qwen3-4B on AIME2025 under different temperatures. When K is small, different temperatures yield similar Pass@ K , since models mainly solve *easy* questions that all temperatures can handle. As K grows, however, temperature-specific advantages emerge: a *hard* question favoring one temperature may eventually be solved there but not at others, leading to larger gaps at Pass@1, 024. In contrast, Figure 2d shows that

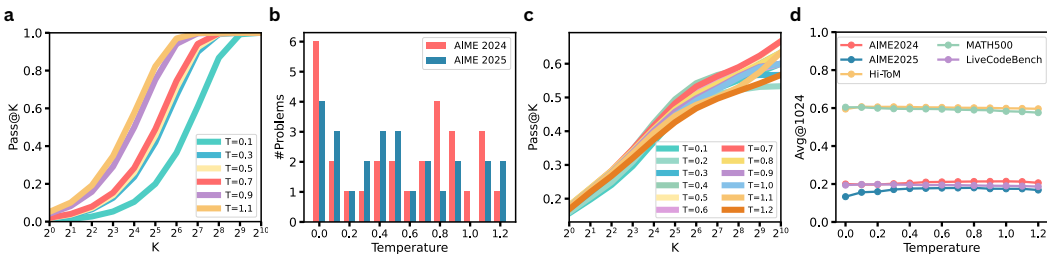


Figure 2: Scaling temperature for test-time compute on Qwen3-4B. (a) Pass@ K curves for different temperatures on AIME 2025 Q22. (b) Distribution of preferred temperatures across AIME 2024/2025. (c) Pass@ K scaling curves on AIME 2025. (d) Avg@1, 024 curves across five datasets.




Figure 3: Comparison of scaling along K and scaling along T . (a) Correlation of Avg@1, 024 across two temperatures on AIME2025, Qwen3-4B. (b) Correlation of Avg@1, 024 across two temperatures on AIME2025, Qwen3-8B. (c) Scaling K vs. scaling T on AIME2025, Qwen3-4B. (d) Temperature scaling vs. RL-trained model on AIME2025, Qwen3-4B, thinking mode.

Avg@1, 024 remains nearly identical across temperatures, as most *easy* questions are solved by all settings and the ability to solve harder ones is drowned out in this metric.

Remarks. Prior work has reported similar Avg@ N curves and argued that, within a range, temperature does not affect LLM reasoning performance (Renze, 2024). Our results suggest this is not the case in the TTS setting: Avg@ N fails to capture the model’s ability to solve *hard* questions with low probability, whereas a good verifier can still identify these sparse correct traces (Zhao et al., 2025).

2.4 SCALING ALONG T VS. SCALING ALONG K

Understanding dataset difficulty distribution and solvability. Figures 3a and 3b plot the Avg@1, 024 of each question under two different temperatures, with each point representing one problem. From this view, the questions can be grouped as:

- *Easy*: questions near the upper-right diagonal, solved by both temperatures with high probability.
- *Medium*: points near the lower-left diagonal, harder but still solved by both temperatures.
- *Hard*: points lying on the axes, solvable at one temperature with low probability but not the other.
- *Impossible*: points at the origin, unsolved regardless of temperature.

We observe that many questions cluster along the upper-right diagonal: these are simple problems that any temperature can solve, explaining why Avg@ N shows little difference across temperatures. Some questions remain at the origin: these are beyond the reach of temperature scaling and likely require training rather than TTS.

Scaling across temperatures unlocks latent capability. The axis points (*hard* problems) illustrate why temperature scaling is effective: sampling across multiple temperatures ensures that these problems enter the solvable set, whereas sampling under a single temperature collapses all solvable *hard* problems from other temperatures back to the origin.

Further scaling the number of samples does not help. As shown in Figure 1b, increasing K from 1,024 to 13,312 does not solve any additional problems. This suggests that the TTS curve has two

216 phases: an initial regime where more samples improve performance, followed by a plateau where
 217 no further gains are observed. In contrast, Figure 3c shows that scaling temperature yields a 6.67%
 218 improvement, demonstrating that temperature provides a meaningful additional dimension for TTS.

219 **Remarks.** When scaling K to 13,312, we initially observe improvements, with the model producing
 220 correct final answers for questions unsolved under a smaller budget. However, after gpt-5-based
 221 trace verification, these cases are found to be invalid: the model is merely more likely to guess the
 222 correct answer with more samples. This highlights that future work on RL and TTS should account
 223 not only for final-answer accuracy but also for the validity of reasoning traces.

225 2.5 COMPARISON WITH RL-TRAINED METHODS

226
 227 **Setup.** We evaluate the RL-trained Polaris-4B-Preview model against its base model, Qwen3-4B.
 228 Following the setup suggested by An et al. (2025), we run the RL model on AIME 2025 with
 229 temperature $T = 1.4$, a maximum context length of 96K tokens, and the “thinking” mode enabled.
 230 To ensure fairness, we apply the same configuration to the base model: scaling temperature from
 231 0.0 to 1.4, enabling “thinking” mode, and setting $N = 128$ samples per question.

232 **Scaling K cannot make the base model comparable to RL.** As shown in Figure 3d, the RL-
 233 trained model consistently outperforms the base model when scaling the number of samples, with a
 234 clear advantage under small budgets (e.g., Pass@1). Although the performance gap narrows as K
 235 increases, the base model never fully catches up.

236 **Further scaling T can make the base model comparable to RL.** As shown in Figure 3d, scaling
 237 across temperatures raises the base model to Pass@All performance on par with the RL-trained
 238 model. The overall success rates are similar, though they differ in which questions remain unsolved:
 239 the base model fails on Q14, Q15, and Q28, while the RL model fails on Q13, Q14, and Q15. Thus,
 240 temperature scaling extends the base model’s reasoning boundary to match RL performance.

241 **Remarks.** There is an active debate on whether RL brings genuinely new capabilities to LLMs (Liu
 242 et al., 2025a). Some prior work argues that scaling K sufficiently can erase or even reverse the
 243 advantage of RL (Yue et al., 2025), while others contend that many of the base model’s large- K
 244 successes are merely guesses, and after verifier filtering the RL-trained model remains stronger (Wen
 245 et al., 2025). Our findings align with the latter: scaling K narrows the gap but does not eliminate it.
 246 However, when we further scale across T , the base model becomes comparable to RL. In this sense,
 247 exploring both K and T dimensions reveals a broader reasoning boundary than either alone.

249 3 DEEP ANALYSIS OF TEMPERATURE SCALING

250
 251 In this section, we first present the entropy dynamics underlying temperature scaling, then provide
 252 a case study to illustrate these behaviors in practice, and finally discuss the strengths and limitations
 253 of temperature scaling.

255 3.1 LEARNING THE ENTROPY DYNAMICS OF TEMPERATURE SCALING

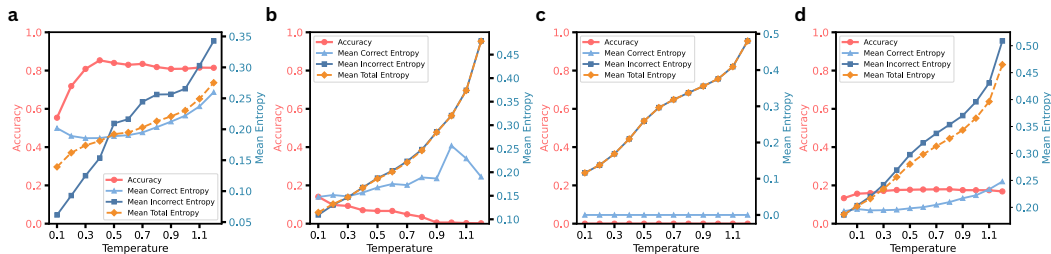
256
 257 **Entropy.** We measure model uncertainty using the entropy of the next-token distribution. At each
 258 step, given logits $f_\theta(x, y_{1:t-1})$, we compute

$$259 H = - \sum_y p(y | x, y_{1:t-1}) \log p(y | x, y_{1:t-1}), \quad p(y | x, y_{1:t-1}) = \text{softmax}(f_\theta(x, y_{1:t-1})).$$

260
 261 We report the average entropy H across all generated tokens. Low H indicates sharp logits and
 262 high confidence, while high H reflects uncertainty. Here, it is always computed from the model’s
 263 untempered softmax. As shown in Figure 4a–c, average H rises with T : higher T increases the
 264 chance of selecting low-probability tokens, raising the sequence-level H .

