NOT EVERY TOKEN IS WORTH DISTILLING: SELECTIVE SUPERVISION FOR DIRECT-OPD

Anonymous authors
Paper under double-blind review

ABSTRACT

Direct On-Policy Distillation (Direct-OPD) transfers reinforcement-learning-induced policy improvements from a small model to a larger student by using the token-level log-ratio between post-RL and pre-RL checkpoints as dense supervision on the student's own rollouts. This transfer rewards the policy shift at every state, yet the log-ratio measures only relative change: it can stay fixed even as the probability mass that both checkpoints assign to the student's candidate tokens vanishes. Through an exact construction, we show that the Direct-OPD reward and its update can remain unchanged while the Jensen–Shannon divergence (JSD) and both KL directions between the checkpoints vanish with this mass, and we note that a small JSD bounds how much the teacher's behavior changed. Motivated by this analysis, we propose Selective Supervision for Direct-OPD (S^2D-OPD), which ranks student-sampled states by their teacher–reference JSD and masks Direct-OPD supervision at low-divergence states, retaining only the top 10% of states per response. Across two teacher pairs and four student models ranging from 1.7B to 8B parameters, S^2D-OPD improves held-out accuracy over dense Direct-OPD on AIME and HMMT benchmarks in seven of eight settings and matches it in the eighth, without extra forward passes. Our code is available at `https://anonymous.4open.science/r/S2D-OPD-8868`.

"No Free Lunch for Supervised Machine Learning." — David H. Wolpert

1 INTRODUCTION

Guided by scaling laws (Kaplan et al., 2020; Hoffmann et al., 2022; Pearce & Song, 2024), recent large language models have continued to scale up pre-training (DeepSeek-AI et al., 2026; Team et al., 2026; GLM-5-Team et al., 2026), which broadens their knowledge and latent capabilities. Post-training, most prominently reinforcement learning (RL) (Yue et al., 2025; Shao et al., 2024; Zhao et al., 2025), is then needed to elicit and refine these capabilities. As model size grows, RL demands more rollout generation, training compute, memory, and infrastructure (Wu et al., 2025), and stable optimization remains difficult to achieve (Wang et al., 2026a). Consequently, the models with the greatest post-training potential are also the most expensive to improve through RL.

One line of work makes large-scale RL more stable through improved training algorithms (Yu et al., 2025; Zheng et al., 2025; MiniMax et al., 2025; Ma et al., 2025; Hou et al., 2026) and more efficient through training and inference systems (Sheng et al., 2025; Fu et al., 2025; Kwon et al., 2023; Narayanan et al., 2021; Zhu et al., 2025). These advances make RL on large models more practical, but its cost still grows with model size. Another line of work uses on-policy distillation (OPD) (Agarwal et al., 2024; Gu et al., 2024; Lu & Lab, 2025), in which a stronger teacher provides token-level supervision on states sampled by the student. OPD transfers capabilities efficiently to a smaller student (Li et al., 2026b; Fu et al., 2026c), but it relies on a teacher stronger than the student, which is unavailable when the target is already the strongest model. Together, these approaches leave a challenge open: how to improve a large model without paying the cost of RL at its scale.

Direct-OPD (Feng et al., 2026) and Proxy-OPD (Fu et al., 2026a) recently proposed a weak-to-strong route around this challenge: running RL on a small model and transferring the result to a larger one. Both methods extract the RL-induced policy shift as the token-level log-ratio between the post-RL teacher and its pre-RL reference, and use it as an OPD reward on states sampled by

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.

the larger student. This turns outcome-level rewards into dense token-level supervision for the large model, without running RL at its scale or requiring a teacher stronger than the student.

However, this apparent free lunch leaves unexamined *whether every token-level reward reflects a meaningful change in the teacher's behavior*. Because the log-ratio reward measures only relative change, it can stay fixed even as the probability mass that both checkpoints assign to the student's top candidates vanishes. Direct-OPD can therefore reward a state where both checkpoints barely support these candidates as strongly as one where RL clearly changed the teacher's behavior. This raises the question that motivates this work: *is every state's policy shift worth distilling for free?*

We make this probability-mass mismatch exact in Sec. 4.1: as the mass that both checkpoints assign to the student's top candidates vanishes, the Direct-OPD reward and gradient can stay fixed, whereas the Jensen–Shannon divergence (JSD) and both directions of KL divergence between the checkpoints vanish with it. Because JSD accounts for the probability mass that the log-ratio ignores, we propose **Selective Supervision for Direct-OPD** (S^2D-OPD), which ranks student-sampled states by teacher–reference JSD and masks Direct-OPD supervision at low-divergence states. Empirically, the answer to our question is no: retaining only the top 10% of states per response, S^2D-OPD improves held-out accuracy over dense Direct-OPD in seven of eight teacher–student settings and matches it in the eighth (mean gain 0.95 points; 95% CI 0.40–1.54), without extra forward passes.

In summary, our contributions are threefold:

- **A Probability-Mass Mismatch in Direct-OPD.** Through an exact construction, we show that the log-ratio reward and its local gradient on the student can remain fixed while the probability mass behind the policy shift vanishes, and with it the teacher–reference JSD and both KL directions. We further show that JSD bounds how much the teacher's behavior can change at every state, which motivates selecting states by divergence rather than by the reward itself.
- **Stable and Effective Selective Transfer.** We propose S^2D-OPD, which masks Direct-OPD supervision at low-divergence states during policy transfer. Across four student scales and two teacher pairs, it improves held-out accuracy over Direct-OPD in seven of eight settings and yields smoother late-stage validation curves under the JustRL teacher pair.
- **Understanding the Gains from Selective Transfer.** Within a fixed teacher–student setting, performance broadly rises with JSD percentile: the top bin outperforms a uniformly sampled 10% subset, whereas the lowest bin degrades the student below its initialization and eventually collapses. Across teacher pairs, the pair with lower overall JSD benefits more from masking, consistent with low-divergence filtering being a source of the improvement over dense Direct-OPD.

2 Related Work

On-Policy Distillation. OPD trains a student on prefixes sampled from its own policy, using the teacher's next-token distributions as dense supervision at every position (Agarwal et al., 2024; Gu et al., 2024; Lu & Lab, 2025). Subsequent work refines it through alternative objectives (Jin et al., 2026; Jia et al., 2026), stabilization strategies (Li et al., 2026b; Fu et al., 2026b), and privileged-context self-distillation (Zhao et al., 2026; Pan et al., 2026). Despite their differences, these methods primarily learn from the teacher's policy itself, limiting transfer to improvements already present in the teacher; recent work therefore targets the teacher's policy shift instead. ExOPD (Yang et al., 2026) extrapolates the teacher's improvement over its reference model to construct a target beyond the teacher. Direct-OPD (Feng et al., 2026) and Proxy-OPD (Fu et al., 2026a) instead transfer the log-ratio between a reward-optimized checkpoint and its pre-RL reference, analogous to the logit shifts induced by fine-tuning studied in CMC (Wu et al., 2024). This targets the RL-induced policy shift and can provide useful supervision even for students already stronger than the post-RL teacher. However, token-level log-ratios capture relative changes but are insensitive to the absolute probability mass supporting these changes. We examine this probability-mass mismatch in Direct-OPD and use teacher–reference divergence to select supervision positions.

Token Selection in Policy Distillation. Selective distillation asks which positions are worth training on, and existing criteria differ mainly in which distributions they read. Some read the student alone, prioritizing positions where it is uncertain (Tavor et al., 2026; Ko et al., 2026); others the teacher alone, weighting by its confidence or local margin (Jin et al., 2026; Zhou et al., 2026); a third group

compares the two, emphasizing teacher–student disagreement, which TIP (Xu et al., 2026) organizes through an entropy–divergence taxonomy and TA-OPD (Wang et al., 2026b) restricts to the student's predictive support. We consider the policy change from a pre-RL reference to a post-RL teacher at student-visited prefixes. A related approach, OPD^2 (Heo et al., 2026), gates sampled-token delta updates by sign agreement between the centered teacher–base and teacher–student log-ratios. Our selection criterion instead ranks positions by teacher–reference JSD, accounting for the probability mass underlying the policy shift while retaining the Direct-OPD update at selected positions.

3 PRELIMINARIES

Setting. We consider three policies: a pre-RL reference π_{ref}, a post-RL teacher π_{T} obtained from π_{ref} by outcome-based RL such as GRPO (Shao et al., 2024), and a larger student π_θ initialized at π_{stu}. In our experiments, both teacher-side checkpoints are publicly released (Sec. 5.1), so we run no RL ourselves. Given a prompt $\boldsymbol{x} \sim \mathcal{D}$ and a response $\boldsymbol{y} = (y_1, \ldots, y_{|\boldsymbol{y}|})$ sampled from the student, position t has state $\boldsymbol{s}_t = (\boldsymbol{x}, \boldsymbol{y}_{<t})$.

Direct-OPD treats the teacher's RL-induced policy shift as a dense reward for the student. For any token v, the reward at state $\boldsymbol{s}_t$ is the teacher–reference log-ratio

$$\Delta_t(v \mid \boldsymbol{s}_t) = \log \frac{\pi_{\mathrm{T}}(v \mid \boldsymbol{s}_t)}{\pi_{\mathrm{ref}}(v \mid \boldsymbol{s}_t)}, \tag{1}$$

which is positive where RL increased the probability of v and negative where it decreased it.

Direct-OPD maximizes this reward on states visited by the student, with KL regularization:

$$J_{\text{Direct-OPD}}(\theta) = \mathbb{E}_{\boldsymbol{x}\sim\mathcal{D},\,(\boldsymbol{s}_t,y_t)\sim\pi_\theta}\left[\Delta_t(y_t \mid \boldsymbol{s}_t) - \alpha D_{\mathrm{KL}}\left[\pi_\theta(\cdot \mid \boldsymbol{s}_t) \| \pi_{\mathrm{stu}}(\cdot \mid \boldsymbol{s}_t)\right]\right]. \tag{2}$$

Here, $\alpha > 0$ controls KL regularization, and the expectation covers all valid response positions.

Top-K implementation. In practice, Direct-OPD evaluates the reward on the student's top-K candidates at each state rather than only on the sampled token:

$$\mathcal{V}_K(\boldsymbol{s}_t) = \mathrm{TopK}(\pi_\theta(\cdot \mid \boldsymbol{s}_t), K), \quad \bar{p}_t(v) = \frac{\pi_\theta(v \mid \boldsymbol{s}_t)}{\sum_{u\in\mathcal{V}_K(\boldsymbol{s}_t)} \pi_\theta(u \mid \boldsymbol{s}_t)}, \tag{3}$$

where $\mathrm{TopK}(\cdot, K)$ returns the set of K tokens with the largest probabilities. Each candidate receives the teacher–reference reward $\Delta_t(v \mid \boldsymbol{s}_t)$, weighted by its renormalized student probability $\bar{p}_t(v)$. With the state and candidate set held fixed, the implemented local reward-gradient contribution is:

$$g_t^R = \sum_{v\in\mathcal{V}_K(\boldsymbol{s}_t)} \mathrm{sg}\left[\bar{p}_t(v) \log \frac{\pi_{\mathrm{T}}(v \mid \boldsymbol{s}_t)}{\pi_{\mathrm{ref}}(v \mid \boldsymbol{s}_t)}\right] \nabla_\theta \log \pi_\theta(v \mid \boldsymbol{s}_t), \tag{4}$$

where sg denotes stop-gradient. The student-anchor KL term supplies a separate regularization gradient. Appendix A describes its implementation and adaptive coefficient.

4 METHOD

Direct-OPD applies its log-ratio reward at every valid position of a student response. Sec. 4.1 shows that this reward can ignore the probability mass behind the teacher's policy shift, and Sec. 4.2 introduces S^{2}D-OPD, which selects states by teacher–reference divergence.

