ATI-OPD: ADAPTIVE TEACHER INTERVENTION FOR STABLE LONG-SEQUENCE ON-POLICY DISTILLATION

Anonymous authors
Paper under double-blind review

ABSTRACT

On-policy distillation (OPD) trains a student on its own responses using token-level feedback from a teacher. As responses grow longer, the student's prefix can depart from the teacher's own reasoning, weakening its feedback on later tokens. Updates based on this feedback can reinforce the same deviations in later responses and destabilize training. Prior approaches such as length control and loss reweighting limit learning from problematic responses but leave their prefixes unchanged. Fixed cutoffs can also discard later reasoning. We propose **Adaptive Teacher Intervention for On-Policy Distillation (ATI-OPD)** to improve subsequent teacher guidance by revising the student's reasoning prefix. ATI-OPD replaces the earliest high-disagreement segment with a short teacher segment, then lets the student resume its own reasoning from the revised text. To handle renewed disagreement, it excludes persistently divergent suffixes from the distillation loss while retaining useful student continuation. Training diagnostics show less teacher-student disagreement, fewer gradient spikes, and fewer extremely long responses. Across five mathematical reasoning benchmarks and two teacher-student pairs, ATI-OPD improves average student-only accuracy over standard OPD.

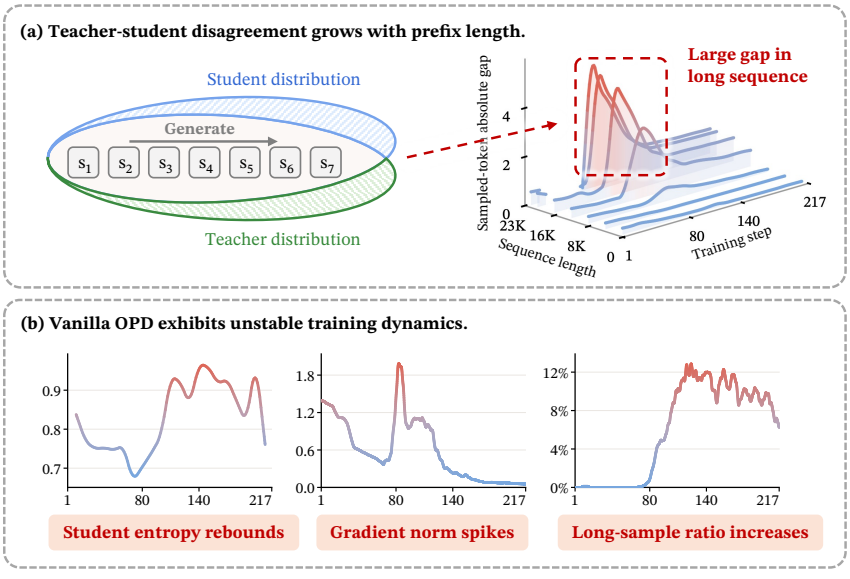


Figure 1: Vanilla OPD on JustRL-DeepSeek exhibits prediction disagreement and training instability. (a) Gap measures the absolute difference between teacher and student log-probabilities for the same sampled token and prefix. (b) Across training steps, entropy rebounds, gradient norms spike, and more responses reach the max token limit.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

On-policy distillation (OPD) trains a student on responses it generates itself, using a teacher's token-level predictions as supervision (Agarwal et al., 2024; Gu et al., 2024). Learning from its own prefixes exposes the student to the decisions it must make at inference. The teacher must therefore provide useful guidance from the prefixes that the student actually generates.

Generating long sequences can weaken this guidance and make training unstable. As the prefix grows, it can depart from the reasoning that the teacher would generate on its own. Even a strong teacher can become less effective when continuing student-written prefixes (Zhu et al., 2026). OPD uses these token-level predictions to update the student, so their quality affects which continuations it learns to generate. Existing works have found that repetitive continuations can receive favorable local updates. As the student samples them more often, they contribute more heavily to subsequent updates and reinforce length inflation (Luo et al., 2026). In long-sequence OPD, declining teacher-student agreement has also been observed alongside greater prediction uncertainty and gradient spikes (Li et al., 2026). We measure disagreement with the *sampled-token absolute gap*: the absolute difference between teacher and student log-probabilities for a sampled token under the same prefix. In vanilla OPD, large gaps appear at later response positions. Over training steps, entropy rebounds, gradient norms spike, and more responses reach the token limit, as shown in Figure 1.

To reduce the impact of unreliable feedback, prior methods control response length or reweight token losses (Ziheng et al., 2026; Zhang et al., 2026b; Xie et al., 2026). These approaches limit learning from problematic continuations but leave the generated prefix unchanged. Fixed cutoffs can also discard useful later reasoning. Teacher participation can instead change the course of generation through token replacement or local takeover (Xu et al., 2025; Koo et al., 2025; Xu et al., 2026). However, improving the prefix does not ensure that the student's later decisions remain useful for learning: disagreement can return after generation is handed back to the student. This calls for local guidance that changes the reasoning trajectory, together with supervision selection based on the regenerated continuation.

To address these challenges, we propose **Adaptive Teacher Intervention for On-Policy Distillation (ATI-OPD)** to revise the student's reasoning prefix and select supervision within the resulting continuation. Its core module, **Adaptive Teacher Intervention**, locates the earliest high-disagreement segment in the student response. The teacher takes over from the preceding prefix and replaces that segment, allowing the student to regenerate what follows. Targeting the earliest high-disagreement segment changes the prefix before that segment influences later decisions. A short teacher segment guides further reasoning while keeping the student's own decisions in the training trajectory.

The student can diverge again as it continues beyond the local teacher rewrite. **Adaptive Supervision Selection** therefore monitors successive windows of the new continuation and excludes persistently divergent suffixes from the distillation loss. This retains useful student supervision beyond the teacher segment while limiting updates on later disagreement.

Our contributions are as follows:

1. We introduce Adaptive Teacher Intervention to address weakened guidance on student-generated prefixes. Replacing a high-disagreement segment with a short teacher segment improves alignment in the student's continuation.
2. We introduce Adaptive Supervision Selection to handle renewed disagreement after intervention. Selective supervision outperforms both learning from the full suffix, even with more interventions, and stopping supervision at the teacher segment.
3. ATI-OPD reduces the area under the absolute-gap curve (AUC) by 47.8% and suppresses gradient spikes and extremely long responses during training. More stable training is accompanied by higher average student-only accuracy across five benchmarks on both model pairs.

2 RELATED WORK

OPD objectives and feedback stability. OPD methods improve feedback by changing divergence objectives (Agarwal et al., 2024; Ko et al., 2024), forming a weighted product of teacher and student distributions as the supervision target (Jang et al., 2026), or adding supervision at high-entropy teacher

positions (Jin et al., 2026). Stable-OPD addresses a feedback loop between repetitive continuations and favorable local updates through off-policy supervision and reference-policy regularization (Luo et al., 2026). ATI-OPD changes trajectories and supervision ranges while retaining the OPD update.

Teacher participation in trajectory generation. Teacher participation ranges from mixing sampling distributions (Gu et al., 2024) and selecting or replacing individual tokens (Xu et al., 2025; Koo et al., 2025) to taking over the remaining suffix (Peng et al., 2025). Teacher rewriting also appears in MOTAB, which joins corrective teacher suffixes to erroneous student prefixes (Wang et al., 2026), and TRD, which asks the teacher to revise completed responses (Jiang et al., 2026). Relay-OPD triggers takeovers when the teacher prefers a reflection token absent from the student's leading predictions and ends the rollout after its final teacher segment (Xu et al., 2026). ATI-OPD instead triggers on segment-averaged divergence and lets the student complete the response after a bounded replacement. Its supervision cutoff depends on persistent divergence in that continuation, rather than on the number of teacher takeovers.

Length control and selective supervision. Fixed-prefix and progressive-length methods restrict rollout lengths without using disagreement in the current continuation (Zhang et al., 2026a; Ziheng et al., 2026; Zhang et al., 2026b). Prune-OPD reduces rewards after low-overlap events and adjusts the next training step's rollout budget using batch-level reliable lengths (Yang et al., 2026), whereas KAT stops and rolls back rollouts after persistent low-KL agreement, calibrating its threshold from recent rollout minima (Xin et al., 2026). IW-OPD reweights compatible prefixes while retaining dense supervision (Xie et al., 2026). AC-OPD retains the original student trajectory and adds teacher-generated branches, selecting each branch's supervision length using uncertainty and accumulated divergence (Zhou et al., 2026). ATI-OPD instead replaces a student segment and selects which tokens in the regenerated student suffix receive OPD supervision, with its threshold calibrated from correct original responses. This selection takes place after full generation and changes only the loss mask.

3 Preliminaries

3.1 On-policy distillation

Let π_θ and π_T denote the student and fixed teacher. For question x, OPD samples $y = (y_1, \ldots, y_N)$ from a student snapshot $\pi_{\theta_{\text{old}}}$ held fixed during batch generation. With $h_t = (x, y_{<t})$ denoting the question and preceding text, the reverse-KL objective is

$$\mathcal{L}_{\text{OPD}}(\theta) = \mathbb{E}_{x,y}\left[\frac{1}{N}\sum_{t=1}^{N} D_{\text{KL}}(\pi_\theta(\cdot \mid h_t) \| \pi_T(\cdot \mid h_t))\right]. \tag{1}$$

Here, N is the response length, and the expectation samples questions from the training distribution and responses from the rollout student. In practice, we query exact teacher log-probabilities for the rollout student's top-16 candidates and use a clipped policy surrogate (Li et al., 2026). See Appendix C.1 for the implementation details.

3.2 Sampled-token divergence

We monitor teacher-student disagreement using a sampled-token estimator (Schulman, 2020):

$$\delta_t = \log \pi_T(y_t \mid h_t) - \log \pi_{\theta_{\text{old}}}(y_t \mid h_t), \qquad g_t = \exp(\delta_t) - \delta_t - 1. \tag{2}$$

Here, δ_t is the log-probability difference and g_t is the nonnegative sampled-token estimator, commonly called K3. With full support and no clipping, $\mathbb{E}_{y_t \sim \pi_{\theta_{\text{old}}}}[g_t] = D_{\text{KL}}(\pi_{\theta_{\text{old}}} \| \pi_T)$. We clip extreme values for numerical stability. Monitoring and local K3 diagnostics use the resulting score $\bar{g}_t$. See Appendix C.6 for details.