265
 266 **Entropy dynamics of correct vs. incorrect traces.** We analyze the average entropy of all traces, as
 267 well as the correct and incorrect subsets. For *easy/medium* questions (Figure 4a), correct-trace en-
 268 tropy increases smoothly with temperature, while incorrect-trace entropy rises much faster. At low
 269 T , correct traces show higher entropy, reflecting greater diversity; at high T , they show lower en-
 269 tropy, reflecting coherence when the model answers correctly. For *impossible* questions (Figure 4c),

270
271
272
273
274
275
276
277
278



279
280
281
282

Figure 4: Entropy dynamics of Qwen3-4B across temperatures on AIME 2025. (a) Q16. (b) Q29. (c) Q11. (d) Across the whole dataset.

283
284
285
286
287
288
289
290
291

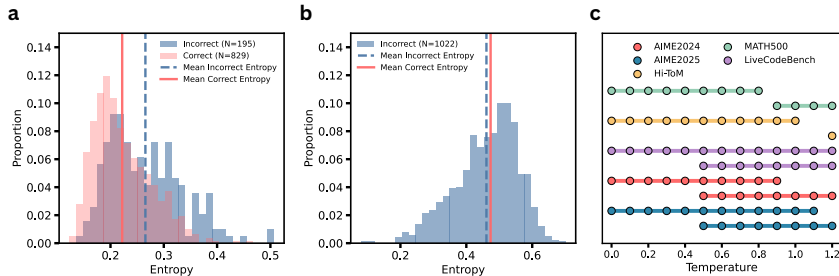


Figure 5: Entropy distributions and temperature subset. (a) An *easy* problem (AIME 2025 Q16, Qwen3-4B). (b) A *hard* problem (AIME 2025 Q25, Qwen3-4B). (c) Temperature minimal subset for Qwen3-4B.

292
293
294
295
296
297
298
299

in contrast, the entropy of all traces increases rapidly as T grows. At the dataset level (Figure 4d), correct traces consistently maintain lower entropy than incorrect ones at higher temperatures, suggesting that *the model seems to “know” when it is answering correctly*. But is this always the case?

300
301
302
303
304
305
306
307

When does the model “know it knows”? The entropy gap between correct and incorrect traces holds reliably only for *easy/medium* questions. As shown in Figure 5a, for an *easy* problem at $T = 1.0$, the entropy distribution of correct traces is clearly shifted to the left compared to incorrect ones. However, for a *hard* problem (Figure 5b), where only 2 out of 1,024 traces are correct, the entropy of correct traces is not always lower than that of incorrect ones. The model “knows it knows” only when the question is *easy*, but not when it is *hard*. This explains why in Figure 4d, where *easy* problems dominate, the aggregate effect makes it appear that correct traces always have lower entropy, drowning out the behavior on *hard* problems.

308
309
310
311
312

Remarks. Recent work has explored uncertainty-guided decoding (Kadavath et al., 2022; Kang et al., 2025), where high-entropy traces are discarded under the assumption they are unlikely correct (Fu et al., 2025). Our findings suggest this assumption holds for *easy/medium* questions, where correct traces typically align with lower entropy. However, for *hard* questions, correct traces do not necessarily exhibit low entropy, indicating uncertainty alone is insufficient as a universal signal.

313
314

3.2 CASE STUDY

315
316
317
318
319

Problem. We illustrate the effect of sampling temperature using AIME 2025 Q24. The task is to determine the number of zeros of the function $f(x) = \sin(7\pi \sin(5x)), 0 < x < 2\pi$. This problem essentially has a single correct reasoning path: one must reduce the condition $f(x) = 0$ to $\sin(5x) = k/7$ with $k \in \{-7, \dots, 7\}$, and count the corresponding solutions over five periods.

320
321
322
323

Case Study. At temperature $T = 0.7$, the model solved the problem in 2 out of 1,024 samples. The successful traces explicitly reduced the condition to $\sin(5x) = k/7$, enumerated the valid k values, and counted their solutions within $(0, 2\pi)$. In contrast, at $T = 0.9$, none of the sampled traces followed this reasoning structure; typical attempts instead guessed the number of zeros by listing approximate intersection points, without connecting them to the $k/7$ values, and thus failed to reach




Figure 6: Overview of the voting algorithm for efficient temperature scaling. Each temperature maintains its own candidate pool and first performs intra-temperature voting with threshold $\tau_{\text{intra}} = 0.8$. Only if all temperatures pass this stage do we proceed to cross-temperature voting, where majority votes are aggregated across temperatures with threshold $\tau_{\text{cross}} = 1.0$. Questions that meet both criteria are marked as *easy* and exit early.

the correct count. This illustrates that a *hard* question may be solvable under one temperature but not another, highlighting the need for temperature scaling.

3.3 STRENGTHS AND LIMITATIONS OF TEMPERATURE SCALING

Strengths. As shown in Figure 1d, the most notable strength of temperature scaling is that it achieves a higher upper bound when K is large, revealing a larger reasoning boundary. A second advantage is that distributing samples across temperatures does not substantially deviate from any single-temperature $\text{Pass}@K$ curve when K is small. This is because improvements at small budgets mainly come from solving *easy* questions, for which accuracy is largely unaffected by the choice of temperature (as also shown in Figures 2d and 3a,b).

Limitations. The main limitation of temperature scaling lies in computational efficiency. For example, using 12 temperatures requires roughly $12\times$ the compute of single-temperature scaling to reach the expanded reasoning boundary. This overhead is less of a concern when a strong verifier is available: most *easy* questions can be solved with little computation and verified early, allowing the budget to be focused on *hard* ones. However, such verifiers are not always accessible. In the next section, we investigate whether some of this cost can be saved even *without* verifiers.

4 DESIGN TEST-TIME METHODS WITH TEMPERATURE SCALING

In this section, we first analyze whether all temperatures and all questions are necessary for temperature scaling. We then design an algorithm for more efficient temperature sampling.

4.1 ARE ALL TEMPERATURES AND QUESTIONS NECESSARY FOR SCALING?

Finding a minimal subset of temperatures. A natural question is whether we must sample from all available temperatures to achieve the expanded upper bound, or whether a smaller subset is sufficient. As shown in Figure 5c, we evaluate subsets formed by gradually adding temperatures from either the low or high end. The results show that starting from higher temperatures requires a smaller subset to reach the upper bound; traces generated at low temperatures can also be obtained at higher ones. Based on this, we select a subset that generalizes across models and datasets, excluding very low temperatures (0.1–0.3).

Early exit for *easy* questions under temperature scaling. Figures 3a and 3b show that *easy* questions require far fewer samples to solve and can be answered reliably at any temperature with high probability. This observation motivates a simple strategy: avoid redundant sampling on such problems by using a voting-based mechanism across temperatures. In the following, we describe how multi-temperature voting can serve as an early-exit method.