4.1 THEORETICAL MOTIVATION: PROBABILITY-MASS MISMATCH

At each state, Direct-OPD weights the teacher–reference rewards $\Delta_t(v \mid \boldsymbol{s}_t)$ by the student's renormalized probabilities $\bar{p}_t(v)$ over its top-K candidates in Eq. 4. These weights reflect the student's preferences, whereas the rewards encode relative changes in the teacher's policy. Neither depends on how much probability mass the teacher and reference place on these candidates: rescaling both by a common factor leaves every log-ratio unchanged. The following exact construction makes this precise: the teacher–reference divergence can vanish while the Direct-OPD update stays fixed.

An exact construction. Fix a state $\boldsymbol{s}$ and a student checkpoint θ_0 with top-K candidate set $\mathcal{V}_K = \mathcal{V}_K(\boldsymbol{s})$. Let $\mathbf{t} \neq \mathbf{q}$ be strictly positive probability vectors on $\mathcal{V}_K$, let $\mathbf{b}$ be a strictly positive probability vector on its complement, and for $0 < \epsilon < 1$ define the teacher and reference:

$$P_\epsilon(v) = \begin{cases} \epsilon t_v, & v \in \mathcal{V}_K, \\ (1-\epsilon) b_v, & v \notin \mathcal{V}_K, \end{cases} \qquad Q_\epsilon(v) = \begin{cases} \epsilon q_v, & v \in \mathcal{V}_K, \\ (1-\epsilon) b_v, & v \notin \mathcal{V}_K. \end{cases} \tag{5}$$

Here, ϵ is the probability mass that each checkpoint assigns to the student's candidates, while the student, candidate set, and conditional distributions stay fixed. Then, for every candidate $v \in \mathcal{V}_K$,

$$\Delta_\epsilon(v \mid \boldsymbol{s}) = \log \frac{t_v}{q_v}, \qquad D_{\mathrm{JS}}(P_\epsilon, Q_\epsilon) = \epsilon \, D_{\mathrm{JS}}(\mathbf{t}, \mathbf{q}), \tag{6}$$

and both directions of KL scale with ϵ in the same way. As $\epsilon \to 0$, all three divergences vanish, whereas every candidate reward, and hence the Direct-OPD update of Eq. 4, stays fixed and nonzero. The construction thus isolates a single degree of freedom, the probability mass behind the policy shift, to which the Direct-OPD update is insensitive but the divergences are not. We state this result formally in Prop. 1 and prove it in App. B, including why the update is nonzero.

Divergence bounds the behavioral change. The construction shows that the Direct-OPD reward can ignore divergence; conversely, divergence bounds how much the teacher's behavior can change. For any two distributions P and Q on a finite set and any event A,

$$|P(A) - Q(A)| \;\le\; D_{\mathrm{TV}}(P, Q) \;\le\; \sqrt{2 \, D_{\mathrm{JS}}(P, Q)}, \tag{7}$$

where $D_{\mathrm{TV}}(P, Q) = \max_A |P(A) - Q(A)|$ is the total variation distance and D_{JS} is measured in nats. The second inequality follows from Pinsker's inequality applied to each term of $D_{\mathrm{JS}}(P, Q) = \frac{1}{2} D_{\mathrm{KL}}(P \| M) + \frac{1}{2} D_{\mathrm{KL}}(Q \| M)$ with $M = (P + Q)/2$, since $D_{\mathrm{TV}}(P, M) = D_{\mathrm{TV}}(Q, M) = \frac{1}{2} D_{\mathrm{TV}}(P, Q)$. Unlike the construction, this bound holds at every state: wherever the teacher–reference JSD is small, the teacher assigns nearly the same probability as the reference to every token and every set of tokens. Together, the two results characterize what low-divergence masking removes: states at which the teacher's behavior provably changed little, yet at which the Direct-OPD update can be as large as anywhere else. They do not show that removing these states improves transfer, which Sec. 5.3 tests by training on JSD percentile bins.

4.2 Divergence-Guided State Selection

Sec. 4.1 shows that the Direct-OPD reward is insensitive to the probability mass behind a policy shift, whereas the teacher–reference JSD bounds how much the teacher's behavior changed at a state. S^{2}D-OPD therefore scores each student-sampled state by this divergence and retains Direct-OPD supervision only at the highest-scoring states within each response (Fig. 1). We use JSD as the default score because, unlike KL, it is symmetric, bounded by $\log 2$, and finite when either checkpoint assigns a token zero probability; the procedure applies unchanged to other divergences.

Scoring on the student's candidates. Rather than the full vocabulary, we evaluate the divergence on the top-K candidate set $\mathcal{V}_K(\boldsymbol{s}_t)$, where Direct-OPD already evaluates both checkpoints, and add a residual token v_{o} that collects all remaining probability mass. For $M \in \{\mathrm{T}, \mathrm{ref}\}$, define

$$\widetilde{\pi}_M(v \mid \boldsymbol{s}_t) = \begin{cases} \pi_M(v \mid \boldsymbol{s}_t), & v \in \mathcal{V}_K(\boldsymbol{s}_t), \\ 1 - \displaystyle\sum_{u \in \mathcal{V}_K(\boldsymbol{s}_t)} \pi_M(u \mid \boldsymbol{s}_t), & v = v_{\mathrm{o}}. \end{cases} \tag{8}$$

This keeps each candidate's probability and the total residual mass without renormalization. By the data-processing inequality, the coarsening cannot increase JSD:

$$\tilde{d}_t \triangleq D_{\mathrm{JS}}(\widetilde{\pi}_{\mathrm{T}}(\cdot \mid \boldsymbol{s}_t) \,\|\, \widetilde{\pi}_{\mathrm{ref}}(\cdot \mid \boldsymbol{s}_t)) \le D_{\mathrm{JS}}(\pi_{\mathrm{T}}(\cdot \mid \boldsymbol{s}_t) \,\|\, \pi_{\mathrm{ref}}(\cdot \mid \boldsymbol{s}_t)) \,. \tag{9}$$

Fig. 1 illustrates two ways in which the score can be small. Either both checkpoints place little mass on the student's candidates, so that both compressed distributions concentrate on v_{o}, which is the regime of the construction in Sec. 4.1; or both place substantial mass on the candidates and agree on how it is distributed. The score treats these cases alike, and it need not separate them: because Eq. 7

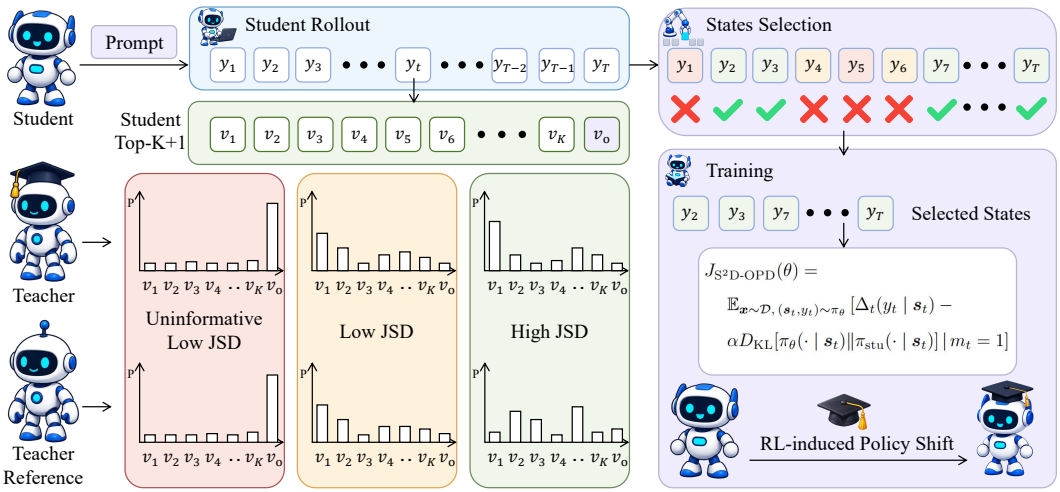


Figure 1: Overview of S^{2}D-OPD. At each student-sampled state, teacher–reference JSD is computed over the student's top-K candidates, with the remaining probability mass grouped into a residual token v_o. The highest-JSD states within each response are retained for the Direct-OPD objective.

holds for any pair of distributions, including the compressed ones, a small score implies in either case that the teacher barely changed its behavior on the student's candidates. Conversely, the score is large only when the checkpoints disagree on the individual candidate probabilities or on the total mass assigned to the candidate set. The score thus measures behavioral change at the resolution of the student's candidates: differences among tail tokens outside the candidate set do not affect it.

Per-response selection. We select states within each response, so that rollouts with different overall divergence levels all contribute supervision. For valid response positions $\mathcal{T}$ and a common retention ratio $\rho \in (0, 1]$, we retain the k states with the largest scores:

$$k = \max\{1, \lceil \rho |\mathcal{T}| \rceil\}, \qquad \mathcal{I} = \mathrm{TopK}_{t \in \mathcal{T}}(\tilde{d}_t, k), \tag{10}$$

where retaining at least one state ensures that every rollout is represented in the objective. Selecting within each response also makes each mask independent of the rest of the batch.

Masked objective. We apply the mask $m_t = \mathbf{1}\{t \in \mathcal{I}\}$ to both terms of the Direct-OPD objective:

$$J_{\text{S}^2\text{D-OPD}}(\theta) = \mathbb{E}_{\boldsymbol{x}\sim\mathcal{D},\,(\boldsymbol{s}_t,y_t)\sim\pi_\theta}\left[\Delta_t(y_t \mid \boldsymbol{s}_t) - \alpha D_{\mathrm{KL}}[\pi_\theta(\cdot \mid \boldsymbol{s}_t)\|\pi_{\mathrm{stu}}(\cdot \mid \boldsymbol{s}_t)] \,|\, m_t = 1\right]. \tag{11}$$

In practice, we hold the selection fixed during optimization and average over all retained states in the batch, which gives the update

$$\mathbf{g}^{\text{S}^2\text{D-OPD}} = \frac{1}{|\mathcal{I}_\mathcal{B}|} \sum_{t \in \mathcal{I}_\mathcal{B}} \Big[g_t^R - \alpha \nabla_\theta D_{\mathrm{KL}}[\pi_\theta(\cdot \mid \boldsymbol{s}_t)\|\pi_{\mathrm{stu}}(\cdot \mid \boldsymbol{s}_t)] \Big], \tag{12}$$

where $\mathcal{I}_\mathcal{B}$ collects the retained states of all responses in the batch and g_t^R is the reward gradient of Eq. 4. Masked states thus receive neither the reward nor the student anchor, and $\rho = 1$ recovers Direct-OPD. The adaptive KL coefficient is still computed from all valid positions; App. A gives the remaining optimization details. Because the score reuses the teacher and reference probabilities that Direct-OPD already computes, S^{2}D-OPD requires no extra forward passes.

5 Experiments

Sec. 4.1 shows which states low-divergence masking removes, but not whether removing them improves transfer. We test this through three research questions (further analyses in App. C, D, E):

RQ1: Does masking low-JSD states improve transfer? (Sec. 5.2)

RQ2: How does JSD relate to supervision utility and selective masking gains? (Sec. 5.3)

RQ3: Are the gains robust across divergence measures and selection schemes? (Sec. 5.4)

Table 1: Main results (Avg@32) across four student models and two teacher–reference pairs, comparing dense Direct-OPD with S^{2}D-OPD using top-10% JSD selection. Checkpoints are selected on AIME24/25 and evaluated on held-out AIME26 and HMMT (Nov. 2025 and Feb. 2026). **Test Avg.** is the average accuracy over 93 held-out problems; parentheses denote gains over the initial student.