4 Adaptive Teacher Intervention for On-Policy Distillation

ATI-OPD links where teacher guidance changes student generation with how far the new continuation contributes to learning. It first locates an intervention segment, rewrites it with the teacher, and then selects subsequent supervision positions, as shown in Figure 2.

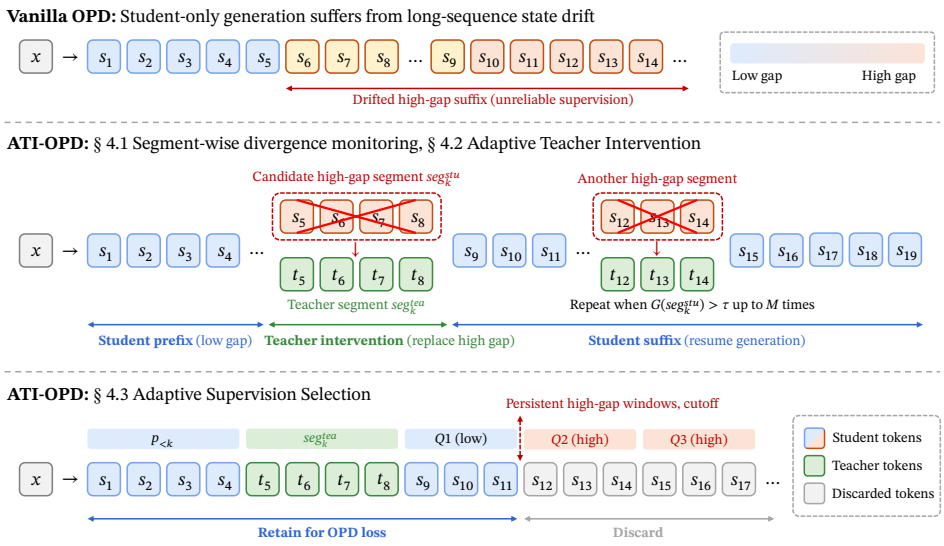


Figure 2: ATI-OPD replaces a high-divergence student segment with a teacher segment and lets the student generate a new continuation. Persistent divergence after the handoff determines a supervision cutoff. Later tokens remain in the generated response but are excluded from distillation.

4.1 Segment-wise Divergence Monitoring

We first generate a complete student response $y^{(0)}$ and score it with the teacher. To locate where their predictions diverge, we divide the response into segments of at most L_S tokens and compute

$$G(\text{seg}_k^{\text{stu}}) = \frac{\sum_{t \in \text{seg}_k^{\text{stu}}} m_t \bar{g}_t}{\sum_{t \in \text{seg}_k^{\text{stu}}} m_t}. \tag{3}$$

Here, $\text{seg}_k^{\text{stu}}$ is student segment k, and m_t marks valid student tokens, excluding padding and teacher tokens. We scan in response order and intervene at the first segment whose mean exceeds threshold τ, if budget remains. The main setting uses $\tau = 0.15$. Appendix B.4 examines its sensitivity.

4.2 Adaptive Teacher Intervention

To give the student a revised starting point, we retain the prefix $p_{<k}$ before the triggering segment, discard that segment and the original suffix, and sample a bounded teacher replacement:

$$\text{seg}_k^{\text{tea}} \sim \pi_T(\cdot \mid x, p_{<k}), \qquad |\text{seg}_k^{\text{tea}}| \le L_T. \tag{4}$$

Here, L_T limits the teacher segment length, and $|\cdot|$ denotes sequence length. Unless the response has ended, the student continues from $(x, p_{<k}, \text{seg}_k^{\text{tea}})$ until EOS or the total response budget $N_{\max}$. The retained prefix, teacher segment, and complete new suffix form a mixed-policy trajectory $\widetilde{y}$.

4.3 Adaptive Supervision Selection

A local intervention does not ensure useful guidance throughout the remaining response. We therefore score the complete new student continuation in windows of length L_F and select its supervision range using

$$Q_j = \frac{1}{|W_j|} \sum_{t \in W_j} \bar{g}_t, \qquad \eta_s = \text{Quantile}_{0.90}(\mathcal{H}_s). \tag{5}$$

Here, W_j contains the valid student tokens in window j, and $|W_j|$ is their count. At training step s, the reference pool $\mathcal{H}_s$ contains window means from earlier original responses with correct final answers. These windows provide an empirical scale of disagreement during successful student generation, allowing threshold η_s to follow changes in the student. We use $\eta_s = \tau$ until enough reference windows are available. See Appendix A for details.

Table 1: Student-only accuracy (%) on five mathematical reasoning benchmarks. ESR uses a fixed-prefix supervision implementation (Ziheng et al., 2026). Subscripts report inference standard errors in 16 independent responses per question. Bold / underline mark the best / second-best results per pair.

Method	AIME24	AIME25	AMC23	AMO	HMMT	Average
JustRL-DeepSeek-1.5B → DeepSeek-R1-Distill-Qwen-1.5B						
Student	$29.58_{\pm 2.08}$	$21.88_{\pm 1.89}$	$72.34_{\pm 1.77}$	$1.44_{\pm 0.48}$	$16.10_{\pm 1.60}$	$28.27_{\pm 0.74}$
+GRPO (Shao et al., 2024)	$31.04_{\pm 2.11}$	$24.58_{\pm 1.97}$	$71.88_{\pm 1.78}$	$1.44_{\pm 0.48}$	$13.83_{\pm 1.50}$	$28.55_{\pm 0.75}$
+OPD (Li et al., 2026)	$\underline{46.25}_{\pm 2.28}$	$\mathbf{33.96}_{\pm 2.16}$	$\underline{87.03}_{\pm 1.33}$	$\underline{2.08}_{\pm 0.57}$	$\underline{22.92}_{\pm 1.83}$	$\underline{38.45}_{\pm 0.78}$
+ESR (Ziheng et al., 2026)	$37.92_{\pm 2.21}$	$28.75_{\pm 2.07}$	$78.44_{\pm 1.63}$	$1.76_{\pm 0.53}$	$16.86_{\pm 1.63}$	$32.74_{\pm 0.77}$
+ATI-OPD	$\mathbf{49.17}_{\pm 2.28}$	$\underline{33.54}_{\pm 2.15}$	$\mathbf{88.59}_{\pm 1.26}$	$\mathbf{2.72}_{\pm 0.65}$	$\mathbf{24.24}_{\pm 1.87}$	$\mathbf{39.65}_{\pm 0.78}$
JustRL-Nemotron-1.5B → OpenMath-Nemotron-1.5B						
Student	$59.17_{\pm 2.24}$	$48.75_{\pm 2.28}$	$90.78_{\pm 1.14}$	$6.09_{\pm 0.96}$	$30.49_{\pm 2.00}$	$47.06_{\pm 0.81}$
+GRPO (Shao et al., 2024)	$59.58_{\pm 2.24}$	$48.96_{\pm 2.28}$	$90.62_{\pm 1.15}$	$5.29_{\pm 0.90}$	$30.87_{\pm 2.01}$	$47.07_{\pm 0.81}$
+OPD (Li et al., 2026)	$\underline{67.08}_{\pm 2.14}$	$\underline{59.38}_{\pm 2.24}$	$\underline{95.00}_{\pm 0.86}$	$\underline{9.78}_{\pm 1.19}$	$32.39_{\pm 2.04}$	$\underline{52.72}_{\pm 0.80}$
+ESR (Ziheng et al., 2026)	$63.96_{\pm 2.19}$	$54.79_{\pm 2.27}$	$93.91_{\pm 0.95}$	$6.57_{\pm 0.99}$	$\underline{32.77}_{\pm 2.04}$	$50.40_{\pm 0.80}$
+ATI-OPD	$\mathbf{68.75}_{\pm 2.12}$	$\mathbf{60.42}_{\pm 2.23}$	$\mathbf{95.63}_{\pm 0.81}$	$\mathbf{10.90}_{\pm 1.25}$	$\mathbf{34.66}_{\pm 2.07}$	$\mathbf{54.07}_{\pm 0.80}$

Two consecutive valid windows above η_s trigger the supervision cutoff:

$$j^\star = \min\{j : Q_j > \eta_s,\ Q_{j+1} > \eta_s\}, \qquad c = \begin{cases} b_{j^\star} - 1, & \text{if such a pair exists,} \\ N, & \text{otherwise.} \end{cases} \tag{6}$$

Here, $j^\star$ indexes the first violating window pair, b_j is the first position of window j, and N is the complete mixed response length. The cutoff position c is the last token retained for supervision. A loss mask excludes both violating windows and all later tokens, without changing the generated text. The retained student prefix, teacher segment, and selected continuation receive the same OPD update as in Section 3.1. Appendix C.6 gives the detailed loss mask implementation and ATI-OPD full algorithm in Algorithm 1.

5 Experiments

We first examine whether ATI-OPD improves distribution alignment, optimization stability, and response lengths, then evaluate the student's independent reasoning. Next, we test the roles of local teacher guidance and supervision selection. Finally, we analyze which interventions improve answers and how teacher participation changes during training.

5.1 Experimental setup

Models and comparison settings. We train on DAPO-Math-17K (Yu et al., 2025) with a maximum response length of $N_{\max} = 23{,}552$ tokens, using two teacher → student pairs: JustRL-DeepSeek-1.5B (He et al., 2025) → DeepSeek-R1-Distill-Qwen-1.5B (Guo et al., 2025), and JustRL-Nemotron-1.5B (He et al., 2025) → OpenMath-Nemotron-1.5B (Moshkov et al., 2025). We refer to these pairs as JustRL-DeepSeek and JustRL-Nemotron, respectively. Our main configuration permits $M = 1$ intervention and applies OPD to retained student and teacher tokens. Vanilla OPD uses the same training settings as ATI-OPD. See Appendix A for detailed configurations.