4.2 ALGORITHM FOR EFFICIENT TEMPERATURE SCALING

Method overview. As illustrated in Figure 6, for each question we maintain per-temperature candidate answer pools. In each round, every temperature contributes one new trace per active question,

Table 2: Results across five datasets. Base reports Pass@1, 024 at $T = 0.6$. $+T$ reports results with full temperature scaling. $+Ours$ denotes our efficient temperature-scaling method. ΔC shows the reduction in computation cost relative to full temperature scaling.

Models	AIME2025				AIME2024				MATH500			
	Base	$+T$	$+Ours$	ΔC	Base	$+T$	$+Ours$	ΔC	Base	$+T$	$+Ours$	ΔC
Qwen3-8B	60.0	66.7	66.7	-31.6%	73.3	80.0	80.0	-31.6%	97.0	97.8	97.8	-54.4%
Qwen3-4B	60.0	73.3	73.3	-33.5%	66.7	76.7	76.7	-31.6%	95.5	98.5	98.5	-49.5%
Qwen3-1.7B	46.7	50.0	50.0	-27.2%	43.3	50.0	50.0	-27.2%	92.5	95.5	95.5	-36.3%
Qwen3-0.6B	20.0	36.7	36.7	-25.0%	23.3	30.0	30.0	-25.0%	88.1	94.0	94.0	-26.0%

Models	LiveCodeBench				Hi-ToM				Average			
	Base	$+T$	$+Ours$	ΔC	Base	$+T$	$+Ours$	ΔC	Base	$+T$	$+Ours$	ΔC
Qwen3-8B	32.6	40.0	40.0	-35.0%	93.0	95.5	94.5	-78.7%	71.2	76.0	75.8	-46.3%
Qwen3-4B	36.0	40.0	40.0	-32.8%	83.0	92.0	86.5	-32.8%	68.2	76.1	75.0	-36.0%
Qwen3-1.7B	29.7	35.4	35.4	-29.9%	37.0	55.0	42.0	-78.3%	49.8	57.2	54.6	-39.8%
Qwen3-0.6B	25.7	31.4	31.4	-26.1%	71.5	82.5	81.5	-32.4%	45.7	54.9	54.7	-26.9%

and the corresponding answers are recorded. We first perform *intra-temperature* voting: for each temperature, the most frequent answer is selected, and if its vote count does not reach the intra-threshold τ_{intra} , that temperature is deemed not yet confident and sampling continues. Only when all temperatures satisfy the intra-threshold do we proceed to *cross-temperature* voting, where the majority answers from each temperature are voted across temperatures. If the winning answer reaches the cross-threshold τ_{cross} , the question is marked as *easy* and exits early. Otherwise, sampling continues in the next round. The complete algorithm is provided in Appendix C.

Remarks. It is well known that the answer most favored by the model is not always correct (Brown et al., 2024). We acknowledge this limitation. However, our goal here is not to aggregate answers for final prediction, but to identify which questions are *easy* and can exit early. Moreover, the use of multi-temperature voting helps smooth out spurious signals, and the relatively high thresholds we adopt ($\tau_{\text{intra}} = 0.8$, $\tau_{\text{cross}} = 1.0$) provide additional robustness against noise.

5 EXPERIMENTS

In this section, we evaluate the proposed method for more efficient temperature sampling. We then analyze the results and discuss directions for future work.

5.1 RESULTS AND ANALYSIS

Main results. As shown in Table 2, our method reduces the computation required for temperature scaling across tasks, while maintaining nearly the same performance. The efficiency gains come from excluding very low temperatures and enabling early exit on *easy* questions. For example, on MATH500, Qwen3-8B achieves a 54.4% reduction in computation with negligible loss in Pass@All. Notably, Hi-ToM shows a different pattern: as a belief-allocation reasoning task, it can be solved by powerful models but also occasionally by weaker ones for spurious reasons, leading to less consistent scaling behavior.

A powerful model is what we need. As shown in Table 2, the overall computation savings also follow a scaling trend. A strong yet compact model, such as Qwen3-8B, is particularly well-suited for temperature scaling: it classifies more questions as *easy*, enabling our method to evict them early for efficiency and apply temperature scaling only to the *hard* ones. By contrast, models that are too small struggle even on *easy* questions, leaving little room for efficient scaling.

5.2 DISCUSSION

Alternative strategies for evicting *easy* questions. Beyond multi-temperature voting, we also explored using entropy signals to identify questions that can be exited early. If the entropy remains very low and changes little across temperatures, the question is likely *easy*; if the entropy is very high, the question is likely *impossible*, and repeated sampling will not help. However, we find that

432 entropy distributions vary significantly across datasets and domains, making it difficult to design a
 433 universal entropy-based eviction strategy.

434 **Temperature sampling in feedback-based workflows.** In this paper, we focus on internal model
 435 signals to evict *easy* questions. In real applications, this is less of a concern: one can simply call
 436 a powerful verifier after each run, and if a trace is deemed correct, the question can be evicted
 437 immediately. This would allow temperature sampling to focus only on unsolved questions, yielding
 438 substantial savings. However, in such workflows the verification cost itself becomes a critical factor.
 439 Regardless, we believe temperature sampling remains a simple yet powerful way to explore the
 440 reasoning boundary of base models. An interesting direction for future work is to study how variable
 441 temperature within a single trace may further influence reasoning ability.

444 6 RELATED WORK

445 **TTS for LLMs.** Multi-trace TTS (Brown et al., 2024; Snell et al., 2025; Schaeffer et al., 2025;
 446 Zhao et al., 2025; Liu et al., 2025b; Khairi et al., 2025) generates multiple candidate completions in
 447 parallel and selects the best one using either a verifier (Wang et al., 2024; Gou et al., 2024; Sun et al.,
 448 2024; 2025; Singh et al., 2025; Guo et al., 2025b; Dorner et al., 2025; Saad-Falcon et al., 2025a) or
 449 voting-based approaches (Wang et al., 2023; 2025c). This approach can be further combined with
 450 search algorithms (Yao et al., 2023; Bi et al., 2025; Chatziveroglou, 2025; Qiu et al., 2024; Zhang
 451 et al., 2024a; Lin et al., 2025), which interleave generation and selection in a step-by-step manner to
 452 progressively improve the final output.

453 Recent studies have also explored the connection between RL and TTS (Sareen et al., 2025; Zuo
 454 et al., 2025; Yue et al., 2025; Wen et al., 2025; Liu et al., 2025a). In parallel, other works investigate
 455 how to teach models to dynamically choose between single-trace and multi-trace reasoning strate-
 456 gies (Yang et al., 2025b; Biju et al., 2025; Jin et al., 2025; Wang et al., 2025f). Recent work also
 457 incorporates TTS into reasoning workflows (Saad-Falcon et al., 2025b; Chakraborty et al., 2025;
 458 Shen et al., 2025).

459 **Sampling Methods for TTS.** A typical sampling pipeline consists of three components: the prompt
 460 input to the model, the stochastic sampling procedure, and post-processing of the logits. Prior work
 461 (Wang et al., 2025b; Liu et al., 2025c) has shown that the design of the prompt can influence the
 462 inference performance. Based on temperature sampling (Ackley et al., 1985), many works have
 463 explored new strategies for sampling (Meister et al., 2023; Li et al., 2024a; Zhu et al., 2024b; Zhang
 464 et al., 2024b; Cai et al., 2024; Dhuliawala et al., 2024; Li et al., 2025b) and logit truncation (Fan
 465 et al., 2018; Holtzman et al., 2020; Basu et al., 2021; Hewitt et al., 2022; Minh et al., 2025) to better
 466 balance diversity and quality. Recently, Renze (2024) reported that temperature has limited impact
 467 on model reasoning capability. Our results suggest this is not the case when using Pass@ K as the
 468 metric. Meanwhile, Du et al. (2025) found that the optimal temperature varies across datasets, a
 469 pattern we also observe in Figure 2d.

