arXiv:2611.32452v1 [cs.LG] 24 Nov 2026

The Teacher Is Already Inside: Self-Distilling Looped Language Models Beyond Their Last Loop

Linghua Zhang

Abstract

Looped language models apply one shared transformer stack several times per token, and every loop is readable through the shared language-model head. Post-training them is hard. Reward-only reinforcement learning needs a verifier and learns from one sparse signal per response; spreading the reward over loops collapses the entropy of every loop; and cross-loop self-distillation, as in LoopOPD, can only teach intermediate loops, because the last loop cannot learn from its own frozen copy: the loss is zero at initialisation. We observe that a looped model already contains a policy that decodes better than its last loop: the next-token distribution extrapolated from the first loop through the last. ELPSD (Extrapolated-Loop Policy Self-Distillation) turns this distribution into a self-distillation target: a frozen copy of the model computes it, and the model's own last loop is distilled towards it on its own rollouts, using the exact reverse KL at every visited prefix. Teacher and student are the same model, so ELPSD needs no reward, no external teacher and no privileged context, and it leaves rollouts and decoding unchanged. We show that the objective pulls the last loop towards the base model's last loop with weight $1 + w$ and away from the base model's first loop with weight w, and that the teacher must stay frozen, since an idealised teacher recomputed from the current model extrapolates without bound. We evaluate ELPSD on Ouro-1.4B-Thinking across five math benchmarks against the base model, GRPO, RLTT and OPSD. Without any reward, ELPSD raises the average avg@8 from 61.15 to 63.69 and is the best method on all five benchmarks, ahead of GRPO (62.32), which uses a verifier, and the self-distillation method OPSD (61.77), which needs reference solutions.

1 Introduction

Looped language models reuse one set of transformer parameters across several passes per token [3, 21]. Ouro [21] applies the same stack $T = 4$ times per token, and because all loops share one language-model head, every loop t can be read out as a next-token distribution $P_\theta^{(t)}$. Rollouts and deployment use the last loop, $\pi_\theta = P_\theta^{(T)}$. Reinforcement learning with verifiable rewards (RLVR), usually GRPO [5, 15], is the main tool for improving the reasoning of such models after pre-training.

Existing methods leave the deployed loop poorly served. GRPO needs a verifier and learns from one scalar per response; a group whose responses are all correct or all wrong gives no gradient. On Ouro-1.4B-Thinking trained on DeepMath problems [6], 100 GRPO steps raise avg@8 over five math benchmarks by only 1.17 points (61.15 to 62.32). RLTT [17] spreads the reward over loops and trains every loop on tokens sampled from the last loop. Running the authors' released code in the same setup, loop-1 token entropy falls to 23% of its initial value and is still falling at step 100, and avg@8 ends below the base model on all five benchmarks, by 10.99 points on average (Section 4.2); LoopOPD [16] also reports RLTT below the model before RL in its own setup. LoopOPD removes the reward and lets the model teach itself: the frozen last loop of the base model is the teacher, and an intermediate loop, loop 2, is the student, trained with reverse KL on rollouts sampled from loop 2. The loop that is deployed, the last loop, therefore improves only indirectly, through the parameters it shares with loop 2, and training requires sampling from an intermediate

loop. LoopOPD cannot simply make the last loop its student: a last loop taught by its own frozen copy starts at zero loss and zero gradient, so training never moves.

Our goal is reward-free self-distillation: a dense signal that comes from the looped model itself and trains the deployed last loop directly on its own rollouts, with no external teacher and no privileged context. Our key observation is that a looped model already contains a policy better than its last loop. Cross-loop contrastive decoding (LoopCD) [12] extrapolates the log-probabilities from loop 1 through loop T, $(1 + w) \log P^{(T)} - w \log P^{(1)}$, and without any training raises the AIME24 pass@1 of Ouro-1.4B-Thinking from 50.83 to 59.17, as reported by Liu et al. [12]. Liu et al. [12] attribute most of this gain to re-ranking tokens rather than to a change of temperature. Internalising this policy into the weights pays off in three ways: the trained model decodes with standard last-loop sampling on any serving stack (LoopCD needs a custom model class in our vLLM implementation and runs without `torch.compile`); the student learns the teacher's behaviour on the states it visits itself; and the distilled model is a new starting point for test-time LoopCD or RL.