	R1-Distill-1.5B → JustRL-1.5B					Nemotron-1.5B → QuestA-1.5B				
	Checkpoint Selection		Held-out Evaluation			Checkpoint Selection		Held-out Evaluation		
Student / Method	**AIME24**	**AIME25**	**AIME26**	**HMMT**	**Test Avg.**	**AIME24**	**AIME25**	**AIME26**	**HMMT**	**Test Avg.**
Qwen3-1.7B	49.3	36.7	37.7	28.2	31.3	49.3	36.7	37.7	28.2	31.3
+ Direct-OPD	59.7	42.9	45.4	32.0	36.4 (+5.1)	58.6	43.4	46.8	31.6	36.5 (+5.2)
+ S^{2}D-OPD	**61.4**	**44.5**	**47.4**	**32.9**	**37.6** (+6.3)	**59.5**	**43.8**	**48.1**	**33.3**	**38.1** (+6.8)
Qwen3-4B	73.2	65.2	64.8	45.9	52.0	73.2	65.2	64.8	45.9	52.0
+ Direct-OPD	**78.1**	70.3	66.8	**47.1**	53.5 (+1.5)	77.6	**70.4**	70.4	47.1	54.6 (+2.6)
+ S^{2}D-OPD	77.0	**70.9**	**68.7**	46.5	**53.7** (+1.7)	**78.0**	68.3	**70.7**	**49.3**	**56.2** (+4.2)
Qwen3-8B	77.3	66.0	67.5	49.7	55.4	77.3	66.0	67.5	49.7	55.4
+ Direct-OPD	77.5	72.2	69.8	**50.2**	56.5 (+1.1)	75.9	**71.7**	69.7	48.6	55.4 (+0.0)
+ S^{2}D-OPD	**78.0**	**73.4**	**70.7**	50.0	**56.7** (+1.3)	**78.1**	**71.7**	**70.9**	**50.6**	**57.1** (+1.7)
R1-Distill-7B	56.7	40.5	48.2	29.5	35.6	56.7	40.5	48.2	29.5	35.6
+ Direct-OPD	63.6	45.7	**56.8**	32.8	**40.5** (+4.9)	**60.7**	**42.3**	49.5	31.9	37.6 (+2.0)
+ S^{2}D-OPD	**64.6**	**47.1**	55.0	**33.6**	**40.5** (+4.9)	60.6	**42.3**	**52.4**	**32.0**	**38.6** (+3.0)

5.1 Experimental Setup

We evaluate S^{2}D-OPD with two teacher pairs, R1-Distill-1.5B → JustRL-1.5B (He et al., 2025a) and Nemotron-1.5B → QuestA-1.5B (Li et al., 2026a), and transfer each policy shift to four students: Qwen3-1.7B, Qwen3-4B, Qwen3-8B[1], and R1-Distill-7B[2]. All students are trained on Skywork-OR1-RL-Data (He et al., 2025b). Following the observation of Meng et al. (2026) that RL-induced policy shifts are sparse, with large divergence concentrated in a small fraction of tokens, we retain the top 10% of states per response ($\rho = 0.1$) by default. We select checkpoints on AIME 2024 and AIME 2025 and evaluate the selected checkpoint on three held-out benchmarks: AIME 2026, HMMT November 2025, and HMMT February 2026. All three postdate the release of every student model and are therefore absent from its training data, including undisclosed post-training data. We report Avg@32 on all benchmarks, define validation accuracy as the mean over AIME 2024 and AIME 2025, and give further training hyperparameter details in App. F.

5.2 Masking Low-Divergence States Improves Transfer

Retaining 10% of states improves held-out transfer. Tab. 1 compares S^{2}D-OPD with dense Direct-OPD across four students and two teacher pairs. Retaining 10% of states achieves comparable validation accuracy and improves the held-out Test Avg. in seven of eight settings, with a tie in the eighth. We assess statistical significance over 93 held-out problems, first averaging each problem's 32 responses and then averaging its paired differences across the eight settings. Using a paired problem bootstrap and an exact one-sided sign-flip test, we find that S^{2}D-OPD improves mean held-out accuracy from 46.36% to 47.31%, a gain of 0.95 points (95% CI: 0.40–1.54; $p = 6.4 \times 10^{-4}$).

Selective supervision preserves early learning while improving late-stage stability. Fig. 2 shows the validation trajectories under the JustRL teacher pair. Despite masking low-divergence states, S^{2}D-OPD keeps pace with dense Direct-OPD during early training, matching or exceeding its accuracy over the first 60 steps for all four students. Discarding 90% of states thus does not slow early learning, suggesting that the supervision driving early improvement is concentrated in the retained high-divergence subset. Later in training, S^{2}D-OPD maintains higher accuracy for all four students and is more stable on Qwen3-4B and Qwen3-8B, where dense Direct-OPD repeatedly falls back from its peaks and, on Qwen3-4B, even drops below the initial student by step 300. These results answer RQ1: masking low-JSD states improves held-out transfer without slowing early learning.

[1] https://huggingface.co/Qwen/Qwen3-{1.7,4,8}B

[2] https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-7B

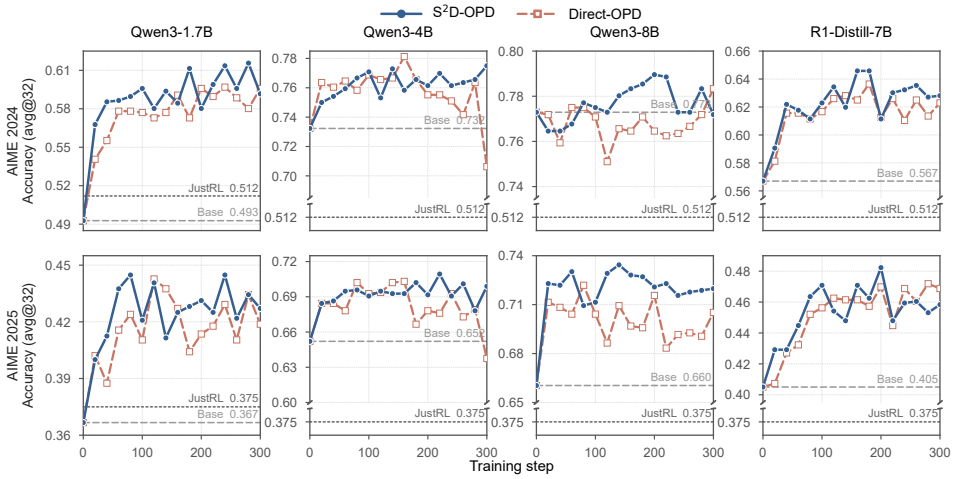
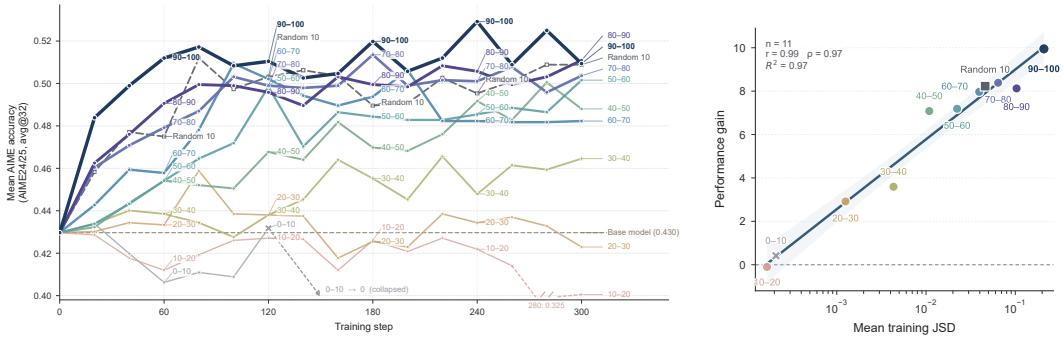


Figure 2: Validation accuracy (Avg@32) under the JustRL teacher pair across four students. Dashed lines mark the initial student; dotted lines mark the JustRL-1.5B teacher.



(a) Performance across JSD percentile bins.

(b) Mean JSD tracks transfer gains.

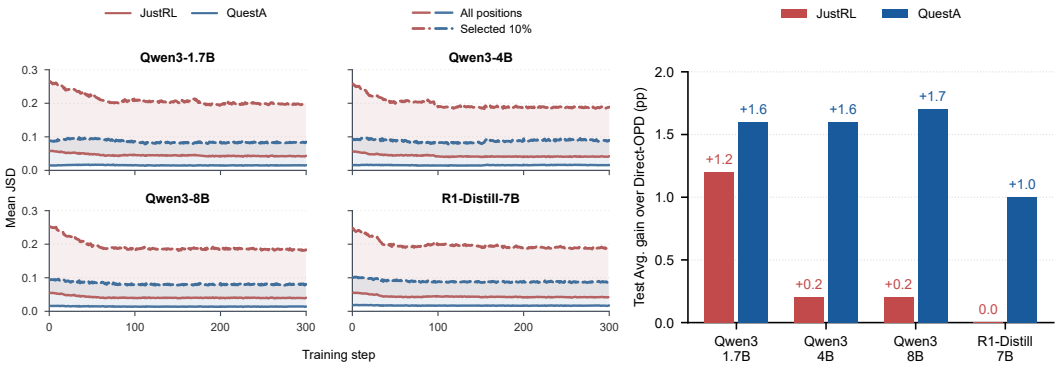
Figure 3: Which states are retained, not how many, determines transfer (Qwen3-1.7B, JustRL teacher pair). **(a)** Validation accuracy when training on each JSD percentile bin or on a random 10% of states. **(b)** Gain in peak validation accuracy over the initial student versus the mean training JSD of the retained states; the line is a log-linear fit.

5.3 Low-Divergence Supervision Is Redundant and Harmful

Sec. 4.1 shows that low-JSD states are those at which the teacher's behavior changed little, yet the Direct-OPD update there can be as large as anywhere else. To test what such supervision contributes, we fix the JustRL teacher pair, the Qwen3-1.7B student, the training data, and the hyperparameters, and vary only which states enter the Direct-OPD objective. Within each response, we rank states by JSD, split them into ten equal-sized percentile bins (0–10, …, 90–100), and train a separate student on each bin. As a control, we train on a uniformly sampled 10% of states, whose update is an unbiased estimate of the dense Direct-OPD update (Fig. 3a).

Dense supervision is largely redundant. The random 10% control reaches a peak validation accuracy of 51.2, on par with dense Direct-OPD (51.3; Tab. 1). Discarding 90% of states at random thus loses little, supporting the view of Fu et al. (2026c) that OPD is "data-overfed but algorithm-starved."

Which states are retained determines transfer. At the same 10% budget, performance rises broadly with JSD percentile, and the top bin reaches 52.9, above both the random control and dense Direct-OPD. The effect is graded: across all eleven subsets, the gain in peak validation accuracy over the initial student is approximately linear in the logarithm of the subset's mean training JSD ($R^2 = 0.97$; Fig. 3b). The random control lies on the same line, so within this setting the mean divergence of the retained states predicts transfer whether they are selected by rank or at random.



(a) All-position and selected-position JSD.

(b) Held-out gains over Direct-OPD.

Figure 4: Teacher-pair comparison with top-10% retention across four students. **(a)** Mean teacher–reference JSD over all valid response states (solid) and retained states (dashed) during S^{2}D-OPD training. **(b)** Held-out Test Avg. gains of S^{2}D-OPD over dense Direct-OPD under each teacher pair, in percentage points.