470 More related work on single-trace TTS, efficient algorithms, and serving systems can be found in
 471 Appendix D.

474 7 CONCLUSION

475 We revisit TTS for reasoning in LLMs and show that scaling the number of samples K alone has di-
 476 minishing returns: once K is large, accuracy plateaus and some questions remain unsolved. In con-
 477 trast, scaling the sampling temperature expands the reasoning boundary, as different temperatures
 478 unlock different *hard* problems while *easy* ones are solved universally. This makes temperature scal-
 479 ing a simple yet powerful complement to traditional TTS, enabling base models to match RL-trained
 480 counterparts without additional training. Through entropy-based analysis and case studies, we fur-
 481 ther characterize the dynamics behind this effect. Finally, we propose efficient multi-temperature
 482 voting methods that cut the overhead of temperature scaling by exiting early on *easy* questions.
 483 Overall, our results highlight temperature scaling as an effective and practical tool to unlock the
 484 latent reasoning capabilities of LLMs.

REFERENCES

- 486
487
488 David H. Ackley, Geoffrey E. Hinton, and Terrence J. Sejnowski. A learning algorithm for boltz-
489 mann machines. *Cognitive Science*, 9(1):147–169, 1985.
- 490
491 Pranjala Aggarwal, Aman Madaan, Yiming Yang, et al. Let’s sample step by step: Adaptive-
492 consistency for efficient reasoning and coding with llms. In *Proceedings of the 2023 Conference*
493 *on Empirical Methods in Natural Language Processing*, pp. 12375–12396, 2023.
- 494
495 Chenxin An, Zhihui Xie, Xiaonan Li, Lei Li, Jun Zhang, Shansan Gong, Ming Zhong, Jingjing Xu,
496 Xipeng Qiu, Mingxuan Wang, and Lingpeng Kong. Polaris: A post-training recipe for scaling
497 reinforcement learning on advanced reasoning models, 2025.
- 498
499 Sourya Basu, Govardana Sachitanandam Ramachandran, Nitish Shirish Keskar, and Lav R Varsh-
500 ney. Mirostat: A neural text decoding algorithm that directly controls perplexity. In *The ninth*
501 *International Conference on Learning Representations*, 2021.
- 502
503 Zhenni Bi, Kai Han, Chuanjian Liu, Yehui Tang, and Yunhe Wang. Forest-of-thought: Scaling test-
504 time compute for enhancing llm reasoning. In *Forty-second International Conference on Machine*
505 *Learning*, 2025.
- 506
507 Emil Biju, Shayan Talaei, Zheming Huang, Mohammadreza Pourreza, Azalia Mirhoseini, and Amin
508 Saberi. Sprint: Enabling interleaved planning and parallelized execution in reasoning models,
509 2025.
- 510
511 Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and
512 Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling,
513 2024.
- 514
515 Chengkun Cai, Xu Zhao, Yucheng Du, Haoliang Liu, and Lei Li. t^2 of thoughts: Temperature tree
516 elicits reasoning in large language models, 2024.
- 517
518 Souradip Chakraborty, Mohammadreza Pourreza, Ruoxi Sun, Yiwen Song, Nino Scherrer, Furong
519 Huang, Amrit Singh Bedi, Ahmad Beirami, Jindong Gu, Hamid Palangi, and Tomas Pfister. On
520 the role of feedback in test-time scaling of agentic ai workflows, 2025.
- 521
522 Giannis Chatziveroglou. A*-decoding: Token-efficient inference scaling, 2025.
- 523
524 Peter Baile Chen, Yi Zhang, Dan Roth, Samuel Madden, Jacob Andreas, and Michael Cafarella.
525 Log-augmented generation: Scaling test-time reasoning with reusable computation, 2025.
- 526
527 Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and
528 Furu Wei. Reasoning with exploration: An entropy perspective, 2025.
- 529
530 Shehzaad Dhuliawala, Ilya Kulikov, Ping Yu, Asli Celikyilmaz, Jason Weston, Sainbayar Sukhbaatar,
531 and Jack Lanchantin. Adaptive decoding via latent preference optimization, 2024.
- 532
533 Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang,
534 Zengmao Wang, Ziwei Liu, Bo Du, Xianglong Liu, and Dacheng Tao. Dynamic parallel tree
535 search for efficient llm reasoning, 2025.
- 536
537 Florian E Dorner, Yatong Chen, André F Cruz, and Fanny Yang. Roc-n-reroll: How verifier imper-
538 fection affects test-time scaling, 2025.
- 539
540 Weihua Du, Yiming Yang, and Sean Welleck. Optimizing temperature for language models with
541 multi-sample inference. In *Forty-second International Conference on Machine Learning*, 2025.
- 542
543 Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In Iryna
544 Gurevych and Yusuke Miyao (eds.), *Proceedings of the 56th Annual Meeting of the Association*
545 *for Computational Linguistics (Volume 1: Long Papers)*, July 2018.
- 546
547 Yichao Fu, Xuefei Wang, Yuandong Tian, and Jiawei Zhao. Deep think with confidence, 2025.

- 540 Aryo Pradipta Gema, Alexander Hagele, Runjin Chen, Andy Arditi, Jacob Goldman-Wetzler, Kit
541 Fraser-Taliente, Henry Sleight, Linda Petrini, Julian Michael, Beatrice Alex, et al. Inverse scaling
542 in test-time compute, 2025.
- 543 In Gim, Guojun Chen, Seung-seob Lee, Nikhil Sarda, Anurag Khandelwal, and Lin Zhong. Prompt
544 cache: Modular attention reuse for low-latency inference. *Proceedings of Machine Learning and
545 Systems*, 6:325–338, 2024.
- 546 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Nan Duan, Weizhu Chen, et al. Critic: Large
547 language models can self-correct with tool-interactive critiquing. In *The Twelfth International
548 Conference on Learning Representations*, 2024.
- 549 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
550 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
551 via reinforcement learning, 2025a.
- 552 Jiaxin Guo, Zewen Chi, Li Dong, Qingxiu Dong, Xun Wu, Shaohan Huang, and Furu Wei. Reward
553 reasoning model, 2025b.
- 554 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn
555 Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset.
556 In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks
557 Track*, 2021.
- 558 John Hewitt, Christopher Manning, and Percy Liang. Truncation sampling as language model
559 desmoothing. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Findings of the Association
560 for Computational Linguistics: EMNLP 2022*, December 2022.
- 561 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text
562 degeneration. In *The Eighth International Conference on Learning Representations*, 2020.
- 563 Coleman Hooper, Sehoon Kim, Suhong Moon, Kerem Dilmen, Monishwaran Maheswaran,
564 Nicholas Lee, Michael W Mahoney, Sophia Shao, Kurt Keutzer, and Amir Gholami. Ets: Ef-
565 ficient tree search for inference-time scaling, 2025.
- 566 Chengsong Huang, Langlin Huang, Jixuan Leng, Jiacheng Liu, and Jiaxin Huang. Efficient test-time
567 scaling via self-calibration, 2025.
- 568 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
569 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
570 evaluation of large language models for code. In *The Thirteenth International Conference on
571 Learning Representations*, 2025.
- 572 Tian Jin, Ellie Y Cheng, Zachary Ankner, Nikunj Saunshi, Blake M Elias, Amir Yazdanbakhsh,
573 Jonathan Ragan-Kelley, Suvinay Subramanian, and Michael Carbin. Learning to keep a promise:
574 Scaling language model decoding parallelism with learned asynchronous decoding. In *Forty-
575 second International Conference on Machine Learning*, 2025.
- 576 Jordan Juravsky, Bradley Brown, Ryan Ehrlich, Daniel Y Fu, Christopher Ré, and Azalia Mirho-
577 seini. Hydragen: High-throughput llm inference with shared prefixes, 2024.
- 578 Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez,
579 Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language mod-
580 els (mostly) know what they know, 2022.
- 581 Zhewei Kang, Xuandong Zhao, and Dawn Song. Scalable best-of-n selection for large language
582 models via self-certainty, 2025.
- 583 Ammar Khairi, Daniel D’souza, Ye Shen, Julia Kreutzer, and Sara Hooker. When life gives you
584 samples: The benefits of scaling up inference compute for multilingual llms, 2025.
- 585 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph
586 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
587 serving with pagedattention. In *Proceedings of the 29th symposium on operating systems princi-
588 ples*, pp. 611–626, 2023.