Training the last loop towards this policy raises three challenges. **C1:** the deployed last loop has no teacher above it inside the model; with $w = 0$ the teacher is its own frozen copy and the objective is degenerate. **C2:** an extrapolated teacher recomputed from the current model extrapolates its own extrapolation; the contrast between loops grows geometrically and there is no stable fixed point. **C3:** the extrapolated teacher amplifies tokens that loop 1 strongly disprefers, which the student rarely uses; forcing the student to cover them pulls it towards states it does not visit.

We propose ELPSD (Extrapolated-Loop Policy Self-Distillation), in which a looped model distills itself: its last loop learns, on its own rollouts, the LoopCD distribution of a frozen copy of the same model. Unlike self-distillation that gives the teacher privileged context, such as a reference solution [20], the teacher here sees exactly the student's input; its advantage comes only from reading the model's own loops differently. With $\ell_t = \log P^{(t)}_{\theta_0}(\cdot \mid s)$ and the loop direction $f = \ell_T - \ell_r$, the teacher is $\log \pi_{\mathrm{T}} = \log \mathrm{softmax}\big((1 + w)\ell_T - w\ell_r\big) = \ell_T + wf - \log Z(s)$, with $w = 0.5$ and $r = 1$, the same formula as test-time LoopCD. Extrapolating beyond the last loop gives the deployed loop a teacher that differs from it at initialisation (C1), and computing the teacher from the frozen base model gives distillation a fixed point, the teacher itself (C2). The objective is the full-vocabulary reverse KL $\mathrm{KL}(\pi_\theta \,\|\, \pi_{\mathrm{T}})$ at every prefix of rollouts sampled from the student, which is mode-seeking and on-policy (C3). Up to a constant, this objective equals $(1 + w)\mathrm{KL}\Big(\pi_\theta \,\Big\|\, P^{(T)}_{\theta_0}\Big) - w\mathrm{KL}\Big(\pi_\theta \,\Big\|\, P^{(1)}_{\theta_0}\Big)$: ELPSD pulls the last loop towards the base model's last loop and pushes it away from the base model's first loop (Proposition 1).

Our contributions are as follows.

- **Method.** ELPSD, on-policy self-distillation for looped language models: the last loop learns the LoopCD distribution of the model's own frozen initial weights; it is reward-free, needs no external teacher or privileged context, and keeps standard last-loop rollouts and decoding (Section 3).
- **Analysis.** A contrastive decomposition of the objective and of the initial gap, which equals a difference between a Rényi divergence and a KL divergence between the base model's last and first loops; a proof that the last loop cannot learn from its own frozen copy; and an argument that the extrapolated teacher must be frozen (Section 3.4).
- **Experiments.** A comparison of ELPSD with the base model, GRPO, RLTT and OPSD on five math benchmarks (Section 4). Without any reward, ELPSD raises the average avg@8 from 61.15 to 63.69 and is the best method on all five benchmarks, ahead of GRPO and the self-distillation baseline OPSD.

2 Related Work

Looped language models and their post-training. Universal Transformers loop one shared block with adaptive halting [3]. Ouro [21] pretrains this design on 7.7T tokens; its exploratory RLVR with DAPO and GRPO gave no significant gain over the SFT checkpoint. RLTT [17] replaces the last-loop log-probability in the policy gradient by a weighted sum over loops with one shared outcome advantage, applied to tokens sampled from the last loop. LSRL [14] adds external-grader scores of decoded intermediate states of the Huginn recurrent model to the outcome reward, but gradients flow only through the final-depth output. LoopOPD [16] removes the reward: an intermediate loop, on its own samples, minimises the full-vocabulary reverse KL to the last loop of the frozen initial model, and D-LoopOPD refreshes this teacher from the current parameters under a stop-gradient. ELPSD is also cross-loop self-distillation, but changes which loop learns and where the teacher sits. Its student is the deployed last loop, trained on unchanged last-loop rollouts, and its teacher lies beyond the last loop rather than at it. Its extrapolated teacher must also stay frozen (Proposition 2), whereas D-LoopOPD can refresh its teacher because that teacher is the terminal loop itself, not an extrapolation of it.