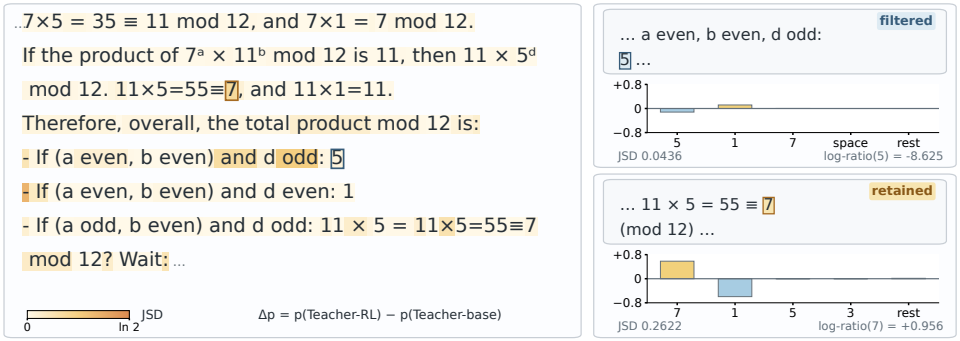
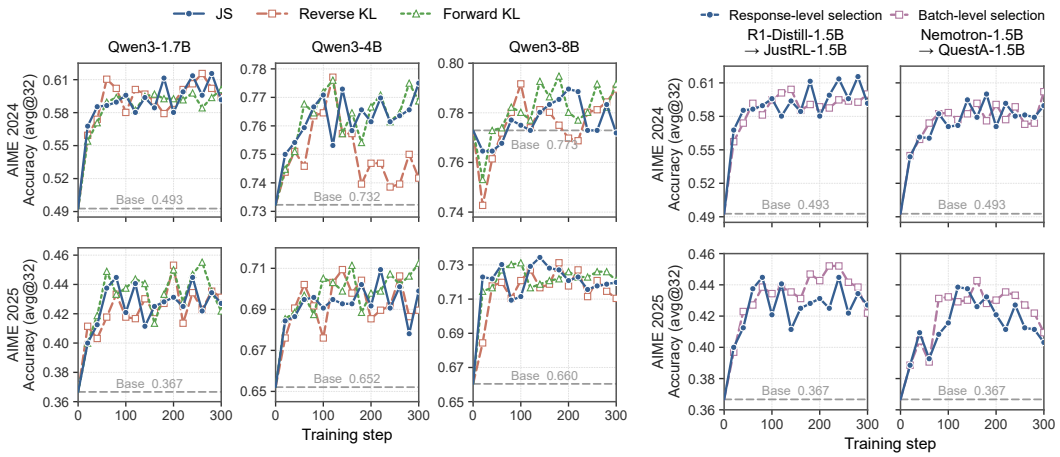


Figure 5: **JSD and token log-ratio rank states differently.** The state with the largest sampled-token log-ratio in the response is masked, while a larger redistribution of probability mass is retained.

Low-divergence supervision is harmful, not merely weak. Trained in isolation, the three lowest bins end below the initial student, and the two lowest collapse. This matches the regime identified in Sec. 4.1, where the Direct-OPD update persists although the teacher barely changed. Top-K overlap suggests a mechanism (App. C): the top bin moves the student toward the teacher while preserving its overlap with the reference, whereas lower bins move it away from the reference without a commensurate gain in teacher alignment.

Divergence and reward magnitude rank states differently. Fig. 5 shows a state from an initial Qwen3-1.7B rollout under the JustRL pair. RL lowers the probability of the sampled token 5 from 0.12 to 2×10^{-5}, giving this state the largest Direct-OPD reward magnitude among the response's 8,192 states, although the teacher's distribution moves by only 0.12 in total variation. Its JSD (0.044) falls below the retention cutoff (0.067), so S^{2}D-OPD masks it while retaining a state with a smaller log-ratio but a larger redistribution of mass (JSD 0.262).

Masking helps more under the lower-divergence teacher pair. Mean JSD is higher for JustRL than for QuestA over both all states and the retained top 10% (Fig. 4a), yet the held-out gains of masking over dense Direct-OPD are larger under QuestA: 1.0–1.7 points versus 0.0–1.2 for JustRL across the four students (Fig. 4b). If the benefit came from concentrating supervision on high absolute divergence, this ordering would be reversed; together with the harm of low-JSD bins, it suggests that masking helps mainly by removing supervision that undermines learning. These results answer RQ2: JSD ranks supervision utility within a teacher–student setting, and masking pays off by filtering low-divergence supervision rather than by maximizing the divergence of what remains.



(a) Divergence measure.

(b) Selection scope.

Figure 6: Robustness of selective masking across design choices. **(a)** Top-10% selection by JSD, reverse KL, and forward KL on Qwen3-1.7B, 4B, and 8B under the JustRL teacher pair. **(b)** Response-level versus batch-level selection on Qwen3-1.7B under both teacher pairs.

5.4 Selective Masking Is Robust Across Design Choices

We vary three design choices while keeping the rest of S^{2}D-OPD fixed: the divergence used for scoring, the scope over which states are ranked, and the retention rate (Fig. 6).

Divergence measure. The construction in Sec. 4.1 shows that both KL directions vanish with the candidate mass just as JSD does, so either should flag the same low-mass states. Under the JustRL pair, selection by forward KL, $D_{\mathrm{KL}}(\pi_{\mathrm{T}}\|\pi_{\mathrm{ref}})$, or reverse KL, $D_{\mathrm{KL}}(\pi_{\mathrm{ref}}\|\pi_{\mathrm{T}})$, attains higher mean validation accuracy than dense Direct-OPD from step 160 onward on all three Qwen3 students, and avoids the late decline of dense Direct-OPD on Qwen3-4B (final accuracy 71.6–74.1 versus 67.2). Forward KL closely tracks JSD at all three scales, whereas reverse KL is weaker on Qwen3-4B.

Selection scope. Ranking states across the whole batch, so that a response may retain none, matches response-level selection on Qwen3-1.7B under both teacher pairs (peak validation accuracy 52.3 versus 52.9 under JustRL and 51.7 versus 51.7 under QuestA). The per-response rule of Sec. 4.2 is thus not needed for accuracy; we keep it so that each mask is independent of other responses.

Retention rate. Varying ρ from 5% to 20% on Qwen3-1.7B changes peak validation accuracy by less than one point (52.2–53.0), and every rate stays above dense Direct-OPD (51.3) and the random 10% control (51.2; App. D). These results answer RQ3: within the tested range, the gains do not hinge on the divergence measure, selection scope, or retention rate.

6 Conclusion

Direct-OPD transfers the policy shift that RL induces in a small teacher to a larger student by rewarding every student-sampled state with the teacher–reference log-ratio. We showed that this reward ignores the probability mass behind the shift: it can stay fixed while the mass on the student's candidates, and with it the teacher–reference divergence, vanishes. Building on this observation, we introduced S^{2}D-OPD, which retained Direct-OPD supervision only at the highest-JSD states of each response and required no extra forward passes. Retaining 10% of states improved held-out accuracy over dense Direct-OPD in seven of eight teacher–student settings and tied in the eighth. Controlled analyses further showed that dense supervision was largely redundant, that transfer depended on which states were retained rather than how many, and that low-divergence supervision could harm the student when trained on alone; the gains persisted across divergence measures, selection scopes, and retention rates. These findings suggest that weak-to-strong policy transfer depends not only on what is transferred but also on where it is applied. Extending state selection to adaptive retention rates and to domains beyond mathematical reasoning is a natural next step.

AI USE STATEMENT

We used generative AI tools solely to improve the readability and clarity of the manuscript, including language editing and suggestions on presentation. All AI-assisted revisions were reviewed by the authors, who take full responsibility for the final content of this work.

ETHICS STATEMENT

We adhere to the ICLR Code of Ethics. This work uses only publicly released models and datasets under their licenses and involves no human subjects or personal data. Because S^2D-OPD selects states by teacher–reference divergence rather than correctness or safety, it may transfer undesirable behaviors to the student, and our math evaluations do not establish the safety of the resulting models.

REPRODUCIBILITY STATEMENT

To facilitate reproducibility, we describe the selective supervision procedure in Section 4 and report the teacher–reference pairs, datasets, training configurations, and evaluation protocols in Section 5.1 and Appendix F. We provide our code at `https://anonymous.4open.science/r/S2D-OPD-8868`.

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*, 2024. URL `https://openreview.net/forum?id=3zKtaqxLhW`.

DeepSeek-AI, Anyi Xu, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, Chenchen Ling, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chengyu Hou, Chenhao Xu, Chenze Shao, Chong Ruan, Conner Sun, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Donghao Li, Dongjie Ji, Erhang Li, Fang Wei, Fangyun Lin, Fangzhou Yuan, Feiyu Xia, Fucong Dai, Guangbo Hao, Guanting Chen, Guoai Cao, Guolai Meng, Guowei Li, Han Yu, Han Zhang, Hanwei Xu, Hao Li, Haofen Liang, Haoling Zhang, Haoming Luo, Haoran Wei, Haotian Yuan, Haowei Zhang, Haowen Luo, Haoyu Chen, Haozhe Ji, Hengqing Zhang, Honghui Ding, Hongxuan Tang, Huanqi Cao, Huazuo Gao, Hui Qu, Hui Zeng, J Yang, JQ Zhu, Jia Luo, Jia Song, Jia Yu, Jialiang Huang, Jialu Cai, Jian Liang, Jiangting Zhou, Jiasheng Ye, Jiashi Li, Jiaxin Xu, Jiewen Hu, Jieyu Yang, Jin Chen, Jin Yan, Jingchang Chen, Jingli Zhou, Jingting Xiang, Jingyang Yuan, Jingyuan Cheng, Jingzi Zhou, Jinhua Zhu, Jiping Yu, Joseph Sun, Jun Ran, Junguang Jiang, Junjie Qiu, Junlong Li, Junmin Zheng, Junxiao Song, Kai Dong, Kaige Gao, Kang Guan, Kexing Zhou, Kezhao Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Wang, Leyi Xia, Li Zhang, Liang Zhao, Lihua Guo, Lingxiao Luo, Linwang Ma, Linyan Zhu, Litong Wang, Liyu Cai, Liyue Zhang, Longhao Chen, MS Di, MY Xu, Max Mei, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Mingxu Zhou, Minmin Han, Ning Wang, Panpan Huang, Panpan Wang, Peixin Cong, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qingyang Li, Qinyu Chen, Qiushi Du, Qiwei Jiang, Rui Tian, Ruifan Xu, Ruijie Lu, Ruiling Xu, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runqian Chen, Runqiu Yin, Runxin Xu, Ruomeng Shen, Ruoyu Zhang, Ruyi Chen, SH Liu, Shanghao Lu, Shangmian Sun, Shangyan Zhou, Shanhuang Chen, Shaofei Cai, Shaoheng Nie, Shaoqing Wu, Shaoyuan Chen, Shengding Hu, Shengyu Liu, Shiqiang Hu, Shirong Ma, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, Shuying Yu, Songyang Zhou, Tao Ni, Tao Yun, Tian Jin, Tian Pei, Tian Ye, Tianle Lin, Tianran Ji, Tianyi Cui, Tianyuan Yue, Tingting Yu, Tun Wang, W Zhang, WL Xiao, Wangding Zeng, Wei An, Weilin Zhao, Wen Liu, Wenfeng Liang, Wenjie Pang, Wenjing Luo, Wenjing Yao, Wenjun Gao, Wenkai Yang, Wenlve Huang, Wenqing Hou, Wentao Zhang, Wenting Ma, Xi Gao, Xiang He, Xiangwen Wang, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaokang Zhang, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingchen Liu, Xingkai Yu, Xingyou Li, Xinyu Yang, Xinyu Zhang, Xu Chen, Xuanyu Wang, Xuecheng Su, Xueyin Chen, Xuheng Lin, Xuwei Fu, YC Yan, YQ Wang, YW Ma, Yanfeng Luo, Yang Zhang, Yanhong Xu, Yanru Ma, Yanwen Huang, Yao Li, Yao Li, Yao Xu, Yao Zhao, Yaofeng Sun, Yaohui Wang,

Yi Qian, Yi Shao, Yi Yu, Yichao Zhang, Yifan Ding, Yifan Shi, Yijia Wu, Yiliang Xiong, Yiling Ma, Ying He, Ying Tang, Ying Zhou, Yingjia Luo, Yinmin Zhong, Yishi Piao, Yisong Wang, Yixiang Zhang, Yixiao Chen, Yixuan Tan, Yixuan Wei, Yiyang Ma, Yiyuan Liu, Yonglun Yang, Yongqiang Guo, Yongtong Wu, Yu Wu, YuKun Li, Yuan Cheng, Yuan Ou, Yuanfan Xu, Yuanhao Li, Yuduan Wang, Yuehan Yang, Yuer Xu, Yuhan Wu, Yuhao Meng, Yuheng Zou, Yukun Zha, Yunfan Xiong, Yupeng Chen, Yuping Lin, Yuqian Cao, Yuqian Wang, Yushun Zhang, Yuting Yan, Yutong Lin, Yuxian Gu, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuxuan Zhou, Yuyang Zhou, Yuzhen Huang, ZF Wu, Zehao Wang, Zehua Zhao, Zehui Ren, Zekai Zhang, Zhangli Sha, Zhe Fu, Zhe Ju, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zheren Gao, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhixian Huang, Zhixuan Chen, Zhiyu Wu, Zhizhou Ren, Zhongyu Wu, Zhuoshu Li, Zhuping Zhang, Zian Xu, Zihao Wang, Zihua Qu, Zihui Gu, Zijia Zhu, Zilin Li, Zipeng Zhang, Ziwei Xie, Ziyi Gao, Ziyi Wan, Zizheng Pan, and Zongqing Yao. Deepseek-v4: Towards highly efficient million-token context intelligence, 2026. URL https://arxiv.org/abs/2606.19348.