- 594 Adrian Lancucki, Konrad Staniszewski, Piotr Nawrot, and Edoardo M Ponti. Inference-time hyper-
595 scaling with kv cache compression, 2025.
596
- 597 Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xiangxi Mo, Eric Tang, Sumanth Hegde, Kourosh
598 Hakhamaneshi, Shishir G Patil, Matei Zaharia, Joseph E Gonzalez, and Ion Stoica. Llms can
599 easily learn to reason from demonstrations structure, not content, is what matters!, 2025a.
- 600 Yiwei Li, Fei Mi, Yitong Li, Yasheng Wang, Bin Sun, Shaoxiong Feng, and Kan Li. Dynamic
601 stochastic decoding strategy for open-domain dialogue generation. In Lun-Wei Ku, Andre Mar-
602 tins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL*
603 *2024*, August 2024a.
- 604 Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Xinglin Wang, Bin Sun, Heda Wang, and
605 Kan Li. Escape sky-high cost: Early-stopping self-consistency for multi-step reasoning. In *The*
606 *Twelfth International Conference on Learning Representations*, 2024b.
- 607 Yiwei Li, Ji Zhang, Shaoxiong Feng, Peiwen Yuan, Xinglin Wang, Jiayi Shi, Yueqi Zhang, Chuyi
608 Tan, Boyuan Pan, Yao Hu, and Kan Li. Revisiting self-consistency from dynamic distributional
609 alignment perspective on answer aggregation. In *Findings of the Association for Computational*
610 *Linguistics: ACL 2025*, July 2025b.
- 611 Qingwen Lin, Boyan Xu, Zijian Li, Zhifeng Hao, Keli Zhang, and Ruichu Cai. Leveraging con-
612 strained monte carlo tree search to generate reliable long chain-of-thought for mathematical rea-
613 soning, 2025.
- 614 Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong.
615 Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models,
616 2025a.
- 617 Runze Liu, Junqi Gao, Jian Zhao, Kaiyan Zhang, Xiu Li, Biqing Qi, Wanli Ouyang, and Bowen
618 Zhou. Can 1b llm surpass 405b llm? rethinking compute-optimal test-time scaling. In *Workshop*
619 *on Reasoning and Planning for Large Language Models at the Thirteenth International Confer-*
620 *ence on Learning Representations*, 2025b.
- 621 Yexiang Liu, Zekun Li, Zhi Fang, Nan Xu, Ran He, and Tieniu Tan. Rethinking the role of prompting
622 strategies in llm test-time scaling: A perspective of probability theory, 2025c.
- 623 Clara Meister, Tiago Pimentel, Gian Wiher, and Ryan Cotterell. Locally typical sampling. *Transac-*
624 *tions of the Association for Computational Linguistics*, 11, 2023.
- 625 Xupeng Miao, Gabriele Oliaro, Zhihao Zhang, Xinhao Cheng, Zeyu Wang, Zhengxin Zhang, Rae
626 Ying Yee Wong, Alan Zhu, Lijie Yang, Xiaoxiang Shi, Chunan Shi, Zhuoming Chen, Daiyaan
627 Arfeen, Reyna Abhyankar, and Zhihao Jia. Specinfer: Accelerating large language model serving
628 with tree-based speculative inference and verification. In *Proceedings of the 29th ACM Interna-*
629 *tional Conference on Architectural Support for Programming Languages and Operating Systems,*
630 *Volume 3*, pp. 932–949, 2024.
- 631 Nguyen Nhat Minh, Andrew Baker, Clement Neo, Allen G Roush, Andreas Kirsch, and Ravid
632 Shwartz-Ziv. Turning up the heat: Min-p sampling for creative and coherent llm outputs. In *The*
633 *Thirteenth International Conference on Learning Representations*, 2025.
- 634 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
635 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
636 scaling, 2025.
- 637 Rui Pan, Yinwei Dai, Zhihao Zhang, Gabriele Oliaro, Zhihao Jia, and Ravi Netravali. Specreason:
638 Fast and accurate inference-time compute via speculative reasoning, 2025a.
- 639 Rui Pan, Zhuang Wang, Zhen Jia, Can Karakus, Luca Zancato, Tri Dao, Yida Wang, and Ravi
640 Netravali. Marconi: Prefix caching for the era of hybrid llms. In *Eighth Conference on Machine*
641 *Learning and Systems*, 2025b.

- 648 Zaifeng Pan, Yitong Ding, Yue Guan, Zheng Wang, Zhongkai Yu, Xulong Tang, Yida Wang, and
649 Yufei Ding. Fasttree: Optimizing attention kernel and runtime for tree-structured llm inference.
650 In *Eighth Conference on Machine Learning and Systems*, 2025c.
- 651
- 652 Ruoyu Qin, Zheming Li, Weiran He, Mingxing Zhang, Yongwei Wu, Weimin Zheng, and Xinran
653 Xu. Mooncake: A kvcache-centric disaggregated architecture for llm serving, 2025.
- 654
- 655 Jiahao Qiu, Yifu Lu, Yifan Zeng, Jiacheng Guo, Jiayi Geng, Huazheng Wang, Kaixuan Huang, Yue
656 Wu, and Mengdi Wang. Treebon: Enhancing inference-time alignment with speculative tree-
657 search and best-of-n sampling, 2024.
- 658 Matthew Renze. The effect of sampling temperature on problem solving in large language models.
659 In *Findings of the association for computational linguistics: EMNLP 2024*, pp. 7346–7356, 2024.
- 660
- 661 Jon Saad-Falcon, Buchanan E Kelly, Mayee F Chen, Tzu-Heng Huang, Brendan McLaughlin, Tanvir
662 Bhathal, Shang Zhu, Ben Athiwaratkun, Frederic Sala, Scott Linderman, Azalia Mirhoseini, and
663 Christopher Ré. Shrinking the generation-verification gap with weak verifiers, 2025a.
- 664
- 665 Jon Saad-Falcon, Adrian Gamarra Lafuente, Shlok Natarajan, Nahum Maru, Hristo Todorov,
666 Etash Kumar Guha, E Kelly Buchanan, Mayee F Chen, Neel Guha, Christopher Re, et al. An
667 architecture search framework for inference-time techniques. In *Forty-second International Con-
668 ference on Machine Learning*, 2025b.
- 669
- 670 Ranajoy Sadhukhan, Zhuoming Chen, Haizhong Zheng, Yang Zhou, Emma Strubell, and Beidi
671 Chen. Kinetics: Rethinking test-time scaling laws, 2025.
- 672
- 673 Kusha Sareen, Morgane M Moss, Alessandro Sordoni, Rishabh Agarwal, and Arian Hosseini.
674 Putting the value back in rl: Better test-time scaling by unifying llm reasoners with verifiers,
2025.
- 675
- 676 Rylan Schaeffer, Joshua Kazdan, John Hughes, Jordan Juravsky, Sara Price, Aengus Lynch, Erik
677 Jones, Robert Kirk, Azalia Mirhoseini, and Sanmi Koyejo. How do large language monkeys get
their power (laws)? In *Forty-second International Conference on Machine Learning*, 2025.
- 678
- 679 Melanie Sclar, Jane Dwivedi-Yu, Maryam Fazel-Zarandi, Yulia Tsvetkov, Yonatan Bisk, Yejin Choi,
680 and Asli Celikyilmaz. Explore theory of mind: program-guided adversarial data generation for
681 theory of mind reasoning. In *The Thirteenth International Conference on Learning Representa-
682 tions*, 2025.
- 683
- 684 Junhong Shen, Hao Bai, Lunjun Zhang, Yifei Zhou, Amrith Setlur, Shengbang Tong, Diego Caples,
685 Nan Jiang, Tong Zhang, Ameet Talwalkar, and Aviral Kumar. Thinking vs. doing: Agents that
reason by scaling test-time interaction, 2025.
- 686
- 687 Anikait Singh, Kushal Arora, Sedrick Keh, Jean Mercat, Tatsunori Hashimoto, Chelsea Finn, and
688 Aviral Kumar. Improving test-time search for llms with backtracking against in-context value
689 verifiers. In *Workshop on Reasoning and Planning for Large Language Models at the Thirteenth
690 International Conference on Learning Representations*, 2025.
- 691
- 692 Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute
693 optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth Inter-
694 national Conference on Learning Representations*, 2025.
- 695
- 696 Jiwon Song, Dongwon Jo, Yulhwa Kim, and Jae-Joon Kim. Reasoning path compression: Com-
pressing generation trajectories for efficient llm reasoning, 2025.
- 697
- 698 Lin Sun, Chuang Liu, Xiaofeng Ma, Tao Yang, Weijia Lu, and Ning Wu. Freeprm: Training process
699 reward models without ground truth process labels, 2025.
- 700
- 701 Zhiqing Sun, Longhui Yu, Yikang Shen, Weiyang Liu, Yiming Yang, Sean Welleck, and Chuang
Gan. Easy-to-hard generalization: Scalable alignment beyond human supervision. *Advances in
Neural Information Processing Systems*, 37:51118–51168, 2024.