Evaluation and diagnostic metrics. We evaluate on AIME 2024 (Codeforces, 2024), AIME 2025 (Zhang & Math-AI, 2025), AMC 2023 (Mathematical Association of America, 2023), AMO-Bench (Liu et al., 2026), and HMMT February 2026 (Dekoninck et al., 2026). We average final-answer correctness over $R = 16$ student-only responses per question, then weight the five tasks equally. Training stability diagnostics include sampled-token absolute gap, the absolute difference between teacher and student log probabilities of each response token, and teacher top-16 token coverage, the fraction of valid tokens among the teacher's 16 most likely choices under the same prefix. Length and optimization diagnostics include the response-length tail, maximum-length clipping rate, gradient norm, and student entropy. Detailed definitions for these metrics are in Appendices C.2 and C.3.

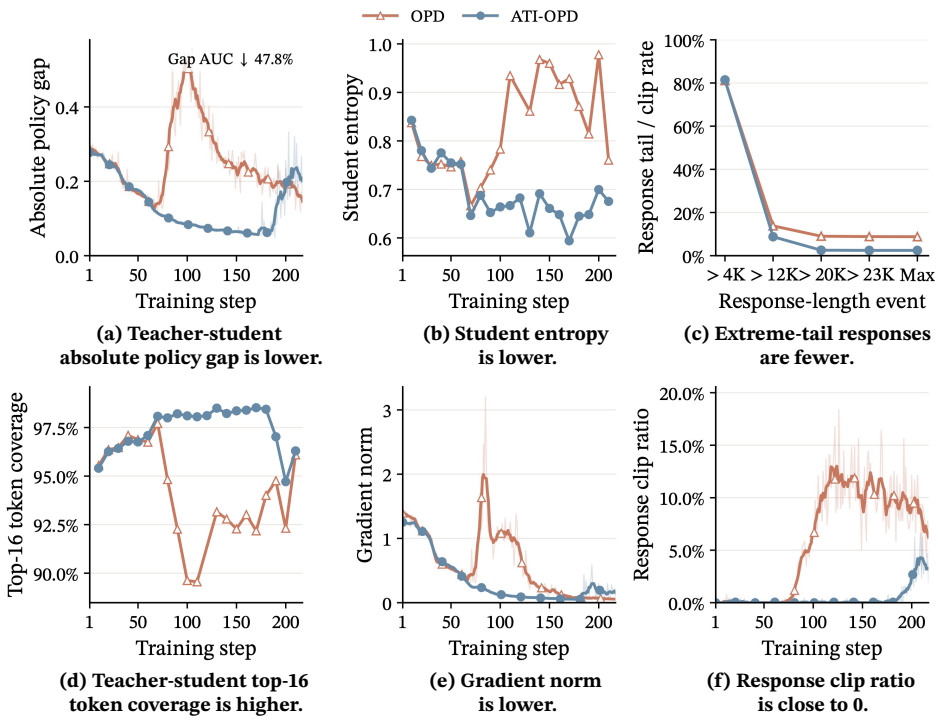


Figure 3: Training diagnostics for JustRL-DeepSeek: absolute gap, entropy, response-length tail (last 32 training steps), top-16 coverage, gradient norm, and clipping rate.

5.2 Training stability and student-only accuracy

Distribution alignment and optimization stability improve together under ATI-OPD. We track prediction disagreement and optimization to examine how teacher guidance affects the stability of training. On JustRL-DeepSeek, the absolute gap initially falls under both methods, then rebounds under vanilla OPD as top-16 coverage drops, as shown in Figure 3(a,d). Entropy increases and gradient spikes occur around the same stage, as shown in Figure 3(b,e). ATI-OPD maintains a lower gap and higher coverage for longer, reducing the raw gap AUC by 47.8%, with steadier gradients and lower, less variable entropy. These improvements support stabilizing training by revising generated prefixes and selecting where to apply supervision.

ATI-OPD concentrates length reduction in the extreme response tail. We compare length thresholds to distinguish suppression of extreme responses from a general loss of long reasoning. The methods have similar fractions above 4K tokens, but 13.82% of OPD responses exceed 12K over the last 32 steps, compared with 8.78% for ATI-OPD, as shown in Figure 3(c). The difference persists further into the tail, and ATI-OPD reaches the token limit less often, as shown in Figure 3(f). These are complete generated responses, including tokens excluded from supervision. ATI-OPD thus preserves substantial multi-thousand-token reasoning while reducing the extreme lengths associated with unstable training.

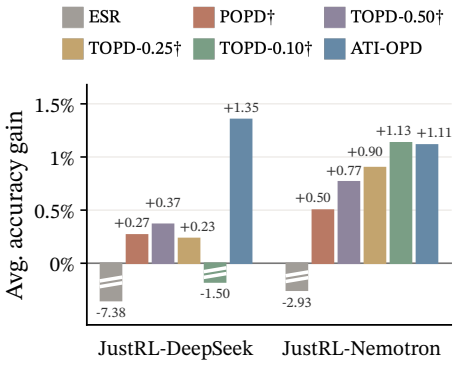
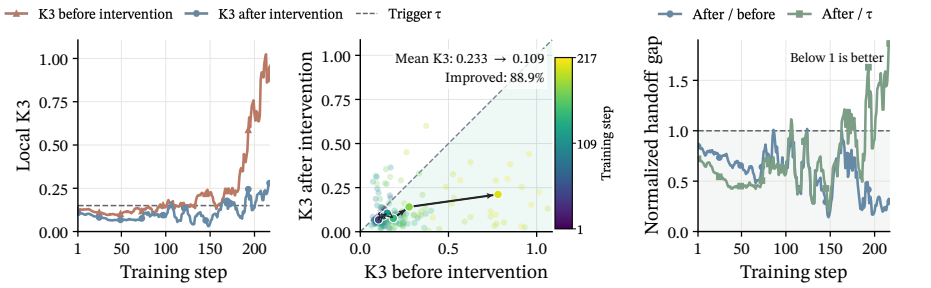


Figure 4: Gains over each OPD baseline on three shared tasks (%). †: published results.

Training-time teacher guidance transfers to student-only reasoning. Without teacher intervention at evaluation, ATI-OPD reaches 39.65% and 54.07% average accuracy on the two pairs, com-



(a) Local K3 decreases in most training steps.

(b) Local gap reduction persists (Blue), while post-handoff gaps rise late in training (Green).

Figure 5: Local handoff diagnostics for JustRL-DeepSeek: (a) student K3 after the original and replacement segments, shown as curves and a scatter plot; (b) post-handoff gap ratios. Appendix C.4 defines the matched student windows.

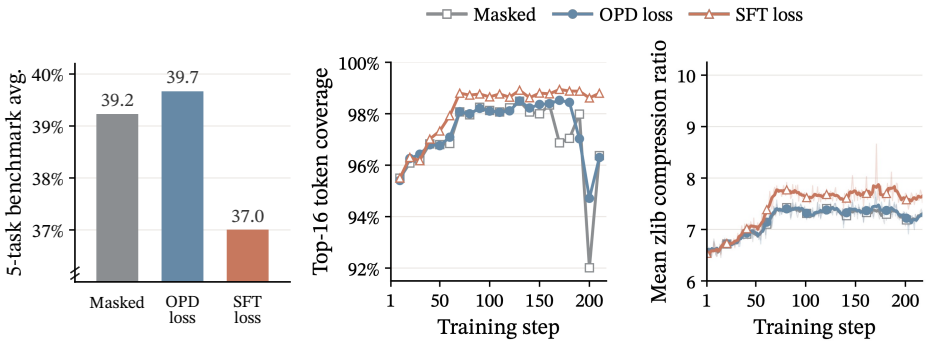
pared with 38.45% and 52.72% for vanilla OPD, as shown in Table 1. The absolute gains are 1.20% and 1.35%, with improvements on four of five tasks for JustRL-DeepSeek and all five for JustRL-Nemotron. ATI-OPD also improves accuracy with a 7B student, as shown in Appendix B.1, demonstrating the scalability of ATI-OPD in terms of model scale. ESR underperforms ATI-OPD on both pairs, suggesting that its fixed supervision cutoff discards useful learning signals from later reasoning. On three shared benchmarks, ATI-OPD gains 1.35% and 1.11% over its own OPD baseline, compared with published TOPD/POPD gains of -1.50% to $+0.37\%$ and $+0.50\%$ to $+1.13\%$, respectively, as shown in Figure 4. It achieves a larger gain on JustRL-DeepSeek and approaches the best gain on JustRL-Nemotron. Detailed results are in Appendix B.2.

5.3 Mechanisms of Intervention and Supervision

Local teacher rewrites improve alignment in the student's own continuation. To test whether guidance extends beyond the inserted teacher tokens, we compare student continuations immediately after the original triggering segment and its replacement, matching up to 16 valid tokens per side. The mean paired decrease in sampled-token K3 is 0.125 across training steps, with a decrease in 88.9% of training steps. This improvement occurs in the student's subsequent decisions, consistent with the rewrite providing a better starting point. The independently measured post-handoff gap nevertheless exceeds τ late in training, as shown in Figure 5, supporting continued monitoring after local guidance. Appendix C.4 gives the paired statistics and the scope of each diagnostic.

OPD loss is more effective than SFT for teacher-segment supervision. Teacher segments provide both context for continuation and demonstration text, making SFT cross-entropy on the teacher's sampled tokens a natural alternative at these positions. To examine how supervision at these positions affects training, we compare OPD, Masked, and SFT loss while keeping all inserted tokens as context. Masked retains teacher context but removes its direct loss, applying OPD only to retained student positions. Masked achieves 39.22% average accuracy across five tasks, above vanilla OPD's 38.45%, as shown in Figure 6(a). Retaining the OPD loss on teacher segments increases accuracy to 39.65%, compared with 36.99% for SFT cross-entropy. SFT loss has higher token coverage than OPD loss ($+2.75\%$), but also a higher mean zlib compression ratio ($+0.39$), indicating a greater tendency to repeat, as shown in Figure 6(b). These results favor using teacher probabilities across candidate tokens over fitting only the sampled teacher token.