Shiyuan Feng, Huan ang Gao, Haohan Chi, Hanlin Wu, Zhilong Zhang, Zheng Jiang, Bingxiang He, Wei-Ying Ma, Ya-Qin Zhang, and Hao Zhou. Weak-to-strong generalization via direct on-policy distillation, 2026. URL https://arxiv.org/abs/2607.05394.

Daocheng Fu, Rong Wu, Yu Yang, Jianbiao Mei, Licheng Wen, Pinlong Cai, Xuemeng Yang, Yong Liu, Botian Shi, and Yu Qiao. Proxy opd: On-policy distillation with transferable relative proxy update, 2026a. URL https://arxiv.org/abs/2607.11505.

Wei Fu, Jiaxuan Gao, Xujie Shen, Chen Zhu, Zhiyu Mei, Chuyi He, Shusheng Xu, Guo Wei, Jun Mei, WANG JIASHU, Tongkai Yang, Binhang Yuan, and Yi Wu. AREAL: A large-scale asynchronous reinforcement learning system for language reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=X9diEuva9R.

Yuqian Fu, Haohuan Huang, Kaiwen Jiang, Jiacai Liu, Zhuo Jiang, Yuanheng Zhu, and Dongbin Zhao. Revisiting on-policy distillation: Empirical failure modes and simple fixes, 2026b. URL https://arxiv.org/abs/2603.25562.

Zixuan Fu, Bingxiang He, Yuxin Zuo, Haohuan Huang, Jinqian Zhang, Ruhang Xiao, Cheng Qian, Qinyu Luo, Huan ang Gao, Yudong Wang, Zhiyuan Liu, Ning Ding, and Chaojun Xiao. Rethinking on-policy distillation of large language models ii: One training example, 2026c. URL https://arxiv.org/abs/2609.04172.

GLM-5-Team, Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, Bin Chen, Da Yin, Chendi Ge, Chenghua Huang, Chengxing Xie, Chenzheng Zhu, Congfeng Yin, Cunxiang Wang, Gengzheng Pan, Hao Zeng, Haoke Zhang, Haoran Wang, Huilong Chen, Jiajie Zhang, Jian Jiao, Jiaqi Guo, Jingsen Wang, Jingzhao Du, Jinzhu Wu, Kedong Wang, Lei Li, Lin Fan, Lucen Zhong, Mingdao Liu, Mingming Zhao, Pengfan Du, Qian Dong, Rui Lu, Shuang-Li, Shulin Cao, Song Liu, Ting Jiang, Xiaodong Chen, Xiaohan Zhang, Xuancheng Huang, Xuezhen Dong, Yabo Xu, Yao Wei, Yifan An, Yilin Niu, Yitong Zhu, Yuanhao Wen, Yukuo Cen, Yushi Bai, Zhongpei Qiao, Zihan Wang, Zikang Wang, Zilin Zhu, Ziqiang Liu, Zixuan Li, Bojie Wang, Bosi Wen, Can Huang, Changpeng Cai, Chao Yu, Chen Li, Chengwei Hu, Chenhui Zhang, Dan Zhang, Daoyan Lin, Dayong Yang, Di Wang, Ding Ai, Erle Zhu, Fangzhou Yi, Feiyu Chen, Guohong Wen, Hailong Sun, Haisha Zhao, Haiyi Hu, Hanchen Zhang, Hanrui Liu, Hanyu Zhang, Hao Peng, Hao Tai, Haobo Zhang, He Liu, Hongwei Wang, Hongxi Yan, Hongyu Ge, Huan Liu, Huanpeng Chu, Jia'ni Zhao, Jiachen Wang, Jiajing Zhao, Jiamin Ren, Jiapeng Wang, Jiaxin Zhang, Jiayi Gui, Jiayue Zhao, Jijie Li, Jing An, Jing Li, Jingwei Yuan, Jinhua Du, Jinxin Liu, Junkai Zhi, Junwen Duan, Kaiyue Zhou, Kangjian Wei, Ke Wang, Keyun Luo, Laiqiang Zhang, Leigang Sha, Liang Xu, Lindong Wu, Lintao Ding, Lu Chen, Minghao Li, Nianyi Lin, Pan Ta, Qiang Zou, Rongjun Song, Ruiqi Yang, Shangqing Tu, Shangtong Yang, Shaoxiang Wu, Shengyan Zhang, Shijie Li, Shuang Li, Shuyi Fan, Wei Qin, Wei Tian, Weining Zhang, Wenbo Yu, Wenjie Liang, Xiang Kuang, Xiangmeng Cheng, Xiangyang Li, Xiaoquan Yan, Xiaowei Hu, Xiaoying Ling, Xing Fan, Xingye Xia, Xinyuan Zhang, Xinze Zhang, Xirui Pan, Xu Zou, Xunkai Zhang, Yadi Liu, Yandong Wu, Yanfu Li, Yidong Wang, Yifan Zhu, Yijun Tan, Yilin Zhou, Yiming Pan, Ying Zhang, Yinpei Su, Yipeng Geng, Yong Yan, Yonglin Tan, Yuean Bi, Yuhan Shen, Yuhao Yang, Yujiang Li, Yunan Liu, Yunqing Wang, Yuntao Li, Yurong Wu, Yutao Zhang, Yuxi Duan, Yuxuan Zhang, Zezhen

Liu, Zhengtao Jiang, Zhenhe Yan, Zheyu Zhang, Zhixiang Wei, Zhuo Chen, Zhuoer Feng, Zijun Yao, Ziwei Chai, Ziyuan Wang, Zuzhou Zhang, Bin Xu, Minlie Huang, Hongning Wang, Juanzi Li, Yuxiao Dong, and Jie Tang. Glm-5: from vibe coding to agentic engineering, 2026. URL https://arxiv.org/abs/2602.15763.

Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. MiniLLM: Knowledge distillation of large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=5h0qf7IBZZ.

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

Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen, Siyuan Li, Liang Zeng, Tianwen Wei, Cheng Cheng, Bo An, Yang Liu, and Yahui Zhou. Skywork open reasoner 1 technical report, 2025b. URL https://arxiv.org/abs/2505.22312.

Byeongho Heo, Jaehui Hwang, Sangdoo Yun, and Dongyoon Han. On-policy delta distillation, 2026. URL https://arxiv.org/abs/2607.15161.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack W. Rae, and Laurent Sifre. Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.

Zhenyu Hou, Yujiang Li, Jie Tang, and Yuxiao Dong. Single-rollout asynchronous optimization for agentic reinforcement learning, 2026. URL https://arxiv.org/abs/2607.07508.

Nan Jia, Haojin Yang, Xing Ma, Jiesong Lian, Shuailiang Zhang, Weipeng Zhang, Ke Zeng, Xunliang Cai, and Zequn Sun. Asymmetric on-policy distillation: Bridging exploitation and imitation at the token level, 2026. URL https://arxiv.org/abs/2605.06387.

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. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=J5i09faOOf.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models, 2020. URL https://arxiv.org/abs/2001.08361.

Jongwoo Ko, Sara Abdali, Young Jin Kim, Tianyi Chen, and Pashmina Cameron. Scaling reasoning efficiently via relaxed on-policy distillation, 2026. URL https://arxiv.org/abs/2603.11137.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, SOSP '23, pp. 611–626, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400702297. doi: 10.1145/3600006.3613165. URL https://doi.org/10.1145/3600006.3613165.

Jiazheng Li, Hongzhou Lin, Hong Lu, Kaiyue Wen, Zaiwen Yang, Jiaxuan Gao, Yi Wu, and Jingzhao Zhang. Questa: Expanding reasoning capacity in LLMs via question augmentation. In *The Fourteenth International Conference on Learning Representations*, 2026a. URL https://openreview.net/forum?id=3MifB0f7qR.

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

Kevin Lu and Thinking Machines Lab. On-policy distillation. *Thinking Machines Lab: Connectionism*, 2025. doi: 10.64434/tml.20251026. https://thinkingmachines.ai/blog/on-policy-distillation.

Wenhan Ma, Hailin Zhang, Liang Zhao, Yifan Song, Yudong Wang, Zhifang Sui, and Fuli Luo. Stabilizing moe reinforcement learning by aligning training and inference routers, 2025. URL https://arxiv.org/abs/2510.11370.

Haoming Meng, Kexin Huang, Shaohang Wei, Chiyu Ma, Shuo Yang, Xue Wang, Guoyin Wang, Bolin Ding, and Jingren Zhou. Sparse but critical: A token-level analysis of distributional shifts in RLVR fine-tuning of LLMs. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=8vWIXno8LW.

MiniMax, Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan, Changqing Yu, Chao Wang, Cheng Zhu, Chengjun Xiao, Chengyu Du, Chi Zhang, Chu Qiao, Chunhao Zhang, Chunhui Du, Congchao Guo, Da Chen, Deming Ding, Dianjun Sun, Dong Li, Enwei Jiao, Haigang Zhou, Haimo Zhang, Han Ding, Haohai Sun, Haoyu Feng, Huaiguang Cai, Haichao Zhu, Jian Sun, Jiaqi Zhuang, Jiaren Cai, Jiayuan Song, Jin Zhu, Jingyang Li, Jinhao Tian, Jinli Liu, Junhao Xu, Junjie Yan, Junteng Liu, Junxian He, Kaiyi Feng, Ke Yang, Kecheng Xiao, Le Han, Leyang Wang, Lianfei Yu, Liheng Feng, Lin Li, Lin Zheng, Linge Du, Lingyu Yang, Lunbin Zeng, Minghui Yu, Mingliang Tao, Mingyuan Chi, Mozhi Zhang, Mujie Lin, Nan Hu, Nongyu Di, Peng Gao, Pengfei Li, Pengyu Zhao, Qibing Ren, Qidi Xu, Qile Li, Qin Wang, Rong Tian, Ruitao Leng, Shaoxiang Chen, Shaoyu Chen, Shengmin Shi, Shitong Weng, Shuchang Guan, Shuqi Yu, Sichen Li, Songquan Zhu, Tengfei Li, Tianchi Cai, Tianrun Liang, Weiyu Cheng, Weize Kong, Wenkai Li, Xiancai Chen, Xiangjun Song, Xiao Luo, Xiao Su, Xiaobo Li, Xiaodong Han, Xinzhu Hou, Xuan Lu, Xun Zou, Xuyang Shen, Yan Gong, Yan Ma, Yang Wang, Yiqi Shi, Yiran Zhong, Yonghong Duan, Yongxiang Fu, Yongyi Hu, Yu Gao, Yuanxiang Fan, Yufeng Yang, Yuhao Li, Yulin Hu, Yunan Huang, Yunji Li, Yunzhi Xu, Yuxin Mao, Yuxuan Shi, Yuze Wenren, Zehan Li, Zelin Li, Zhanxu Tian, Zhengmao Zhu, Zhenhua Fan, Zhenzhen Wu, Zhichao Xu, Zhihang Yu, Zhiheng Lyu, Zhuo Jiang, Zibo Gao, Zijia Wu, Zijian Song, and Zijun Sun. Minimax-m1: Scaling test-time compute efficiently with lightning attention, 2025. URL https://arxiv.org/abs/2506.13585.

Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. Efficient large-scale language model training on gpu clusters using megatron-lm. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '21, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384421. doi: 10.1145/3458817.3476209. URL https://doi.org/10.1145/3458817.3476209.

Leyi Pan, Shuchang Tao, Yunpeng Zhai, Lingzhe Zhang, Zhaoyang Liu, Bolin Ding, Aiwei Liu, and Lijie Wen. Rlcsd: Reinforcement learning with contrastive on-policy self-distillation, 2026. URL https://arxiv.org/abs/2606.11709.

Tim Pearce and Jinyeop Song. Reconciling kaplan and chinchilla scaling laws. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=NLoaLyuUUF. Reproducibility Certification.

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

Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, pp. 1279–1297, New

York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400711961. doi: 10.1145/3689031.3696075. URL https://doi.org/10.1145/3689031.3696075.

Almog Tavor, Itay Ebenspanger, Neil Cnaan, and Mor Geva. Rethinking selective knowledge distillation, 2026. URL https://arxiv.org/abs/2602.01395.

Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, S. H. Cai, Yuan Cao, Y. Charles, H. S. Che, Cheng Chen, Guanduo Chen, Huarong Chen, Jia Chen, Jiahao Chen, Jianlong Chen, Jun Chen, Kefan Chen, Liang Chen, Ruijue Chen, Xinhao Chen, Yanru Chen, Yanxu Chen, Yicun Chen, Yimin Chen, Yingjiang Chen, Yuankun Chen, Yujie Chen, Yutian Chen, Zhirong Chen, Ziwei Chen, Dazhi Cheng, Minghan Chu, Jialei Cui, Jiaqi Deng, Muxi Diao, Hao Ding, Mengfan Dong, Mengnan Dong, Yuxin Dong, Yuhao Dong, Angang Du, Chenzhuang Du, Dikang Du, Lingxiao Du, Yulun Du, Yu Fan, Shengjun Fang, Qiulin Feng, Yichen Feng, Garimugai Fu, Kelin Fu, Hongcheng Gao, Tong Gao, Yuyao Ge, Shangyi Geng, Chengyang Gong, Xiaochen Gong, Zhuoma Gongque, Qizheng Gu, Xinran Gu, Yicheng Gu, Longyu Guan, Yuanying Guo, Xiaoru Hao, Weiran He, Wenyang He, Yunjia He, Chao Hong, Hao Hu, Jiaxi Hu, Yangyang Hu, Zhenxing Hu, Ke Huang, Ruiyuan Huang, Weixiao Huang, Zhiqi Huang, Tao Jiang, Zhejun Jiang, Xinyi Jin, Yu Jing, Guokun Lai, Aidi Li, C. Li, Cheng Li, Fang Li, Guanghe Li, Guanyu Li, Haitao Li, Haoyang Li, Jia Li, Jingwei Li, Junxiong Li, Lincan Li, Mo Li, Weihong Li, Wentao Li, Xinhang Li, Xinhao Li, Yang Li, Yanhao Li, Yiwei Li, Yuxiao Li, Zhaowei Li, Zheming Li, Weilong Liao, Jiawei Lin, Xiaohan Lin, Zhishan Lin, Zichao Lin, Cheng Liu, Chenyu Liu, Hongzhang Liu, Liang Liu, Shaowei Liu, Shudong Liu, Shuran Liu, Tianwei Liu, Tianyu Liu, Weizhou Liu, Xiangyan Liu, Yangyang Liu, Yanming Liu, Yibo Liu, Yuanxin Liu, Yue Liu, Zhengying Liu, Zhongnuo Liu, Enzhe Lu, Haoyu Lu, Zhiyuan Lu, Junyu Luo, Tongxu Luo, Yashuo Luo, Long Ma, Yingwei Ma, Shaoguang Mao, Yuan Mei, Xin Men, Fanqing Meng, Zhiyong Meng, Yibo Miao, Minqing Ni, Kun Ouyang, Siyuan Pan, Bo Pang, Yuchao Qian, Ruoyu Qin, Zeyu Qin, Jiezhong Qiu, Bowen Qu, Zeyu Shang, Youbo Shao, Tianxiao Shen, Zhennan Shen, Juanfeng Shi, Lidong Shi, Shengyuan Shi, Feifan Song, Pengwei Song, Tianhui Song, Xiaoxi Song, Hongjin Su, Jianlin Su, Zhaochen Su, Lin Sui, Jinsong Sun, Junyao Sun, Tongyu Sun, Flood Sung, Yunpeng Tai, Chuning Tang, Heyi Tang, Xiaojuan Tang, Zhengyang Tang, Jiawen Tao, Shiyuan Teng, Chaoran Tian, Pengfei Tian, Ao Wang, Bowen Wang, Chensi Wang, Chuang Wang, Congcong Wang, Dingkun Wang, Dinglu Wang, Dongliang Wang, Feng Wang, Hailong Wang, Haiming Wang, Hengzhi Wang, Huaqing Wang, Hui Wang, Jiahao Wang, Jinhong Wang, Jiuzheng Wang, Kaixin Wang, Linian Wang, Qibin Wang, Shengjie Wang, Shuyi Wang, Si Wang, Wei Wang, Xiaochen Wang, Xinyuan Wang, Yao Wang, Yejie Wang, Yipu Wang, Yiqin Wang, Yucheng Wang, Yuzhi Wang, Zhaoji Wang, Zhaowei Wang, Zhengtao Wang, Zhexu Wang, Zihan Wang, Zizhe Wang, Chu Wei, Ming Wei, Chuan Wen, Zichen Wen, Chengjie Wu, Haoning Wu, Junyan Wu, Rucong Wu, Wenhao Wu, Yuefeng Wu, Yuhao Wu, Yuxin Wu, Zijian Wu, Chenjun Xiao, Jin Xie, Xiaotong Xie, Yuchong Xie, Yifei Xin, Bowei Xing, Boyu Xu, Jianfan Xu, Jing Xu, Jinjing Xu, L. H. Xu, Lin Xu, Suting Xu, Weixin Xu, Xinbo Xu, Xinran Xu, Yangchuan Xu, Yichang Xu, Yuemeng Xu, Zelai Xu, Ziyao Xu, Junjie Yan, Yuzi Yan, Guangyao Yang, Hao Yang, Junwei Yang, Kai Yang, Ningyuan Yang, Ruihan Yang, Xiaofei Yang, Xinlong Yang, Ying Yang, Yi Yang, Yi Yang, Zhen Yang, Zhilin Yang, Zonghan Yang, Haotian Yao, Dan Ye, Wenjie Ye, Zhuorui Ye, Bohong Yin, Chengzhen Yu, Longhui Yu, Tao Yu, Tianxiang Yu, Enming Yuan, Mengjie Yuan, Xiaokun Yuan, Yang Yue, Weihao Zeng, Dunyuan Zha, Haobing Zhan, Dehao Zhang, Hao Zhang, Jin Zhang, Puqi Zhang, Qiao Zhang, Rui Zhang, Xiaobin Zhang, Y. Zhang, Yadong Zhang, Yangkun Zhang, Yichi Zhang, Yizhi Zhang, Yongting Zhang, Yu Zhang, Yushun Zhang, Yutao Zhang, Yutong Zhang, Zheng Zhang, Chenguang Zhao, Feifan Zhao, Jinxiang Zhao, Shuai Zhao, Xiangyu Zhao, Yikai Zhao, Zijia Zhao, Huabin Zheng, Ruihan Zheng, Shaojie Zheng, Tengyang Zheng, Junfeng Zhong, Longguang Zhong, Weiming Zhong, M. Zhou, Runjie Zhou, Xinyu Zhou, Zaida Zhou, Jinguo Zhu, Liya Zhu, Xinhao Zhu, Yuxuan Zhu, Zhen Zhu, Jingze Zhuang, Weiyu Zhuang, Ying Zou, and Xinxing Zu. Kimi k2.5: Visual agentic intelligence, 2026. URL https://arxiv.org/abs/2602.02276.

Xiaoxuan Wang, Han Zhang, Haixin Wang, Yidan Shi, Ruoyan Li, Kaiqiao Han, Chenyi Tong, Haoran Deng, Alexander K Taylor, Renliang Sun, Yanqiao Zhu, Jason Cong, Yizhou Sun, and Wei Wang. ARLArena: A unified framework for stable agentic reinforcement learning. In *Forty-third International Conference on Machine Learning*, 2026a. URL https://openreview.net/forum?id=90kxFi9VGP.

Yuanyi Wang, Su Lu, Yanggan Gu, Pengkai Wang, Yifan Yang, Zhaoyi Yan, Congkai Xie, Jianmin Wu, and Hongxia Yang. Not all disagreement is learnable: Token teachability in on-policy distillation, 2026b. URL https://arxiv.org/abs/2605.26844.

Bo Wu, Sid Wang, Yunhao Tang, Jia Ding, Eryk Helenowski, Liang Tan, Tengyu Xu, Tushar Gowda, Zhengxing Chen, Chen Zhu, Xiaocheng Tang, Yundi Qian, Beibei Zhu, and Rui Hou. Llamarl: A distributed asynchronous reinforcement learning framework for efficient large-scale llm training, 2025. URL https://arxiv.org/abs/2505.24034.

Jiayi Wu, Hao Sun, Hengyi Cai, Lixin Su, Shuaiqiang Wang, Dawei Yin, Xiang Li, and Ming Gao. Cross-model control: Improving multiple large language models in one-time training. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=YPqHSTSoFs.

Yuanda Xu, Hejian Sang, Zhengze Zhou, Ran He, Zhipeng Wang, and Alborz Geramifard. Tip: Token importance in on-policy distillation, 2026. URL https://arxiv.org/abs/2604.14084.

Wenkai Yang, Weijie Liu, Ruobing Xie, Kai Yang, Saiyong Yang, and Yankai Lin. Learning beyond teacher: Generalized on-policy distillation with reward extrapolation, 2026. URL https://arxiv.org/abs/2602.12125.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan 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, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL https://arxiv.org/abs/2503.14476.

Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao Huang. Does reinforcement learning really incentivize reasoning capacity in LLMs beyond the base model? In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=4OsgYD7em5.

Rosie Zhao, Alexandru Meterez, Sham M. Kakade, Cengiz Pehlevan, Samy Jelassi, and Eran Malach. Echo chamber: RL post-training amplifies behaviors learned in pretraining. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=dp4KWuSDzj.

Siyan Zhao, Zhihui Xie, Mengchen Liu, Jing Huang, Guan Pang, Feiyu Chen, and Aditya Grover. Self-distilled reasoner: On-policy self-distillation for large language models, 2026. URL https://arxiv.org/abs/2601.18734.

Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. Group sequence policy optimization, 2025. URL https://arxiv.org/abs/2507.18071.

Yuhang Zhou, Lizhu Zhang, Yifan Wu, Mingyi Wang, Bo Peng, Jiayi Liu, Xiangjun Fan, and Zhuokai Zhao. Sage-opd: Selective agent-guided intervention for multi-turn on-policy distillation, 2026. URL https://arxiv.org/abs/2606.19659.

Zilin Zhu, Chengxing Xie, Xin Lv, and slime Contributors. slime: An llm post-training framework for rl scaling. https://github.com/THUDM/slime, 2025. GitHub repository. Corresponding author: Xin Lv.

A Optimization Details

This section details the adaptive KL control inherited from Direct-OPD and the selective position-wise aggregation used in S^2D-OPD.