- 702 Ghosal Soumya Suvra, Souradip Chakraborty, Avinash Reddy, Yifu Lu, Mengdi Wang, Dinesh
703 Manocha, Furong Huang, Mohammad Ghavamzadeh, and Amrit Singh Bedi. Does thinking more
704 always help? understanding test-time scaling in reasoning models, 2025.
705
- 706 Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
707 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Pro-
708 ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume
709 1: Long Papers)*, pp. 9426–9439, 2024.
- 710 Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,
711 Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive
712 effective reinforcement learning for llm reasoning, 2025a.
713
- 714 Tianchun Wang, Zichuan Liu, Yuanzhou Chen, Jonathan Light, Haifeng Chen, Xiang Zhang, and
715 Wei Cheng. Diversified sampling improves scaling llm inference, 2025b.
- 716 Weiqin Wang, Yile Wang, and Hui Huang. Ranked voting based self-consistency of large language
717 models, 2025c.
718
- 719 Xinglin Wang, Yiwei Li, Shaoxiong Feng, Peiwen Yuan, Yueqi Zhang, Jiayi Shi, Chuyi Tan, Boyuan
720 Pan, Yao Hu, and Kan Li. Every rollout counts: Optimal resource allocation for efficient test-time
721 scaling, 2025d.
- 722 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha
723 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language
724 models. In *The Eleventh International Conference on Learning Representations*, 2023.
725
- 726 Yiming Wang, Pei Zhang, Siyuan Huang, Baosong Yang, Zhuosheng Zhang, Fei Huang, and Rui
727 Wang. Sampling-efficient test-time scaling: Self-estimating the best-of-n sampling in early de-
728 coding, 2025e.
- 729 Yizhou Wang, Lingzhi Zhang, Yue Bai, Mang Tik Chiu, Zhengmian Hu, Mingyuan Zhang, Qihua
730 Dong, Yu Yin, Sohrab Amirghodsi, and Yun Fu. Cautious next token prediction. In Wanxiang
731 Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the
732 Association for Computational Linguistics: ACL 2025*, July 2025f.
733
- 734 Youkang Wang, Jian Wang, Rubing Chen, Xiao-Yong Wei, and Qing Li. Probabilistic optimality for
735 inference-time scaling, 2025g.
- 736 Zhenglin Wang, Jialong Wu, Yilong Lai, Congzhi Zhang, and Deyu Zhou. Seed: Accelerating
737 reasoning tree construction via scheduled speculative decoding. In *Proceedings of the 31st Inter-
738 national Conference on Computational Linguistics*, pp. 4920–4937, 2025h.
739
- 740 Zhibin Wang, Rui Ning, Chao Fang, Zhonghui Zhang, Xi Lin, Shaobo Ma, Mo Zhou, Xue Li,
741 Zhongfeng Wang, Chengying Huan, Rong Gu, Kun Yang, Guihai Chen, Sheng Zhong, and Chen
742 Tian. Flashforge: Ultra-efficient prefix-aware attention for llm decoding, 2025i.
- 743 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
744 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
745 neural information processing systems*, 35:24824–24837, 2022.
746
- 747 Xumeng Wen, Zihan Liu, Shun Zheng, Zhijian Xu, Shengyu Ye, Zhirong Wu, Xiao Liang, Yang
748 Wang, Junjie Li, Ziming Miao, Jiang Bian, and Mao Yang. Reinforcement learning with verifiable
749 rewards implicitly incentivizes correct reasoning in base llms, 2025.
- 750 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws:
751 An empirical analysis of compute-optimal inference for llm problem-solving. In *The Thirteenth
752 International Conference on Learning Representations*, 2025a.
753
- 754 Yufan Wu, Yinghui He, Yilin Jia, Rada Mihalcea, Yulong Chen, and Naihao Deng. Hi-ToM: A
755 benchmark for evaluating higher-order theory of mind reasoning in large language models. In
Findings of the Association for Computational Linguistics: EMNLP 2023, December 2023.