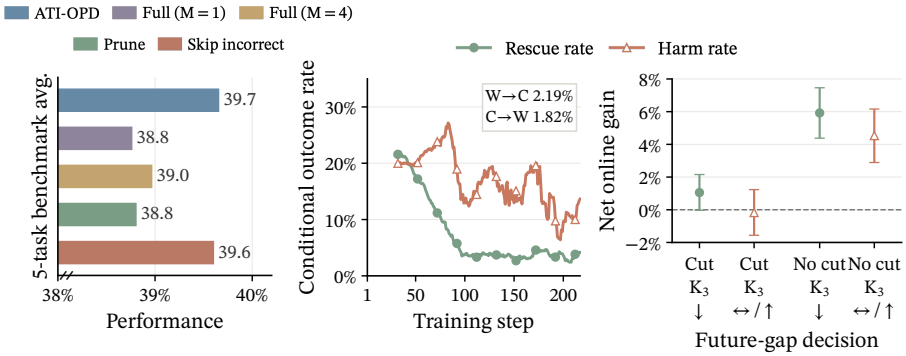
Student continuation provides useful supervision when retained selectively. We compare ATI-OPD with Full, which supervises the entire continuation, and Prune, which ends supervision at the teacher segment, to test the value of selecting later student decisions. With $M = 1$, ATI-OPD reaches 39.65% five-task accuracy, compared with 38.76% for Full and 38.80% for Prune, as shown in Figure 7(a). Increasing Full's budget to $M = 4$ yields 38.96%. The two extremes both perform worse, indicating that useful student supervision remains after the handoff and that additional interventions do not replace selecting its extent. Short teacher segments retain the gains:



(a) OPD supervision improves autonomous accuracy.

(b) OPD loss combines high token coverage with lower output compressibility.

Figure 6: Teacher-segment supervision for JustRL-DeepSeek: (a) five-task average accuracy; (b) top-16 coverage and token-zlib compression ratio (higher values indicate greater compressibility). Both diagnostics include teacher tokens for OPD, but exclude them for Masked and SFT.



(a) Adaptive truncation outperforms Full and Prune.

(b) Rescues outweigh harms across all rollouts.

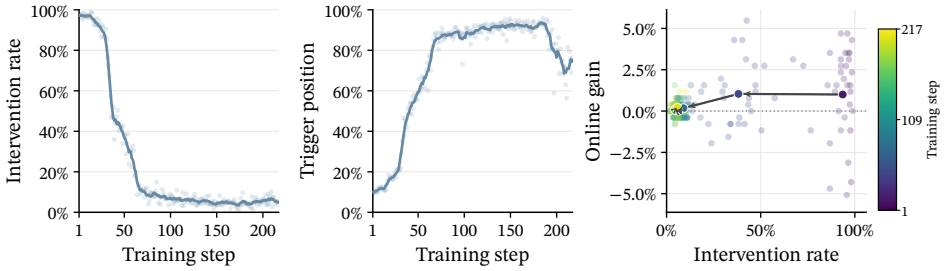
(c) Suffix-gate groups show distinct online gains.

Figure 7: JustRL-DeepSeek ablations and outcomes: (a) benchmark accuracy; (b) rescue/harm rates among originally incorrect/correct intervened responses, respectively; inset percentages use all rollouts. (c) Online gain (%) by supervision cutoff (Cut/No cut) and 128-token local K3 change (↓: decrease; ↔ / ↑: no decrease). Error bars are step-cluster bootstrap 95% confidence intervals; see Appendix C.5 for details.

$L_T = 128$ and 512 achieve 39.65% and 39.34%, and the tested trigger thresholds span 0.32%. See Appendices B.3 and B.4 for details.

5.4 Outcomes and Dynamics of Intervention

Retaining responses with incorrect final answers still preserves student-only accuracy. We compare answer transitions and correctness-based filtering to assess whether final-answer errors justify discarding an entire training trajectory. Rescues (W→C) and harms (C→W) account for 2.19% and 1.82% of all rollouts, respectively, yielding a net immediate accuracy change, or online gain, of +0.37%, as shown in Figure 7(b). Skip incorrect retains ATI-OPD's supervision selection but excludes an entire intervened trajectory from training if its final answer is incorrect. Its 39.60% five-task accuracy is close to ATI-OPD's 39.65%, as shown in Figure 7(a). Final-answer correctness evaluates the whole response, whereas OPD learns from selected token predictions within it. See Appendix C.5 for definitions. A response can therefore retain useful token-level supervision even when its final answer is incorrect.



(a) Intervention rate decreases. **(b) Intervention position shifts later.** **(c) Online gain shrinks as interventions become sparse.**

Figure 8: Intervention dynamics for JustRL-DeepSeek: (a) intervention rate; (b) first trigger position relative to the original response length; (c) intervention rate versus online gain over all rollouts (%). Colors mark training steps; arrows connect consecutive 32-step phase centers.

Persistent suffix divergence separates intervention gains more clearly than immediate local K3 changes. To assess whether later student decisions add information beyond the local change, we group responses by supervision cutoff and the 128-token local K3 change, as shown in Figure 7(c). Without a cutoff, online gain is $+5.92\%$ when local K3 decreases and $+4.54\%$ when it does not; the corresponding values with a cutoff are $+1.05\%$ and -0.17%. Differences between cutoff groups (4.7%–4.9%) exceed those between local K3 categories (1.2%–1.4%). Higher gains concentrate on continuations without persistent disagreement, making later divergence a more informative diagnostic of the intervention outcome.

Teacher guidance evolves from broad participation to selective intervention on failed trajectories. We track frequency, location, and original correctness to examine how the same intervention rule changes its focus during training. Between the first and last 32 steps, intervention rates fall from 93.57% to 5.51%, while original accuracy among intervened responses drops from 49.34% to 9.76%, as shown in Figure 8(a). The first trigger shifts from 15.1% to 79.0% of the original response length, as shown in Figure 8(b). Guidance thus increasingly targets a small subset of originally incorrect responses at later positions. Late observations also concentrate at smaller overall online gains: as interventions become sparse, fewer responses can benefit from an immediate answer change, as shown in Figure 8(c).

6 Limitations

ATI-OPD adds teacher decoding, student regeneration, and scoring costs. Selecting supervision after full generation does not reduce these decoding costs. Our main evaluation covers two 1.5B teacher–student pairs in mathematical reasoning, with one training run per setting. An exploratory 7B-student comparison is limited to four tasks, see Appendix B.1 for details, leaving broader model and task coverage for future work.

7 Conclusion

We introduce ATI-OPD to improve the stability of long-sequence on-policy distillation. Adaptive teacher intervention changes the prefix for subsequent student reasoning, and adaptive supervision selection limits learning from renewed disagreement after the handoff. ATI-OPD reduces distribution mismatch, gradient spikes, and extremely long responses. Local and supervision ablations support learning selectively from the student's continuation after guidance. ATI-OPD also improves average student-only accuracy across five mathematical reasoning benchmarks on both model pairs.

AI USE STATEMENT

We used GPT-5.6 Sol to give experimental design suggestions, implement the code, write visualization scripts, and interpret experimental results. GPT-6 Astra helped with manuscript translation, drafting, and editing. DeepSeek provided additional feedback on exposition and manuscript organization.

We did not use AI to develop theoretical models or conceptual frameworks. AI use for synthetic dataset generation, dataset cleaning or reformatting, hypothesis formulation or refinement, formulating mathematical claims, providing key proof elements, or writing proofs was not applicable. We manually reviewed AI-assisted work, including implementation code, visualization scripts, and experimental analyses, and take responsibility for the final paper's content and artifacts.

ETHICS STATEMENT

This study trains and evaluates released models on public mathematical reasoning data, with evaluation based on model-generated solutions. Models, datasets, and outputs are used solely for academic research. Distillation may transfer teacher errors or biases to the student. Deployment in specific applications requires further assessment of output reliability, safety, and potential misuse.

REPRODUCIBILITY STATEMENT

Section 4 describes ATI-OPD's divergence monitoring, adaptive teacher intervention, and supervision selection rules. Appendix C.6 gives the complete trajectory construction procedure for the main configuration in Algorithm 1. Section 5.1 and Appendix A specify the model and data sources, training and evaluation settings, and ablation variants. Appendix C details the optimization objective, evaluation and diagnostic metrics, and aggregation across training steps.

REFERENCES

Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=3zKtaqxLhW.

MAA Codeforces. American invitational mathematics examination-aime 2024, 2024.

Jasper Dekoninck, Nikola Jovanović, Tim Gehrunger, Kári Rögnvaldsson, Ivo Petrov, Chenhao Sun, and Martin Vechev. Beyond benchmarks: Matharena as an evaluation platform for mathematics with llms, 2026. URL https://arxiv.org/abs/2605.00674.

Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. Minillm: Knowledge distillation of large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=5h0qf7IBZZ.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638, 2025.

Bingxiang He, Zekai Qu, Zeyuan Liu, Yinghao Chen, Yuxin Zuo, Cheng Qian, Kaiyan Zhang, Weize Chen, Chaojun Xiao, Ganqu Cui, et al. Justrl: Scaling a 1.5 b llm with a simple rl recipe, 2025. URL https://arxiv.org/abs/2512.16649.

Ijun Jang, Jewon Yeom, Juan Yeo, Hyunggyu Lim, and Taesup Kim. Stable on-policy distillation through adaptive target reformulation. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 42217–42227, San Diego, California, United States, 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.2094. URL https://aclanthology.org/2026.findings-acl.2094/.

Li Jiang, Haoran Xu, Yichuan Ding, and Amy Zhang. Trajectory-refined distillation. *arXiv preprint arXiv:2606.08432*, 2026.

Woogyeol Jin, Taywon Min, Yongjin Yang, Dennis Wei, Yi Zhou, Swanand Ravindra Kadhe, Nathalie Baracaldo, and Kimin Lee. Entropy-aware on-policy distillation of language models, 2026. URL https://arxiv.org/abs/2603.07079.

Jongwoo Ko, Sungnyun Kim, Tianyi Chen, and Se-Young Yun. Distillm: Towards streamlined distillation for large language models. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. URL https://openreview.net/forum?id=lsHZNNoC7r.