A.1 Adaptive KL Control

The KL coefficient controls the strength of the student anchor relative to the policy-shift reward. We update this coefficient using Direct-OPD's sign-based controller:

$$\alpha_{m+1} = \mathrm{clip}(\alpha_m \left[1 + \epsilon \ \mathrm{sgn}(\bar{r}_m)\right], \alpha_{\min}, \alpha_{\max}), \tag{13}$$

where $\bar{r}_m$ denotes the mean student-weighted policy shift $\bar{p}_t(v)\Delta_t(v \mid \boldsymbol{s}_t)$ over valid response positions and their top-K candidates at iteration m. The controller increases α when the mean weighted policy shift is positive and decreases it when negative, subject to the bounds. We use $\alpha_0 = 2.5$, $\epsilon = 0.01$, and $[\alpha_{\min}, \alpha_{\max}] = [0.5, 2.5]$. The updated coefficient α_{m+1} is used in the subsequent actor update.

A.2 Detailed Loss Aggregation

Direct-OPD averages the local update over all valid response positions. S^{2}D-OPD keeps the local update unchanged but restricts this average to selected positions. Let $\mathcal{I}$ collect the positions selected independently within each response by Eq. 10 across the batch. Giving each retained position equal weight yields the local ascent direction

$$g_\theta^{\mathrm{S^2D\text{-}OPD}} = \frac{1}{|\mathcal{I}|} \sum_{t \in \mathcal{I}} \left[\sum_{v \in \mathcal{V}_K(\boldsymbol{s}_t)} \mathrm{sg}[\bar{p}_t(v)\Delta_t(v \mid \boldsymbol{s}_t)] \, \nabla_\theta \log \pi_\theta(v \mid \boldsymbol{s}_t) - \alpha_{m+1} \nabla_\theta \widehat{d}_t(\theta) \right]. \tag{14}$$

Here, sg denotes stop-gradient, and $\widehat{d}_t(\theta)$ is the per-position KL penalty estimator toward the initial student π_{stu}, computed using `verl`'s `low_var_kl` implementation. Both terms are applied to the same selected response positions. The sampled prefixes, candidate sets, selection mask, and KL coefficient are held fixed during differentiation.

B Proof of the Mass-Invariance Construction

Proposition 1 (Mass invariance of the Direct-OPD update)**.** *Fix a state $\boldsymbol{s}$ and a student checkpoint θ_0, and write $p(v) = \pi_{\theta_0}(v \mid \boldsymbol{s}) > 0$ for all $v \in \mathcal{V}$. Let $\mathcal{V}_K(\boldsymbol{s}) \subsetneq \mathcal{V}$ be the top-K candidate set, with $K \geq 2$, and define $\bar{p}(v) = p(v)/\sum_{u \in \mathcal{V}_K(\boldsymbol{s})} p(u)$ for $v \in \mathcal{V}_K(\boldsymbol{s})$. Let $\mathbf{t} = (t_v)_{v \in \mathcal{V}_K(\boldsymbol{s})}$ and $\mathbf{q} = (q_v)_{v \in \mathcal{V}_K(\boldsymbol{s})}$ be distinct, strictly positive probability vectors on the candidate set, and let $\mathbf{b} = (b_v)_{v \in \mathcal{V} \setminus \mathcal{V}_K(\boldsymbol{s})}$ be a common strictly positive probability vector on its complement. For $0 < \epsilon < 1$, define the teacher and reference distributions:*

$$P_\epsilon(v) = \begin{cases} \epsilon t_v, & v \in \mathcal{V}_K(\boldsymbol{s}), \\ (1-\epsilon) b_v, & v \notin \mathcal{V}_K(\boldsymbol{s}), \end{cases} \qquad Q_\epsilon(v) = \begin{cases} \epsilon q_v, & v \in \mathcal{V}_K(\boldsymbol{s}), \\ (1-\epsilon) b_v, & v \notin \mathcal{V}_K(\boldsymbol{s}). \end{cases} \tag{15}$$

Here, ϵ is the total probability mass that each checkpoint assigns to the student's candidates. Both checkpoints vary with ϵ, while the student, candidate set, and conditional distributions remain fixed.

Then, for every $\epsilon \in (0, 1)$:

(i) Reward invariance. *For every candidate $v \in \mathcal{V}_K(\boldsymbol{s})$, the Direct-OPD reward (Eq. 1) with $\pi_{\mathrm{T}} = P_\epsilon$ and $\pi_{\mathrm{ref}} = Q_\epsilon$ satisfies*

$$\Delta_\epsilon(v \mid \boldsymbol{s}) = \log \frac{P_\epsilon(v)}{Q_\epsilon(v)} = \log \frac{t_v}{q_v}. \tag{16}$$

(ii) Fixed, nonzero update. *The local reward gradient of Direct-OPD (Eq. 4) at θ_0,*

$$\mathbf{g}_\epsilon = \sum_{v \in \mathcal{V}_K(\boldsymbol{s})} \bar{p}(v) \log \frac{t_v}{q_v} \nabla_\theta \log \pi_\theta(v \mid \boldsymbol{s}) \Big|_{\theta = \theta_0} = \mathbf{g}, \tag{17}$$

is independent of ϵ and nonzero with respect to the student's logits.

(iii) Divergences scale with the mass. *The teacher–reference divergences satisfy*

$$D_{\mathrm{JS}}(P_\epsilon, Q_\epsilon) = \epsilon D_{\mathrm{JS}}(\mathbf{t}, \mathbf{q}), \ D_{\mathrm{KL}}(P_\epsilon \| Q_\epsilon) = \epsilon D_{\mathrm{KL}}(\mathbf{t} \| \mathbf{q}), \ D_{\mathrm{KL}}(Q_\epsilon \| P_\epsilon) = \epsilon D_{\mathrm{KL}}(\mathbf{q} \| \mathbf{t}). \tag{18}$$

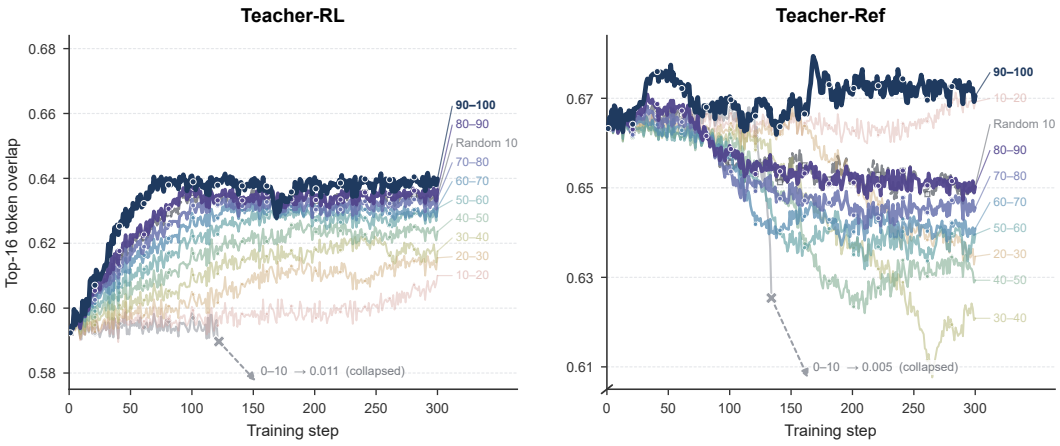


Figure 7: Top-16 token overlap over training for Qwen3-1.7B under the JustRL teacher pair, using JSD deciles or random 10% selection. Panels compare the student with the post-RL teacher (left) and pre-RL reference (right). Overlap is the fraction of shared tokens between top-16 sets.

Consequently, as $\epsilon \to 0$, all three divergences vanish, whereas the reward and its gradient remain fixed and nonzero.

Proof. (i) The common factor ϵ cancels in the ratio. (ii) By (i), each term of $\mathbf{g}_\epsilon$ is independent of ϵ, since $\mathcal{V}_K(\boldsymbol{s})$ and $\bar{p}$ depend only on θ_0. For softmax logits $\mathbf{z}$, $\nabla_{\mathbf{z}} \log \pi(v \mid \boldsymbol{s}) = \mathbf{e}_v - \pi(\cdot \mid \boldsymbol{s})$, so $\mathbf{g} = \mathbf{w} - (\sum_u w_u)\, p$, where $w_v = \bar{p}(v) \log(t_v/q_v)$ on $\mathcal{V}_K(\boldsymbol{s})$ and $w_v = 0$ elsewhere. Since $\mathcal{V}_K(\boldsymbol{s}) \subsetneq \mathcal{V}$ and $p > 0$, $\mathbf{g} = \mathbf{0}$ would require $\mathbf{w} = \mathbf{0}$, i.e., $\mathbf{t} = \mathbf{q}$. (iii) Off the candidate set, P_ϵ, Q_ϵ, and their mixture coincide, so the shared tail contributes zero to each divergence; on the candidate set, the factor ϵ cancels inside each logarithm and factors out of the sum. □

C Token Overlap Across Divergence Ranks

Prior work associates successful OPD with increasing overlap between the student's and teacher's high-probability token sets (Li et al., 2026b). Direct-OPD instead transfers a teacher–reference log-ratio, motivating us to examine overlap with both checkpoints. We track how these overlaps evolve when supervision is restricted to different teacher–reference JSD deciles.

For a student policy π_S and a comparison policy π_C at state s_t, let

$$\mathcal{V}_k^S(\boldsymbol{s}_t) = \mathrm{TopK}\big(\pi_S(\cdot \mid \boldsymbol{s}_t), k\big), \tag{19}$$

$$\mathcal{V}_k^C(\boldsymbol{s}_t) = \mathrm{TopK}\big(\pi_C(\cdot \mid \boldsymbol{s}_t), k\big), \tag{20}$$

so that the per-state overlap ratio is

$$\mathrm{Overlap}_k(S, C; \boldsymbol{s}_t) = \frac{\big|\mathcal{V}_k^S(\boldsymbol{s}_t) \cap \mathcal{V}_k^C(\boldsymbol{s}_t)\big|}{k}. \tag{21}$$

which we report with $k = 16$ against both the post-RL teacher ($C = T$) and the reference ($C = T_{\mathrm{ref}}$). Every bin starts from the same student initialization, so the two panels record how training on a given decile moves the student relative to each checkpoint.

Higher-JSD bins generally attain greater overlap with the post-RL teacher (Figure 7, left), broadly tracking the performance ordering in Figure 3a. This resembles the overlap growth observed in successful standard OPD (Li et al., 2026b), here under an objective that transfers the teacher's policy shift. The reference panel adds a complementary observation: the highest-JSD bin increases teacher overlap while maintaining reference overlap above its initial level. Reference overlap alone does not follow the performance ordering; for example, the 10–20 bin recovers close to its initial reference overlap despite much smaller gains in teacher overlap.

The strongest transfer occurs without a trade-off between overlap with the two checkpoints: the highest-JSD bin increases teacher overlap while sustaining reference overlap. This suggests that

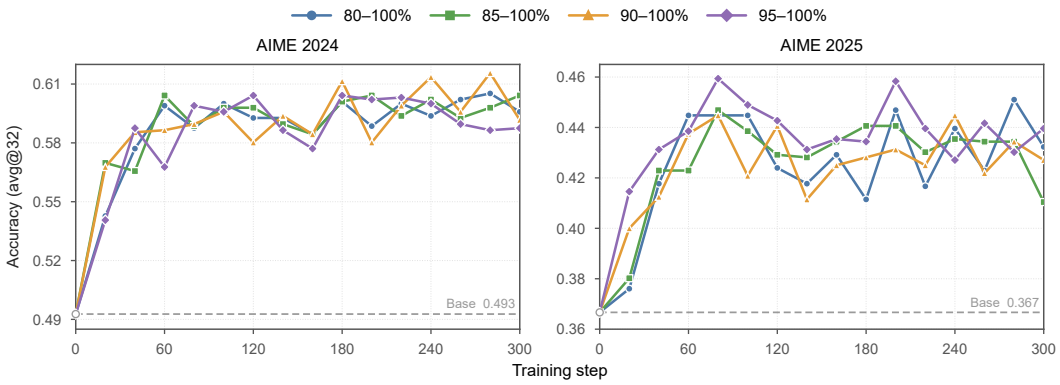


Figure 8: Sensitivity to the per-response retention ratio for Qwen3-1.7B under the JustRL teacher pair. The four settings retain the top 20%, 15%, 10%, and 5% of positions ranked by teacher–reference JSD. Dashed lines indicate base-model accuracy.

transferring the RL-induced policy shift need not entail moving from the reference's high-probability candidate set. The teacher–reference log-ratio can favor a token that remains highly ranked under both checkpoints, allowing transfer through changes in relative preference among shared candidates. Related evidence from standard OPD shows that supervision restricted to tokens shared by the student and teacher recovers nearly the full benefit of student top-k supervision (Li et al., 2026b). Our JSD criterion identifies positions with substantial changes in candidate probabilities or total mass relative to the tail (Eq. 8). Together, these observations suggest a view of selective transfer as learning substantial changes in probability allocation within largely overlapping candidate spaces.