- 756 Yuheng Wu, Wentao Guo, Zirui Liu, Heng Ji, Zhaozhuo Xu, and Denghui Zhang. How large lan-
757 guage models encode theory-of-mind: a study on sparse parameter patterns. *npj Artificial Intelli-*
758 *gence*, 1(1):20, 2025b.
- 759 Yuheng Wu, Jianwen Xie, Denghui Zhang, and Zhaozhuo Xu. Del-tom: Inference-time scaling for
760 theory-of-mind reasoning via dynamic epistemic logic, 2025c.
- 761 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
762 Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report, 2025a.
- 763 Xinyu Yang, Yuwei An, Hongyi Liu, Tianqi Chen, and Beidi Chen. Multiverse: Your language
764 models secretly decide how to parallelize and merge generation, 2025b.
- 765 Zhihe Yang, Xufang Luo, Zilong Wang, Dongqi Han, Zhiyuan He, Dongsheng Li, and Yunjian Xu.
766 Do not let low-probability tokens over-dominate in rl for llms, 2025c.
- 767 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
768 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-*
769 *vances in neural information processing systems*, 36:11809–11822, 2023.
- 770 Lu Ye, Ze Tao, Yong Huang, and Yang Li. Chunkattention: Efficient self-attention with prefix-aware
771 kv cache and two-phase partition. In *Proceedings of the 62nd Annual Meeting of the Association*
772 *for Computational Linguistics (Volume 1: Long Papers)*, pp. 11608–11620, 2024.
- 773 Zihao Ye, Lequn Chen, Ruihang Lai, Wuwei Lin, Yineng Zhang, Stephanie Wang, Tianqi Chen,
774 Baris Kasikci, Vinod Grover, Arvind Krishnamurthy, et al. Flashinfer: Efficient and customiz-
775 able attention engine for llm inference serving. In *Eighth Conference on Machine Learning and*
776 *Systems*, 2025.
- 777 Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. Orca: A
778 distributed serving system for Transformer-Based generative models. In *16th USENIX Symposium*
779 *on Operating Systems Design and Implementation (OSDI 22)*, pp. 521–538, July 2022.
- 780 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao
781 Huang. Does reinforcement learning really incentivize reasoning capacity in llms beyond the
782 base model?, 2025.
- 783 Ted Zadouri, Hubert Strauss, and Tri Dao. Hardware-efficient attention for fast decoding, 2025.
- 784 Dan Zhang, Sining Zhoubian, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. Rest-mcts*:
785 Llm self-training via process reward guided tree search. In *Proceedings of the 38th International*
786 *Conference on Neural Information Processing Systems*, pp. 64735–64772, 2024a.
- 787 Kexun Zhang, Shang Zhou, Danqing Wang, William Yang Wang, and Lei Li. Scaling llm inference
788 efficiently with optimized sample compute allocation. In *Proceedings of the 2025 Conference of*
789 *the Nations of the Americas Chapter of the Association for Computational Linguistics: Human*
790 *Language Technologies (Volume 1: Long Papers)*, pp. 7959–7973, Albuquerque, New Mexico,
791 April 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.naacl-long.404.
- 792 Shimao Zhang, Yu Bao, and Shujian Huang. Edt: Improving large language models’ generation by
793 entropy-based dynamic temperature sampling, 2024b.
- 794 Eric Zhao, Pranjal Awasthi, and Sreenivas Gollapudi. Sample, scrutinize and scale: Effective
795 inference-time search by scaling verification. In *Forty-second International Conference on Ma-*
796 *chine Learning*, 2025.
- 797 Lianmin Zheng, Liangsheng Yin, Zhiqiang Xie, Chuyue Livia Sun, Jeff Huang, Cody Hao Yu, Shiyi
798 Cao, Christos Kozyrakis, Ion Stoica, Joseph E Gonzalez, et al. Sglang: Efficient execution of
799 structured language model programs. *Advances in neural information processing systems*, 37:
800 62557–62583, 2024.
- 801 Lei Zhu, Xinjiang Wang, Wayne Zhang, and Rynson Lau. Relayattention for efficient large language
802 model serving with long system prompts. In *Proceedings of the 62nd Annual Meeting of the*
803 *Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4945–4957, 2024a.

810 Yuqi Zhu, Jia Li, Ge Li, YunFei Zhao, Jia Li, Zhi Jin, and Hong Mei. Hot or cold? adaptive
 811 temperature sampling for code generation with large language models. In *Proceedings of the*
 812 *Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innova-*
 813 *tive Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances*
 814 *in Artificial Intelligence*, pp. 437–445, 2024b.

815 Yuxuan Zhu, Ali Falahati, David H Yang, and Mohammad Mohammadi Amiri. Sentencekv: Effi-
 816 cient llm inference via sentence-level semantic kv caching, 2025.

817
 818 Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen
 819 Zhang, Xinwei Long, Ermo Hua, Biqing Qi, Youbang Sun, Zhiyuan Ma, Lifan Yuan, Ning Ding,
 820 and Bowen Zhou. Ttrl: Test-time reinforcement learning, 2025.

821 822 823 APPENDIX

824 825 A DATASETS, EVALUATION, AND VERIFICATION

826 827 A.1 DATASETS AND EVALUATION METHODS

828
 829 **Mathematical reasoning.** We use AIME 2024, AIME 2025, and MATH500 (Hendrycks et al.,
 830 2021) to evaluate mathematical reasoning. AIME is a high school mathematics competition that
 831 features 30 challenging problems each year, and we include all 30 problems from both the 2024 and
 832 2025 editions. MATH500 is a subset of the MATH benchmark comprising challenging problems
 833 across multiple topics; we select all problems of difficulty level 5, which is the highest level in
 834 MATH500, yielding 134 questions. For evaluation, we prompt the model to place its final answer
 835 inside a `\boxed{ }` expression. For AIME, where answers are integers, the boxed content can be
 836 directly parsed and compared to the reference. For MATH500, where boxed expressions can be more
 837 complex, we use `sympy` to check mathematical equivalence between the predicted and reference
 838 answers.

839 **Code generation.** We use LiveCodeBench v6 (Jain et al., 2025), which consists of recently collected
 840 programming problems from AtCoder and LeetCode. Version 6 covers 175 problems released be-
 841 tween January 2025 and May 2025. To evaluate model outputs, we run each generated program
 842 against a large set of private test cases, and a solution is considered correct only if it passes all test
 843 cases.

844 **Social and logical reasoning.** Theory-of-mind (ToM) refers to the ability to infer others’ mental
 845 states such as beliefs (Sclar et al., 2025; Wu et al., 2025b). A common ToM evaluation format
 846 is the unexpected transfer task, which can be viewed as a form of commonsense dynamic logical
 847 reasoning (Wu et al., 2025c). For example, in a multi-step story: Alice and Bob are in a room with a
 848 chocolate in a box; after Alice leaves, Bob moves the chocolate to the table. A correct solver should
 849 infer that Alice still believes the chocolate is in the box. Hi-ToM (Wu et al., 2023) contains 200
 850 problems that evaluate multi-agent ToM reasoning over long, temporally structured scenarios. To
 851 evaluate model outputs, we generate process-level labels, and a prediction is considered correct only
 852 if all beliefs are correctly inferred at every step.

853 854 A.2 EVALUATION PROMPTS

855 The prompts for AIME, MATH500, and LiveCodeBench are shown in Figure 7, Figure 8, and
 856 Figure 9, respectively. For Hi-ToM, we use the same one-shot prompt format as in Wu et al. (2025c).

857 858 A.3 VERIFICATION PIPELINE

859
 860 For AIME, we apply an additional reasoning-trace verification step to reduce the risk of spurious
 861 correctness. Specifically, for every model and every AIME problem, if among 1,024 sampled traces
 862 the number of correct answers is ≤ 32 (i.e., $\text{Avg}@1024 \leq 3\%$), we perform gpt-5-assisted valida-
 863 tion. Gpt-5 is provided with multiple human-written reference solutions and asked to judge whether
 each model-generated reasoning trace is logically valid and leads to the correct answer. Only traces

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

```
Prompt for AIME

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

{Question}
```

Figure 7: Prompt used for AIME.

```
Prompt for MATH500

Answer the following math question step by step, given in LaTeX
format, clearly and concisely, and present the final answer as
\boxed{x}, where X is the fully simplified solution.

Example:
Question: \int_0^1 (3x^2 + 2x) \, dx
Solution: \int (3x^2 + 2x) \, dx = x^3 + x^2 + C
Evaluating from 0 to 1: (1^3 + 1^2) - (0^3 + 0^2) = 1 + 1 - 0 = 2

\boxed{2}

Now, solve the following question step by step.

{Question}
```

Figure 8: Prompt used for MATH500.

judged as correct are retained; all others are filtered out. The validation prompt is shown in Figure 10.

B LLM USAGE

We used GPT-5 as a general-purpose assist tool for language refinement and proofreading. In addition, GPT-5 was employed to assist in the verification of AIME reasoning traces, as described in Appendix A.

C ALGORITHM

The algorithm for the efficient temperature scaling can be found in Algorithm 1.

D EXTENDED RELATED WORK

Single-trace TTS. In single-trace TTS, the goal is to encourage deeper and more deliberate reasoning within a single inference path (Wei et al., 2022; Muennighoff et al., 2025). This can be achieved by RL (Guo et al., 2025a; Yang et al., 2025c; Wang et al., 2025a; Cheng et al., 2025; An et al., 2025) or by distilling reasoning traces from a stronger teacher model (Li et al., 2025a). Compared to single-trace approaches, multi-trace TTS has been shown to produce more stable results (Suvra et al., 2025; Gema et al., 2025). This work focuses on the multi-trace setting.