Jahyun Koo, Yerin Hwang, Yongil Kim, Taegwan Kang, Hyunkyung Bae, and Kyomin Jung. SWITCH: Studying with teacher for knowledge distillation of large language models. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 3733–3746, Albuquerque, New Mexico, 2025. Association for Computational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findings-naacl.206. URL https://aclanthology.org/2025.findings-naacl.206/.

Yaxuan Li, Yuxin Zuo, Bingxiang He, Jinqian Zhang, Chaojun Xiao, Cheng Qian, Tianyu Yu, Huanang Gao, Wenkai Yang, Zhiyuan Liu, et al. Rethinking on-policy distillation of large language models: Phenomenology, mechanism, and recipe, 2026. URL https://arxiv.org/abs/2604.13016.

Junlin Liu, Shengnan An, Shuang Zhou, Dan Ma, Yehao Lin, Xinxuan Lv, Xuanlin Wang, Xiaoyu Li, Ziwen Wang, Xuezhi Cao, and Xunliang Cai. AMO-bench: Large language models still struggle in high school math competitions. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 2120–2137, San Diego, California, United States, 2026. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.101. URL https://aclanthology.org/2026.findings-acl.101/.

Feng Luo, Yu-Neng Chuang, Guanchu Wang, Zicheng Xu, Xiaotian Han, Tianyi Zhang, and Vladimir Braverman. Demystifying opd: Length inflation and stabilization strategies for large language models, 2026. URL https://arxiv.org/abs/2604.08527.

Mathematical Association of America. American Mathematics Competitions (AMC 2023), 2023. URL https://maa.org/student-programs/amc/.

Ivan Moshkov, Darragh Hanley, Ivan Sorokin, Shubham Toshniwal, Christof Henkel, Benedikt Schifferer, Wei Du, and Igor Gitman. Aimo-2 winning solution: Building state-of-the-art mathematical reasoning models with openmathreasoning dataset, 2025. URL https://arxiv.org/abs/2504.16891.

Jingyu Peng, Maolin Wang, Hengyi Cai, Yuchen Li, Kai Zhang, Shuaiqiang Wang, Dawei Yin, and Xiangyu Zhao. Adaswitch: Balancing exploration and guidance in knowledge distillation via adaptive switching, 2025. URL https://arxiv.org/abs/2510.07842.

John Schulman. Approximating kl divergence. *John Schulman's Homepage*, 5, 2020.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

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

Bing Wang, Shaotian Yan, Chen Shen, Sinan Fan, Ximing Li, Rui Miao, Xiaosong Yuan, Zhanming Shen, Jieping Ye, et al. Backtracking when it strays: Mitigating dual exposure biases in llm reasoning distillation, 2026. URL https://arxiv.org/abs/2605.19433.

Yan Xie, Sijie Zhu, Tiansheng Wen, Bo Chen, and Yifei Wang. On the position bias of on-policy distillation, 2026. URL https://arxiv.org/abs/2606.22600.

Haoran Xin, Anhao Zhao, Ying Sun, Jin Li, Xiaoyu Shen, and Hui Xiong. Escaping the kl agreement trap in on-policy distillation. *arXiv preprint arXiv:2606.09471*, 2026.

Haolei Xu, Xiaowen Xu, Haiwen Hong, Zixuan Ni, Hongxing Li, Yiwen Qiu, Weiming Lu, and Yongliang Shen. Pass the baton: Trajectory-relayed on-policy distillation, 2026. URL https://arxiv.org/abs/2607.26057.

Wenda Xu, Rujun Han, Zifeng Wang, Long T. Le, Dhruv Madeka, Lei Li, William Yang Wang, Rishabh Agarwal, Chen-Yu Lee, and Tomas Pfister. Speculative knowledge distillation: Bridging the teacher-student gap through interleaved sampling. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=EgJhwYR2tB.

Zhicheng Yang, Zhijiang Guo, Yifan Song, Minrui Xu, Yongxin Wang, Yiwei Wang, Xiaodan Liang, and Jing Tang. Prune-opd: Efficient and reliable on-policy distillation for long-horizon reasoning. *arXiv preprint arXiv:2605.07804*, 2026.

Deheng Ye, Zhao Liu, Mingfei Sun, Bei Shi, Peilin Zhao, Hao Wu, Hongsheng Yu, Shaojie Yang, Xipeng Wu, Qingwei Guo, Qiaobo Chen, Yinyuting Yin, Hao Zhang, Tengfei Shi, Liang Wang, Qiang Fu, Wei Yang, and Lanxiao Huang. Mastering complex control in MOBA games with deep reinforcement learning. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 6672–6679. AAAI Press, 2020. URL https://aaai.org/ojs/index.php/AAAI/article/view/6144.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Juncai Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Ru Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Yonghui Wu, and Mingxuan Wang. DAPO: an open-source LLM reinforcement learning system at scale. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/a4277440d50f1f15d2cb4c14f7e0c0d2-Abstract-Conference.html.

Dongxu Zhang, Zhichao Yang, Sepehr Janghorbani, Jun Han, Andrew Ressler II, Qian Qian, Gregory D Lyng, Sanjit Singh Batra, and Robert E. Tillman. Fast and effective on-policy distillation from reasoning prefixes. In Maria Liakata, Viviane P. Moreira, Jiajun Zhang, and David Jurgens (eds.), *Findings of the Association for Computational Linguistics: ACL 2026*, pp. 25553–25569, San Diego, California, United States, 2026a. Association for Computational Linguistics. ISBN 979-8-89176-395-1. doi: 10.18653/v1/2026.findings-acl.1276. URL https://aclanthology.org/2026.findings-acl.1276/.

Yaocheng Zhang, Jiajun Chai, Yuqian Fu, Songjun Tu, Xiaohan Wang, Wei Lin, Guojun Yin, Qichao Zhang, Yuanheng Zhu, and Dongbin Zhao. Are full rollouts necessary for on-policy distillation?, 2026b. URL https://arxiv.org/abs/2605.31490.

Yifan Zhang and Team Math-AI. American invitational mathematics examination (aime) 2025. *Hugging Face dataset*, 2025.

Jin-Gang Zhou, Yu-Yi Zhou, Haiyang Guo, Xu-Kai Wang, Shuai Feng, Si-Rui Gao, Jian Xu, Qing-Pei Guo, and Xu-Yao Zhang. Teacher should think ahead: Adaptive continuations for reliable on-policy distillation. 2026. URL https://api.semanticscholar.org/CorpusID:292205757.

Siqi Zhu, Xuyan Ye, Hongyu Lu, Weiye Shi, and Ge Liu. The many faces of on-policy distillation: Pitfalls, mechanisms, and fixes, 2026. URL https://arxiv.org/abs/2605.11182.

Zhou Ziheng, Jiaqi Li, Huacong Tang, Ying Nian Wu, and Demetri Terzopoulos. Less is more: Early stopping rollout for on-policy distillation. *arXiv preprint arXiv:2605.27028*, 2026.

A HYPERPARAMETERS AND MODEL COVERAGE

Table 2 summarizes the detailed settings for ATI-OPD. Vanilla OPD shares all settings except those for teacher intervention and supervision selection. Here, ι_T controls OPD on teacher tokens.

Table 2: Main settings for ATI-OPD training, student-only evaluation, and local intervention.

Component	Setting
Training data	DAPO-Math-17K
Max prompt length	1,024 tokens
Max response length	$N_{\max} = 23{,}552$ tokens
Batch size	64
Rollouts per prompt	4
Training steps	217
Optimizer	AdamW, constant learning rate 10^{-6}
Weight decay	0.01
Gradient clip	1.0
Student rollout	temperature 1.0, top-p 1.0
Evaluation	$R = 16$, temperature 0.7, top-p 0.95
OPD candidate set	Student top-16
Student monitor segment	$L_S = 128$
Teacher intervention segment	$L_T = 128$
Trigger	Mean clipped divergence score, $\tau = 0.15$
Teacher intervention sampling	temperature 0.6, top-p 0.95, top-k 20
Teacher samples per intervention	1
Intervention budget	$M = 1$
Future window	$L_F = 64$ tokens, at least 32 valid tokens
Reference pool	2,048 window means from correct original responses
Reference contribution	At most 8 windows per correct response
Initial reference threshold	$\eta_s = \tau$ until 128 windows are available
Future cutoff	η_s: delayed 90th percentile, persistence 2
Teacher-segment objective	OPD loss, $\iota_T = 1$

Ablation settings. Teacher-segment OPD loss is the main method, applying distribution supervision to both retained student and teacher tokens ($\iota_T = 1$). Teacher-segment masking (Masked) sets $\iota_T = 0$, retaining teacher tokens as context but removing their direct loss. The remaining retained student tokens still receive OPD supervision. The SFT loss variant replaces OPD on teacher tokens with teacher-forcing cross-entropy weighted by $\lambda_{\mathrm{SFT}} = 0.1$; Appendix C.1 gives the objective. Full supervises the complete continuation after intervention, with intervention budgets of $M = 1$ and $M = 4$; Prune restricts supervision to positions through the teacher segment. Both variants still generate the complete continuation. The current adaptive supervision implementation requires $M = 1$. Skip incorrect follows ATI-OPD's intervention and supervision selection, then excludes an entire intervened trajectory from training if its final answer is incorrect.

ESR implementation. For ESR, we disable teacher intervention and set a fixed 100-token OPD supervision horizon (Ziheng et al., 2026). The student still generates responses up to the 23,552-token limit. A separate mask restricts teacher-derived token rewards and the actor policy loss to the first 100 valid response tokens.

B ADDITIONAL COMPARISONS AND INTERVENTION ANALYSIS

B.1 EXPLORATORY EVALUATION WITH A LARGER STUDENT

Exploratory results with a larger student. To examine whether the main result extends beyond the 1.5B students, we compare OPD and ATI-OPD with DeepSeek-R1-Distill-Qwen-7B as the student and JustRL-DeepSeek-1.5B as the teacher. The teacher is unchanged in size. ATI-OPD is higher on all four available tasks, with a four-task average of 44.72% compared with 43.40% for OPD.

Table 3: Exploratory student-only accuracy (%) with a 7B student. Evaluation uses temperature 1.0, top-$p = 1.0$, no top-k truncation, with 16 responses per question. Avg. equally weights the four available tasks.