D Sensitivity to Retention Ratio

Our main experiments retain the 10% of positions with the highest teacher–reference JSD within each response. Figure 8 compares retention ratios of 5%, 10%, 15%, and 20% for Qwen3-1.7B under the JustRL teacher pair. All four settings yield broadly similar learning curves and substantial gains over the base model on AIME 2024 and AIME 2025. No retention ratio consistently dominates across both benchmarks and training checkpoints.

Together with the decile analysis in Section 5.3, these results reveal an asymmetry in supervision selection. Training on low-JSD bins can substantially degrade transfer, whereas varying the retained fraction within the high-JSD region has a comparatively small effect. Moreover, the most selective setting does not consistently outperform the broader subsets. This pattern supports using JSD to screen supervision, without requiring that larger divergence always imply greater transfer utility. It also accords with our theoretical motivation: small divergence bounds the observable policy change, while large divergence alone does not establish the value of transferring that change. The similar performance across retention ratios suggests that effective selection admits a broad operating range in this setting, and we adopt 10% as a common default.

E Case Study: What Policy Changes Does Selection Preserve?

To examine what the mask keeps and removes, we inspect states from 20 rollouts of the initial Qwen3-1.7B student on AIME 2026 under the JustRL pair. At each state, we compare the teacher and reference next-token distributions over the student's top-K candidates and the residual token, and we mark the state as retained if its JSD falls within the top 10% of its response, following the per-response rule of Sec. 4.2. We write Δp for the change in a token's probability from the reference to the teacher. We selected the four cases below by hand for interpretability, so they illustrate how the score behaves rather than estimate how often each pattern occurs.

Selection acts on states, not on sampled tokens. Fig. 9 contrasts two states at which the student samples 3. At the first, 3 completes a valid calculation, and the teacher and reference distributions are nearly identical (JSD $= 0.0001$). At the second, 3 sets an incorrect boundary: $k = 3$ still

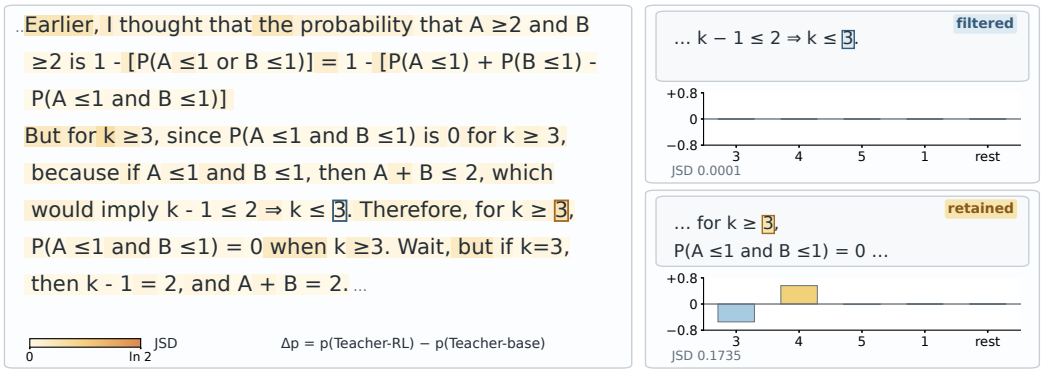


Figure 9: **Same sampled token, different mask decisions.** The mask filters a calculation on which the teacher and reference agree but retains supervision for a boundary correction.

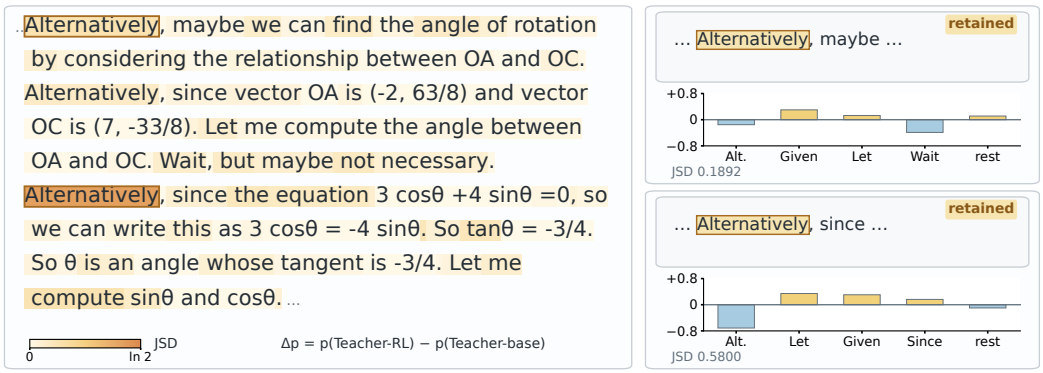


Figure 10: **Retained shifts encode continuation preferences.** At two retained states, the teacher shifts probability away from `Alternatively`; the later state favors `Let`, `Given`, and `Since`.

permits $A = B = 1$, so excluding their simultaneous occurrence requires $k \geq 4$. Here RL moves probability from 3 to 4 ($\Delta p \approx -0.55$ and $+0.56$), and the state is retained (JSD $= 0.1735$). The same token thus receives opposite mask decisions, because the score depends on how the teacher's distribution changed at the state rather than on which token was sampled.

Retained shifts include preferences over how reasoning continues. Fig. 10 shows two retained states at which the student opens a new passage with `Alternatively`. At both, RL lowers the probability of this token; at the later state the decrease is large ($\Delta p \approx -0.71$), and the teacher instead favors `Let`, `Given`, and `Since`. Together with the arithmetic cases, this shows that the retained supervision covers both local corrections and preferences over how the reasoning proceeds, consistent with a score that measures behavioral change rather than correctness.

JSD does not judge correctness. In Fig. 11, the student describes `R1 = 6` as `odd`. RL reinforces this error, raising the probability of `odd` ($\Delta p \approx +0.41$) and lowering that of `even` ($\Delta p \approx -0.12$). The JSD of 0.1197 exceeds the response's retention cutoff of 0.0744, so this adverse shift is retained. At the next state, conditioned on the erroneous statement, the teacher favors `Wait` over `Therefore` (from 0.01 to 0.71 and from 0.49 to 0.02), a preference compatible with reconsideration, but this later shift does not cancel the reward for `odd` at the earlier state. The score thus selects states by the size of the teacher's change, not by whether the change is correct.

Summary. These cases separate where supervision is applied from what it encourages. JSD decides which states are retained, while Direct-OPD's log-ratio rewards decide which tokens are encouraged at those states. Because the score measures how much the teacher's distribution changed on the student's candidates, the retained states include numerical corrections, continuation preferences, and local errors alike. Masking therefore concentrates Direct-OPD supervision on the states where RL changed the teacher most, without filtering shifts by their correctness.

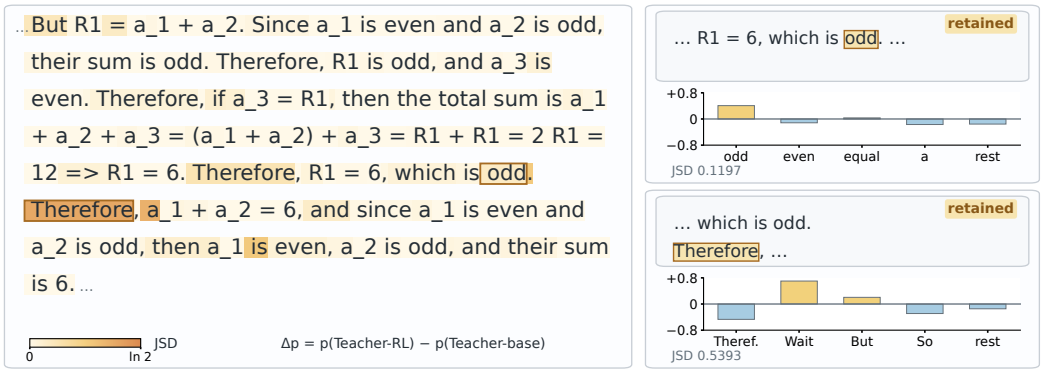


Figure 11: **A later reconsideration cue does not undo an earlier adverse shift.** The mask retains a shift toward an incorrect parity judgment, followed by a teacher preference for reconsideration.

Table 2: Default training and evaluation configuration.

Setting	**Training**	**Evaluation**
Framework	`verl`	–
Hardware	8× NVIDIA H200	–
Global batch size	128	–
Mini-batch size	128	–
Rollout n	4	–
Max. prompt length	1,024	–
Max. response length	2,048	31,744
Samples per problem	–	32
Sampling temperature	1.0	0.7
OPD support size Top-K	16	–
Top-p sampling	1.0	0.95
Learning rate	1×10^{-6}	–
Training steps	300	–
KL coefficient α	Adaptive	–
Controller ϵ	0.01	–
$[\alpha_{\min}, \alpha_{\max}]$	$[0.5, 2.5]$	–
Checkpoint selection	–	AIME 24/25
Held-out evaluation	–	AIME 26, HMMT Nov.25 / Feb.26

F Training Details

Data and prompt. All runs use the math subset of Skywork-OR1-RL-Data (He et al., 2025b) and apply the following prompt template.

```
Solve the following math problem step by step.
The last line of your response should be of the form
Answer:  $Answer (without quotes) where $Answer is the
answer to the problem.

{Question}

Remember to put your answer on its own line after "Answer:".
```

Training and evaluation. All experiments are implemented with `verl` and run on 8 NVIDIA H200 GPUs. Table 2 summarizes the default training and evaluation configuration used throughout our experiments. Owing to limited compute, each configuration is trained with a single run; the confidence interval in Sec. 5.2 resamples held-out problems and does not reflect variation across training seeds.

G LIMITATIONS

Our work has four main limitations. First, JSD is a proxy for how much RL changed the teacher, not a judge of whether the change is correct. Because it scores only the magnitude of the teacher–reference difference, S^{2}D-OPD retains a large shift toward an error as readily as a large correction (Fig. 11), so the quality of the retained supervision still depends on the teacher's RL. Selectors that also account for the direction or correctness of a shift, for example through verifier or outcome signals, could filter such states. Second, JSD is blind to what the student needs. The score depends on the student only through its top-K candidate set: a state where the student already follows the teacher's shift consumes the same budget as one where it does not, and a low-divergence shift that the student lacks is discarded. Combining teacher–reference divergence with student-side signals, in the spirit of the student-side calibration in Proxy-OPD (Fu et al., 2026a), is a natural extension. Third, our evidence is limited in scale and scope. We transfer from 1.5B teachers to students of up to 8B parameters, whereas weak-to-strong transfer is most valuable for much larger students, where RL is most costly. All experiments use mathematical reasoning with verifiable rewards; whether selective supervision helps in code generation or agentic tasks, where RL-induced shifts may be distributed differently across states, remains untested. Finally, owing to limited compute, each configuration is trained once, so our confidence interval reflects variation across held-out problems rather than across training seeds; repeated runs would strengthen the per-setting comparisons.