Efficient Algorithms for TTS. One line of work focuses on resource allocation, aiming to allocate more computation to difficult examples and less to easier ones (Aggarwal et al., 2023; Li et al., 2024b; Huang et al., 2025; Wu et al., 2025a; Wang et al., 2025d;e;g; Zhang et al., 2025; Fu et al.,

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

```
Prompt for LiveCodeBench

You are an expert Python programmer.
You will be given a programming problem and must generate a correct
Python solution that matches the specification and passes all
tests.

{Question}

Format:
You will use the following starter code to write the solution
and enclose your code within backticks.

```python
class Solution:
 def solve(self, ...):
 pass

Answer:
```

Figure 9: Prompt used for LiveCodeBench.

2025). Another line of work incorporates system-level techniques into TTS, such as compressing or reusing the KV cache to avoid redundant computation (Chen et al., 2025; Song et al., 2025; Ding et al., 2025; Hooper et al., 2025), or integrating with speculative decoding to reduce latency (Wang et al., 2025h; Pan et al., 2025a). These methods improve test-time efficiency by combining better search with system-level optimizations.

**Efficient Serving System for TTS.** Recent work has proposed various serving systems to reduce decoding latency. These include optimized attention kernels (Ye et al., 2025; Zadouri et al., 2025), speculative decoding methods (Miao et al., 2024), and KV cache compression techniques (Lancucki et al., 2025). Other efforts improve system throughput via techniques like continuous batching (Yu et al., 2022) and separating the prefilling and decoding stages (Qin et al., 2025).

Parallel decoding-based TTS can be viewed as a tree-structured decoding problem, where a shared prefix is expanded into multiple completions. In this setting, many queries access the same KV cache from the shared prefix, making efficient KV cache management a key challenge (Gim et al., 2024; Pan et al., 2025b; Zhu et al., 2025; Sadhukhan et al., 2025). vLLM addresses this by introducing PagedAttention (Kwon et al., 2023), which reduces memory usage and improves throughput. Building on this, SGLang proposes RadixAttention (Zheng et al., 2024), allowing programmatic control over KV cache reuse. In addition, to address the I/O overhead of reading and writing the KV cache, recent works (Juravsky et al., 2024; Ye et al., 2024; Zhu et al., 2024a; Pan et al., 2025c; Wang et al., 2025i) develop new methods to more efficiently compute both the shared prefix and its multiple completions.

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

```

Validation Prompt for AIME

You are a rigorous grading instructor for math contest solutions.
You are given the original problem text, the student's full written
solution (whose final numeric answer is known to be correct),
and one or more reference solutions (provided only as guidance;
exact verbatim matching is NOT required).

Your task is to judge whether the student's *reasoning process* is
logically valid and self-contained, instead of merely checking
the final answer.

[Original problem]
{ProblemText}

[Student's full solution]
{ModelTrace}

[Reference solution(s) - for guidance only]
Solution 1: ...
Solution 2: ...
Solution 3: ...

Judging criteria:
1) Are key derivations justified with sufficient intermediate steps,
 without circular reasoning or illicit assumptions?
2) Minor arithmetic/notation slips that are explicitly corrected
 later and do not affect the logic may still be acceptable.
3) If there is a fundamental flaw (misused theorem, false equality,
 missing essential conditions) such that the correct final answer
 could be a coincidence, the solution should be judged incorrect
 .
4) Different approaches from the reference are permitted as long as
 the argument forms a logically sound and complete proof.
Please reason step by step before you decide.

Output requirement:
After your analysis, the VERY LAST LINE must be exactly one of the
following:
[CORRECT]
or
[INCORRECT]
Do not output anything after that final line.

```

Figure 10: Validation prompt used for gpt-5-assisted verification on AIME problems.

1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

---

**Algorithm 1** Efficient Temperature Scaling with Two-Stage Voting
 

---

**Require:** Questions  $\mathcal{Q}$ ; temperatures  $\mathcal{T} = \{T_1, \dots, T_M\}$ ; max rounds  $R$ ; intra-temp threshold  $\tau_{\text{intra}}$ ; cross-temp threshold  $\tau_{\text{cross}}$ ; sampler  $\text{SAMPLE}(q, T)$ ; extractor  $\text{ANS}(y)$

**Ensure:** For each  $q \in \mathcal{Q}$ , per-temperature pools  $\{\mathcal{S}_{q,T}\}$  and their answers

- 1: **for all**  $q \in \mathcal{Q}$  **do**
- 2:   **for all**  $T \in \mathcal{T}$  **do**
- 3:      $\mathcal{S}_{q,T} \leftarrow \emptyset$  ▷ trace pool at temperature  $T$
- 4:      $\mathcal{A}_{q,T} \leftarrow \emptyset$  ▷ multiset of answers for  $T$
- 5:   **end for**
- 6:    $\text{ACTIVE}(q) \leftarrow \text{true}$
- 7: **end for**
- 8: **for**  $r \leftarrow 1$  **to**  $R$  **do**
- 9:   **for all**  $q \in \mathcal{Q}$  **with**  $\text{ACTIVE}(q) = \text{true}$  **do** ▷ Round- $r$  sampling
- 10:    **for all**  $T \in \mathcal{T}$  **do**
- 11:      $y \leftarrow \text{SAMPLE}(q, T)$
- 12:      $a \leftarrow \text{ANS}(y)$
- 13:     Append  $(T, y, a)$  to  $\mathcal{S}_{q,T}$ ; insert  $a$  into  $\mathcal{A}_{q,T}$
- 14:    **end for**
- 15:     $\text{all\_passed} \leftarrow \text{true}; \mathcal{V}_q \leftarrow \emptyset$  ▷ Stage 1: intra-temperature voting
- 16:    **for all**  $T \in \mathcal{T}$  **do** ▷  $\mathcal{V}_q$ : one vote per  $T$
- 17:     Build  $h_{q,T}(a) \leftarrow \#\{a' \in \mathcal{A}_{q,T} : a' = a\}$
- 18:      $v_{\max}^{(T)}(q) \leftarrow \max_a h_{q,T}(a)$
- 19:     **if**  $v_{\max}^{(T)}(q) < \tau_{\text{intra}}$  **then**
- 20:        $\text{all\_passed} \leftarrow \text{false}$  ▷ this  $T$  not confident yet
- 21:     **else**
- 22:        $a_T \leftarrow \arg \max_a h_{q,T}(a)$  ▷ temperature- $T$  majority
- 23:       Append  $a_T$  to  $\mathcal{V}_q$  ▷ one vote from this  $T$
- 24:     **end if**
- 25:    **end for**
- 26:    **if**  $\text{all\_passed} = \text{false}$  **then**
- 27:      **continue** ▷ skip cross-temp vote; keep sampling next round
- 28:    **end if**
- 29:    Build cross-temp tally  $H_q(b) \leftarrow \#\{a_T \in \mathcal{V}_q : a_T = b\}$  ▷ Stage 2: cross-temperature voting
- 30:     $V_{\max}(q) \leftarrow \max_b H_q(b)$
- 31:    **if**  $V_{\max}(q) \geq \tau_{\text{cross}}$  **then**
- 32:       $\text{ACTIVE}(q) \leftarrow \text{false}$  ▷ early exit for  $q$
- 33:    **end if**
- 34: **end for**
- 35: **if**  $\text{ALLINACTIVE}(\mathcal{Q})$  **then**
- 36:    **break**
- 37: **end if**
- 38: **end for**
- 39: **return**  $\{\mathcal{S}_{q,T} \mid q \in \mathcal{Q}, T \in \mathcal{T}\}$  ▷ downstream verifier/BoN picks final answers

---