Method	AIME24	AIME25	AMC23	AMO-Bench	Avg.
OPD	46.46	37.71	86.41	3.04	43.40
ATI-OPD	48.54	39.17	87.66	3.53	44.72

B.2 Contextual comparison with published length control

We compare ATI-OPD with published length control methods (Zhang et al., 2026b). Table 4 reports gains over each study's own vanilla OPD baseline on three shared benchmarks. The published results average the final ten evaluations, whereas ours use the final training checkpoint. The original response limit is 15,360 tokens, compared with 23,552 in our experiments. The TOPD and POPD results show that controlling training trajectory length can improve distillation, with gains varying across tasks and configurations. ATI-OPD achieves a larger average gain on JustRL-DeepSeek and an average gain close to that of the best length control configuration on JustRL-Nemotron.

B.3 Teacher segment length sensitivity

Short teacher segments retain the strongest average performance in the tested range. We vary the teacher's segment length to test whether more teacher-generated text is needed after an intervention, keeping the student segment length fixed at $L_S = 128$. The five-task averages for $L_T = 128$, 512, and 1024 are 39.65%, 39.34%, and 38.87%, respectively (Table 5). Longer replacements therefore bring no average improvement in these runs. The 512-token configuration performs best on AIME25, so the effect is task dependent; the average result supports returning control to the student after a short teacher segment.

B.4 Trigger threshold sensitivity

The five-task mean changes little across the three tested trigger thresholds. The threshold controls how much local disagreement is tolerated before teacher intervention. We compare $\tau = 0.10$, 0.15, and 0.30 to assess whether the main result depends on a narrowly chosen value. Their five-task averages are 39.94%, 39.65%, and 39.63%, respectively (Table 6). The mean decreases slightly as the threshold increases, with a total difference of 0.32%, although individual tasks vary more. With one training run per setting, this comparison describes the observed sensitivity rather than establishing a statistically optimal threshold.

C Implemented objectives and metric definitions

Our training implementation follows Li et al. (2026). This section defines the implemented OPD objective, training stability and intervention metrics, and mixed-trajectory construction.

C.1 Exact Top-K OPD surrogate

To provide distribution supervision over the student's high-probability candidates, the OPD surrogate weights candidate log-probability differences by student probabilities:

$$w_t(v) = \frac{\pi_{\theta_{\text{old}}}(v \mid h_t)}{\sum_{u \in \mathcal{V}_t} \pi_{\theta_{\text{old}}}(u \mid h_t)}, \tag{7}$$

$$A_t(v) = w_t(v) \left[\log \pi_T(v \mid h_t) - \log \pi_{\theta_{\text{old}}}(v \mid h_t)\right]. \tag{8}$$

Here, u, v are candidate tokens, and $\mathcal{V}_t$ contains candidates with finite student and teacher log-probabilities, drawn from $\mathcal{C}_t = \text{TopK}(\pi_{\theta_{\text{old}}}(\cdot \mid h_t), K)$. TopK returns the K highest-probability tokens given the prefix h_t; we use $K = 16$. $w_t(v)$ is the student probability renormalized over $\mathcal{V}_t$, and $A_t(v)$ is the corresponding weighted log-probability difference. This is the `student_p`

Table 4: Absolute accuracy gains over each method's own OPD baseline on three shared benchmarks (%). The number following TOPD denotes the truncation length as a fraction of the original maximum response length. † denotes gains calculated from the accuracies published in the corresponding work (Zhang et al., 2026b).

Method	ΔAIME24	ΔAIME25	ΔAMC23	ΔAvg.
JustRL-DeepSeek-1.5B → DeepSeek-R1-Distill-Qwen-1.5B				
POPD†	+1.70	+0.50	−1.40	+0.27
TOPD-0.50†	+1.00	**+0.80**	−0.70	+0.37
TOPD-0.25†	+2.00	−0.10	−1.20	+0.23
TOPD-0.10†	+1.40	−1.80	−4.10	−1.50
ATI-OPD	**+2.92**	−0.42	**+1.56**	**+1.35**
JustRL-Nemotron-1.5B → OpenMath-Nemotron-1.5B				
POPD†	+1.30	+0.40	−0.20	+0.50
TOPD-0.50†	+1.30	+0.70	+0.30	+0.77
TOPD-0.25†	+1.40	+0.60	**+0.70**	+0.90
TOPD-0.10†	**+2.00**	**+1.10**	+0.30	**+1.13**
ATI-OPD	+1.67	+1.04	+0.63	+1.11

Table 5: Sensitivity to teacher segment length for JustRL-DeepSeek (accuracy, %), with $L_S = 128$. The first four benchmarks use the training-step-217 evaluation; HMMT February 2026 uses the corresponding standalone checkpoint evaluation. Avg. equally weights the five tasks.

L_T	AIME24	AIME25	AMC23	AMO-Bench	HMMT Feb26	Avg.
128	49.17	33.54	88.59	2.72	24.24	39.65
512	48.33	36.88	87.19	2.72	21.59	39.34
1024	46.46	36.46	85.94	2.56	22.92	38.87

weighting used in the implementation. The teacher scores every student top-K candidate, whether or not it also belongs to the teacher's top-K set.

To compute candidate probability ratios between the current student and the rollout student, we first clip the log-ratio for numerical stability:

$$\rho_t(v) = \exp(\mathrm{clip}[\log \pi_\theta(v \mid h_t) - \log \pi_{\theta_{\mathrm{old}}}(v \mid h_t), -20, 20]), \tag{9}$$

where $\rho_t(v)$ is the candidate probability ratio, and $\mathrm{clip}(z, a, b) = \min(\max(z, a), b)$ bounds its input within $[a, b]$. Here, the log-ratio is clipped to $[-20, 20]$.

To limit the magnitude of policy updates, we use the probability-ratio clipping of Proximal Policy Optimization (PPO) (Schulman et al., 2017) and its dual-clip form (Ye et al., 2020). The loss for each candidate token is

$$\ell_t(v) = \begin{cases} \min\{-c_{\mathrm{clip}} A_t(v), \max[-\rho_t(v)A_t(v), -\bar{\rho}_t(v)A_t(v)]\}, & A_t(v) < 0, \\ \max[-\rho_t(v)A_t(v), -\bar{\rho}_t(v)A_t(v)], & A_t(v) \geq 0, \end{cases} \tag{10}$$

where $\bar{\rho}_t(v) = \mathrm{clip}(\rho_t(v), 1 - \epsilon_\ell, 1 + \epsilon_h)$ is the clipped probability ratio. ϵ_ℓ and ϵ_h control the clipping ranges below and above 1, and c_{clip} provides an additional clipping coefficient for negative advantages. The main setting uses $\epsilon_\ell = \epsilon_h = 0.2$ and $c_{\mathrm{clip}} = 3$; c_{clip} is distinct from the trajectory cutoff position c.

To update only the selected supervision positions, we sum the candidate losses at each position and average over tokens selected for OPD:

$$\mathcal{L}_{\mathrm{OPD}} = \frac{\sum_t m_t^{\mathrm{OPD}} \sum_{v \in \mathcal{V}_t} \ell_t(v)}{\sum_t m_t^{\mathrm{OPD}}}. \tag{11}$$

Here, m_t^{OPD} is the supervision mask defined in Equation 41, which determines whether the current position contributes to distillation.

To compare supervision objectives on the teacher segment, the ablations change its mask or add an objective. OPD sets $\iota_T = 1$ for teacher-generated tokens. Masked sets $\iota_T = 0$ while retaining these

Table 6: Sensitivity to the intervention trigger threshold for JustRL-DeepSeek (accuracy, %). Avg. is the equal-weight mean of five benchmarks.

τ	AIME24	AIME25	AMC23	AMO-Bench	HMMT Feb26	Avg.
0.10	47.71	37.50	86.72	3.37	24.43	39.94
0.15	49.17	33.54	88.59	2.72	24.24	39.65
0.30	47.71	34.58	88.91	2.88	24.05	39.63

tokens as context for subsequent student tokens. SFT also excludes teacher-generated tokens from the OPD, task-reward, entropy, and KL masks, and adds a cross-entropy objective:

$$\mathcal{L}_{\mathrm{SFT}} = -\frac{1}{\sum_t m_t^T} \sum_t m_t^T \log \pi_\theta(\widetilde{y}_t \mid h_t), \tag{12}$$

$$\mathcal{L} = \mathcal{L}_{\mathrm{OPD}} + \lambda_{\mathrm{SFT}} \mathcal{L}_{\mathrm{SFT}}, \tag{13}$$

where $m_t^T \in \{0, 1\}$ marks only the inserted teacher tokens, and $h_t = (x, \widetilde{y}_{<t})$ is the mixed-trajectory prefix. λ_{SFT} weights the cross-entropy term, and $\mathcal{L}$ is the combined training loss. The SFT ablation uses $\lambda_{\mathrm{SFT}} = 0.1$.

C.2 Token- and trajectory-level stability metrics

Response masks and token sources. The logged absolute gap, token coverage, and token-zlib compression ratio use the response mask supplied by the training configuration. Vanilla OPD contains only student-generated tokens. In the main ATI-OPD configuration, the mask includes student tokens and inserted teacher tokens in the complete mixed response. It also includes positions beyond the supervision cutoff: that cutoff changes the distillation loss mask, while the response mask and generated text remain unchanged. Masked and the teacher-segment SFT variant exclude teacher tokens from this diagnostic mask. Thus, the aggregate curves describe the tokens selected by each configuration, and their differences should not be attributed entirely to changes in the student's own predictions. The separate local comparison in Appendix C.4 excludes teacher tokens in all cases.

Sampled-token absolute gap. To measure teacher–student divergence at the prefixes actually visited, we average the absolute sampled-token log-probability difference over valid positions:

$$d_{it}^{\mathrm{abs}} = \left| \log \pi_T(y_{it} \mid h_{it}) - \log \pi_{\theta_{\mathrm{old}}}(y_{it} \mid h_{it}) \right|, \tag{14}$$

$$\mathrm{GapAbs} = \frac{\sum_{i,t} m_{it}^{\mathrm{gap}} d_{it}^{\mathrm{abs}}}{\sum_{i,t} m_{it}^{\mathrm{gap}}}. \tag{15}$$

Here, i and t index the response and token position, respectively, y_{it} is the actual sampled token, and $h_{it} = (x_i, y_{i,<t})$ consists of prompt x_i and the preceding response tokens. d_{it}^{abs} is the absolute difference at one position, and GapAbs is the aggregate metric. The validity mask $m_{it}^{\mathrm{gap}} \in \{0, 1\}$ is 1 only when the response mask is valid and both log-probabilities are finite, no greater than 10^{-6}, and at most 10^4 in absolute value.

Top-16 token coverage. To measure how often actual response tokens belong to the teacher's high-probability candidates, we compute teacher top-16 token coverage under the same question and preceding text:

$$\mathrm{Cov}_{16} = \frac{\sum_{i,t} m_{it}^{16} \mathbf{1}[y_{it} \in \mathcal{T}_{it}^{16}]}{\sum_{i,t} m_{it}^{16}}. \tag{16}$$

Here, $\mathcal{T}_{it}^{16}$ contains the teacher's 16 highest-probability tokens under prefix h_{it}, and $\mathbf{1}[\cdot]$ is the indicator function. m_{it}^{16} marks valid positions, requiring a valid response token and at least one finite teacher top-16 log-probability at that position.

Student entropy. We use student entropy to track uncertainty in the sampling distribution. To reduce communication in the distributed cluster, we approximate entropy using the top-16 probabilities and

a residual bucket containing the remaining probability mass:

$$\widehat{H}_{it}^{16} = -\sum_{v \in \mathcal{S}_{it}^{16}} p_{itv} \log p_{itv} - r_{it} \log r_{it}. \tag{17}$$

Here, $\mathcal{S}_{it}^{16}$ contains the student's 16 highest-probability tokens under prefix h_{it}, and $p_{itv} = \pi_{\theta_{\text{old}}}(v \mid h_{it})$ is the probability of candidate token v. The remaining probability mass is $r_{it} = \text{clip}(1 - \sum_{v \in \mathcal{S}_{it}^{16}} p_{itv}, 0, 1)$, with $0 \log 0 := 0$.

Response tail and clipping. To characterize the tail of the response-length distribution, we measure the fraction of responses exceeding each length threshold and the clipping rate at the maximum token budget:

$$S(q) = \frac{1}{B} \sum_{i=1}^{B} \mathbf{1}[N_i > q], \tag{18}$$

$$R_{\text{clip}} = \frac{1}{B} \sum_{i=1}^{B} \mathbf{1}[N_i = N_{\max}]. \tag{19}$$

Here, B is the number of responses in a batch, N_i is the valid length of response i, q is the length threshold, and $N_{\max}$ is the maximum token budget. For ATI-OPD, these are the complete mixed responses after the intervention, including positions beyond the supervision cutoff. Their lengths exclude discarded text from the initial full rollout and therefore do not measure the total number of tokens generated during training. $S(q)$ is the survival rate beyond length q, and R_{clip} is the fraction of responses that reach $N_{\max}$. The "$N > 12$K ratio" in the main text uses $q = 12000$. The tail plot reports the arithmetic mean of the step-level rates over the final 32 training steps. Its Max entry uses R_{clip}, whereas other length thresholds use $S(q)$; the former includes responses equal to the length limit.

Token-zlib compression ratio. To help diagnose repetition in responses, the supervision-objective ablation uses the logged mean token-zlib compression ratio over the last 4,096 valid tokens of each response, after applying the response mask. Responses with fewer than 4,096 valid tokens have no compression-ratio observation and are excluded from this mean. A higher value indicates that the token sequence is more compressible. The loss-ablation differences in coverage and compressibility are means over the final 32 training steps, using the available observations for each metric. Coverage is logged less frequently, so its late-stage mean uses three observations in these runs; the compression statistic uses step-level observations.

Gradient norm. To monitor gradient spikes during optimization, we record the global ℓ_2 norm before gradient clipping:

$$G_s = \left(\sum_{p} \|\nabla_{\theta_p} \mathcal{L}_s\|_2^2 \right)^{1/2}. \tag{20}$$

Here, G_s is the gradient norm at training step s, p indexes parameter tensors, θ_p denotes the corresponding parameters, and $\mathcal{L}_s$ is the training loss at that step.

Student-only accuracy. To evaluate the trained student's independent reasoning ability, we measure benchmark accuracy without teacher intervention and average across tasks:

$$\text{Acc}_b = \frac{1}{R|\mathcal{Q}_b|} \sum_{q \in \mathcal{Q}_b} \sum_{r=1}^{R} \mathbf{1}[\text{judge}(y_{qr}) = 1], \tag{21}$$

$$\text{Avg} = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \text{Acc}_b \,. \tag{22}$$

Here, $\mathcal{B}$ is the set of evaluation benchmarks, and $\mathcal{Q}_b$ is the question set for benchmark b. We sample $R = 16$ independent responses per question. y_{qr} is the r-th response to question q, and judge returns a correctness label based on the final answer (1 for correct and 0 for incorrect). Acc_b is the accuracy

on benchmark b, and Avg is the mean accuracy across benchmarks. The main evaluation set contains AIME24, AIME25, AMC23, AMO-Bench, and HMMT February 2026. For the small subscripts in Table 1, we use the conditional Bernoulli plug-in standard error

$$\widehat{\mathrm{SE}}_b = \sqrt{\frac{\mathrm{Acc}_b(1-\mathrm{Acc}_b)}{R|\mathcal{Q}_b|}}, \tag{23}$$

$$\widehat{\mathrm{SE}}_{\mathrm{Avg}} = \frac{1}{|\mathcal{B}|}\left(\sum_{b\in\mathcal{B}} \widehat{\mathrm{SE}}_b^2\right)^{1/2}. \tag{24}$$

Both quantities are multiplied by 100 when reported in percentage points. This calculation treats the sampled responses as independent conditional on a fixed checkpoint and question set; it does not estimate variability across training runs or question sets. The archived benchmark summaries do not contain the per-question outcomes needed for a paired significance analysis, so we report no benchmark significance tests.

C.3 Across-step aggregation and plotting conventions

To compare overall divergence during training, we compute the area under the curve (AUC), normalized by the training-step interval, and report the reduction of ATI-OPD relative to OPD:

$$\mathrm{AUC}(z) = \frac{1}{s_J - s_1}\sum_{j=1}^{J-1}\frac{z_j + z_{j+1}}{2}(s_{j+1} - s_j), \tag{25}$$

$$\mathrm{Red}_{\mathrm{AUC}} = 100\frac{\mathrm{AUC}_{\mathrm{OPD}} - \mathrm{AUC}_{\text{ATI-OPD}}}{\mathrm{AUC}_{\mathrm{OPD}}}. \tag{26}$$

Here, the J observations correspond to training steps $s_1 < \cdots < s_J$, and z_j is the raw metric value at observation j. $\mathrm{AUC}_{\mathrm{OPD}}$ and $\mathrm{AUC}_{\text{ATI-OPD}}$ are the AUC values for the two methods, and $\mathrm{Red}_{\mathrm{AUC}}$ is the relative reduction (%). The absolute-gap AUC reduction integrates the raw step-level values from Equation 15. We integrate over the actual recorded training steps without smoothing or filling missing observations; the trapezoidal rule connects adjacent available points. For the main comparison, the integration interval is steps 1–217.

C.4 Intervention location, local effect, and supervised suffix length

To characterize teacher participation and intervention locations, we measure the batch intervention rate, the mean number of interventions per response, and the first trigger position relative to the original response length:

$$r = \frac{1}{B}\sum_{i=1}^{B}\mathbf{1}[K_i > 0], \tag{27}$$

$$\overline{K} = \frac{1}{B}\sum_{i=1}^{B}K_i, \tag{28}$$

$$P_{\mathrm{first}} = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}}\frac{s_i^{(1)}}{\max(N_i^{\mathrm{raw}}, 1)}, \tag{29}$$

where K_i counts completed replacements in response i, and r and $\overline{K}$ are the batch intervention rate and mean intervention count, respectively. $s_i^{(1)}$ is the zero-based token offset of the first intervention, N_i^{raw} is the length of the complete original student response, and $\mathcal{I} = \{i : K_i > 0\}$ indexes responses that receive an intervention. P_{first} is averaged only over this set and is undefined when no valid intervened response is available. An intervention is counted when a replacement is completed, regardless of whether the final answer changes from incorrect to correct.

To measure local divergence changes after a teacher replacement, we pair the original student continuation immediately after the triggering segment with the new student continuation immediately

after the teacher segment. We compare means of the clipped divergence score $\bar{g}_t$ from Equation 43 over up to 16 valid student tokens at the start of each region:

$$G_{\text{before}} = \frac{1}{n}\sum_{j=1}^{n} b_j, \tag{30}$$

$$G_{\text{after}} = \frac{1}{n}\sum_{j=1}^{n} a_j, \tag{31}$$

$$\Delta G = G_{\text{before}} - G_{\text{after}}. \tag{32}$$

Here, $(b_1, \ldots, b_{n_b})$ is the sequence of finite divergence scores following the original triggering segment, $(a_1, \ldots, a_{n_a})$ is the sequence of finite scores in the student continuation after the teacher segment, and $n = \min(16, n_b, n_a)$. G_{before} and G_{after} are the local means before and after intervention, respectively, and ΔG is their difference. Teacher-generated tokens are excluded from both windows. The two windows follow different versions of the preceding segment; this is a paired descriptive comparison of different continuations. An intervention contributes only when $n > 0$.

For each training step s, the logged paired decrease ΔG_s is the mean of ΔG over valid intervention pairs. We report its mean over valid steps and the fraction of steps with a positive decrease:

$$\frac{1}{|\mathcal{S}_{\text{valid}}|}\sum_{s\in\mathcal{S}_{\text{valid}}} \mathbf{1}[\Delta G_s > 0], \tag{33}$$

where $\mathcal{S}_{\text{valid}}$ is the set of training steps with valid paired statistics. The mean paired decrease is 0.125 over 217 steps; the 88.9% decrease rate corresponds to 193 of those steps, not to a fraction of individual interventions. The paired-decrease curve uses a centered 7-step mean for display; the numerical summary and decrease rate use unsmoothed step values.

We separately track the post-handoff score over up to 16 valid student tokens, without requiring a valid original comparison window. For comparison with the trigger threshold, we average this score within each step, smooth the step means with a centered 7-step mean, and divide by τ. This diagnostic can use a different set of interventions and different window lengths from the paired statistic. It therefore cannot be combined with ΔG_s to reconstruct a paired original mean or an after-to-before ratio. The run contains 13,186 intervention boundaries, with 12,949 valid pairs and 13,017 valid independent post-handoff windows. Of the independent windows, 68 have no original comparison window; eight paired comparisons also use fewer tokens than their independent post-handoff windows. Paired before/after means and their ratios must therefore be computed from the same matched windows.

To quantify how much of the student suffix receives supervision under Equation 6, we compute the supervised suffix length, its fraction of the generated suffix, and an indicator for no supervision cutoff:

$$H_i = c_i - e_i, \tag{34}$$

$$F_i = \begin{cases} (c_i - e_i)/(N_i^{\text{gen}} - e_i), & N_i^{\text{gen}} > e_i, \\ 1, & N_i^{\text{gen}} = e_i, \end{cases} \tag{35}$$

$$U_i = \mathbf{1}[c_i = N_i^{\text{gen}}]. \tag{36}$$

Here, e_i is the final position of the teacher segment, N_i^{gen} is the complete generated response length, and c_i is the final supervised position. H_i is the supervised suffix length, F_i is its fraction of the generated suffix, and U_i is the no-cut indicator. When a cutoff is selected, both detection windows and all later tokens remain in the generated response but are excluded from distillation. We average F_i and U_i separately over intervened responses. The ratio F_i describes supervision coverage, not a saving in generation cost.

C.5 Conditional outcomes and online gain

To distinguish the effects of teacher intervention on originally incorrect and correct responses, we score final answers before and after intervention within the same aggregation window and compute

the rescue and harm rates:

$$R_{\text{rescue}} = \frac{n_{WC}}{N_{I,W}}, \tag{37}$$

$$R_{\text{harm}} = \frac{n_{CW}}{N_{I,C}}. \tag{38}$$

Here, W and C denote incorrect and correct final answers, respectively, and n_{WC} and n_{CW} count transitions from incorrect to correct and from correct to incorrect. $N_{I,W}$ and $N_{I,C}$ count intervened responses whose original answers are incorrect and correct, respectively. These are the respective denominators of R_{rescue} and R_{harm}; the original accuracy of the intervened subset is $N_{I,C}/(N_{I,W} + N_{I,C})$. The curves aggregate counts over a trailing 32-step window before computing the ratios. A ratio is undefined when its denominator is zero.

To measure the net accuracy change from rescue and harm, we compute online gain for both the intervened subset and all rollouts:

$$\Delta A_I = 100\frac{n_{WC} - n_{CW}}{N_I}, \tag{39}$$

$$\Delta A_{\text{all}} = 100\frac{n_{WC} - n_{CW}}{B}. \tag{40}$$

Here, $N_I = N_{I,W} + N_{I,C}$ is the number of intervened responses, and B is the total number of rollouts in the current aggregation window. ΔA_I and ΔA_{all} are the respective gains, both expressed in %. Responses without intervention remain unchanged and contribute zero to the net change. Thus, $\Delta A_{\text{all}} = (N_I/B)\Delta A_I$, where N_I/B is the intervention rate.

To show how intervention gains change during training, the Online gain scatter plot in Figure 8 uses ΔA_{all}. Each point represents a training step, colored by step; arrows connect the centers of consecutive 32-step phases. Online metrics assess changes in the current training responses, while student-only accuracy after training is evaluated separately by Equation 22.

For the four groups in Figure 7(c), each response is grouped by whether the continuation's supervision was cut short and whether its 128-token local divergence mean decreased. Correctness is evaluated on the complete generated response, including any tokens excluded from distillation. This local grouping statistic differs from the paired first-16-token diagnostic in Equation 32. Each group's online gain uses that group's response count as its denominator. The groups contain 12,949 intervened responses with complete grouping information; 237 interventions lack a valid local comparison and are excluded. Error bars are 95% percentile intervals from 10,000 bootstrap resamples of training steps, retaining all responses within each sampled step (random seed 20260921). These intervals describe variation in the recorded training outcomes, not variability across independent training runs or uncertainty in the final benchmark scores.

C.6 TRAJECTORY CONSTRUCTION ALGORITHM AND IMPLEMENTATION DETAILS

To construct a mixed trajectory for distillation, Algorithm 1 combines divergence monitoring, teacher intervention, and supervision selection into the complete procedure for the main setting $M = 1$. Its inputs follow the definitions in Section 4.

Algorithm 1 Adaptive teacher intervention and supervision selection in ATI-OPD ($M = 1$)

Require: prompt x; student $\pi_{\theta_{\text{old}}}$; teacher π_T
Require: lengths L_S, L_T, L_F; thresholds τ, η_s; budget $N_{\max}$
Ensure: mixed response $\widetilde{y}$, token owners o, final supervised position c

sample a raw response $y^{(0)} \sim \pi_{\theta_{\text{old}}}(\cdot \mid x)$ under $N_{\max}$
score eligible segments; $\mathcal{K} \leftarrow \{k : G(\text{seg}_k^{\text{stu}}) > \tau\}$
if $\mathcal{K} = \emptyset$ **then**
    return $y^{(0)}, S^{|y^{(0)}|}, |y^{(0)}|$
end if
$k^\star \leftarrow \min \mathcal{K}$; $p \leftarrow p_{<k^\star}$; discard the original suffix
sample $z \sim \pi_T(\cdot \mid x, p)$ with at most $\min(L_T, N_{\max} - |p|)$ tokens
$u \leftarrow \emptyset$
if $p \oplus z$ has not ended and $|p \oplus z| < N_{\max}$ **then**
    sample a complete suffix $u \sim \pi_{\theta_{\text{old}}}(\cdot \mid x, p, z)$ under the remaining budget
end if
$\widetilde{y} \leftarrow p \oplus z \oplus u$; $o \leftarrow S^{|p|} \oplus T^{|z|} \oplus S^{|u|}$; $c \leftarrow |\widetilde{y}|$
score $\widetilde{y}$; split u into windows of at most L_F tokens
for successive windows $j = 1, 2, \ldots$ in u **do**
    compute Q_j if window j is valid; let b_j be its first response position
    if $j \geq 2$, both adjacent windows are valid, and $Q_{j-1} > \eta_s,\ Q_j > \eta_s$ **then**
        $c \leftarrow b_{j-1} - 1$
        break ▷ stop the cutoff search; keep the generated response
    end if
end for
return $\widetilde{y}, o, c$

Here, S^n and T^n denote n identical token-owner markers, and $\oplus$ denotes sequence concatenation. The algorithm returns the complete mixed response $\widetilde{y}$, token-owner sequence $o = (o_t)$, and final supervised position c. To apply OPD only at selected positions, we use the following mask over the complete response:

$$m_t^{\text{OPD}} = m_t^{\text{base}} \mathbf{1}[t \leq c] \left(\mathbf{1}[o_t = S] + \iota_T \mathbf{1}[o_t = T]\right). \tag{41}$$

Here, m_t^{OPD} marks positions used for distillation, and $m_t^{\text{base}} \in \{0, 1\}$ marks valid response tokens. $o_t \in \{S, T\}$ records whether a token was generated by the student or teacher. $\iota_T \in \{0, 1\}$ controls whether teacher tokens receive OPD supervision.

Valid segments and windows. To make divergence statistics reflect the student's generation state, monitoring uses only student response tokens with valid log-probabilities, excluding padding and teacher tokens. For numerical stability, we clip the log-probability difference and the resulting K3 score:

$$\bar{\delta}_t = \text{clip}(\delta_t, -20, 20), \tag{42}$$

$$\bar{g}_t = \text{clip}\big(\exp(\bar{\delta}_t) - \bar{\delta}_t - 1, -10, 10\big). \tag{43}$$

Here, $\bar{\delta}_t$ is the bounded log-probability difference and $\bar{g}_t$ is the clipped sampled-token K3 score. Segment triggers, continuation windows, and local intervention diagnostics all use this clipped score; the absolute-gap diagnostic is defined separately in Equation 15. Both the sums and denominators of segment and window means count only valid tokens, while starting positions retain their original response offsets. Suffix windows must meet the minimum valid-token count in Table 2; both adjacent windows in Equation 6 must be valid and exceed the threshold. A single threshold crossing does not trigger the cutoff. The reference pool stores valid window means from original student responses with correct final answers, retains the latest values up to the capacity in Table 2, and sets the threshold to the delayed 90th percentile. Each correct response contributes at most eight window means, sampled at evenly spaced indices when more are available. The threshold is τ until the pool contains 128 windows. Updates from a completed rollout batch affect only the following batch, so the current response cannot change its own reference threshold.

Budgets and termination. To satisfy both the local intervention limit and the total response budget, the teacher generates at most $\min(L_T, N_{\max} - |p_{<k}|)$ tokens per intervention. If the teacher segment ends the response, the student suffix is $u^{\mathrm{stu}} = \emptyset$; otherwise, the student continues until the response ends or the total budget is exhausted. All valid generated positions remain eligible for supervision if no cutoff is selected, and the original student rollout is used if no intervention occurs. The adaptive supervision implementation uses one intervention. In Full with multiple interventions, the next trigger is sought only after the latest teacher segment, leaving earlier teacher segments intact.

Supervision masking and compute cost. The implementation applies the cutoff only to the distillation loss mask, preserving the complete response and its token-owner sequence. Equation 6 uses one-based token positions to set $c = b_{j^\star} - 1$; the implementation admits zero-based positions strictly below c. Both windows that exceed the threshold and all subsequent positions are excluded from the loss, as in Equation 41. The complete new continuation has already incurred generation and scoring costs. We therefore do not interpret the supervision cutoff or the excluded-token fraction as decoding savings.