On-policy distillation and self-distillation. On-policy distillation (OPD) [1] trains a student on its own rollouts towards a teacher's next-token distribution. KDRL [18] adds reverse-KL distillation from a fixed external teacher to GRPO as $J_{\mathrm{GRPO}} - \beta D^{k2}_{\mathrm{KL}}$. ExOPD [19] casts OPD as dense KL-regularised RL with equal reward and KL weights; a reward scale $\lambda > 1$ extrapolates the optimum beyond the teacher, giving a target of the form $\pi \propto \pi_{\mathrm{ref}}(\pi_{\mathrm{teacher}}/\pi_{\mathrm{ref}})^{\lambda}$. D^2-OPD [4] bounds the extrapolated target's distance from the teacher by a per-prefix KL constraint, and EOPD [8] adds a top-k forward KL at high-entropy teacher positions. All four need a separate teacher model. Self-distillation removes the second model. OPSD [20] uses privileged information instead: the frozen initial model, given the problem together with its reference solution, teaches the same model without the solution, through a full-vocabulary divergence on the student's rollouts. It therefore needs a worked solution for every training problem. Our teacher has the form of ExOPD's extrapolated optimum with $\pi_{\mathrm{ref}} = P^{(r)}_{\theta_0}$, $\pi_{\mathrm{teacher}} = P^{(T)}_{\theta_0}$ and $\lambda = 1 + w$ (Equation (3)). It extrapolates between two loops of one frozen model instead of between two models, so both readouts come from a single forward pass, and it needs neither a second model nor reference solutions. ELPSD is thus self-distillation whose privileged view is extra computation, the model's own loops read through extrapolation, rather than extra context.

Contrastive and extrapolated decoding. Contrastive Decoding [10] corrects a large expert's logits with a small amateur's, and DoLa [2] contrasts the final layer with an earlier layer chosen per token. LoopCD [12] contrasts loops of a looped model, decoding from $(1+w)z_T - wz_1$ where z_t are loop-t logits, and attributes most of its gain to re-ranking tokens rather than to a temperature change; on Ouro-1.4B-Thinking it raises AIME24 pass@1 from 50.83 to 59.17 without training. These methods change only decoding, so every generated token needs the contrast and the serving stack must support it. ELPSD trains the last loop towards the LoopCD distribution of the frozen base model, so the trained model decodes with standard last-loop sampling.

3 Method

ELPSD is on-policy self-distillation for looped language models. The last loop is trained, on its own rollouts, towards a fixed teacher computed by the same model: the cross-loop contrastive decoding (LoopCD) distribution of its frozen initial weights. The objective contains no reward.

Table 1: Post-training methods for looped language models. "Loop 2" is LoopOPD's intermediate student loop $k = 2$; "last loop" is loop T, the loop used at deployment.

Method	Reward / verifier	Teacher	Student	Rollouts from	Teacher fixed
GRPO [15]	yes	none	last loop	last loop	n/a
RLTT [17]	yes	none (reward spread over loops by exit distribution)	all loops	last loop	n/a
LoopOPD [16]	no	frozen base, last loop	loop 2	loop 2	yes
D-LoopOPD [16]	no	current terminal loop, stop-gradient	loop 2	loop 2	no (refreshed)
OPSD [20]	no	frozen base, reference solution in context	last loop	last loop	yes
ELPSD (ours)	no	frozen base, LoopCD (loop 1 $\rightarrow$ loop T, $w = 0.5$)	last loop	last loop	yes

3.1 Preliminaries

Looped language models. A looped language model applies one shared stack of transformer layers T times to every token [3, 21]; Ouro [21] uses $T = 4$. With $h^{(t)}(s)$ the normalised hidden state after loop t at the last position of a prefix s, the shared LM head W gives one next-token distribution per loop:

$$P_\theta^{(t)}(\cdot \mid s) = \mathrm{softmax}\big(W\, h^{(t)}(s)\big), \qquad t = 1, \ldots, T. \tag{1}$$

We disable early exit, so the policy, which we also call the student, is the last loop, $\pi_\theta := P_\theta^{(T)}$, the distribution sampled at deployment. We write θ_0 for the initial parameters and $\ell_t := \log P_{\theta_0}^{(t)}(\cdot \mid s)$ for the base model's loop-t log-probabilities over the vocabulary $\mathcal{V}$.

On-policy self-distillation. The policy before the update, $\pi_{\theta_{\mathrm{old}}}$, samples responses y_i to a prompt x, which define the visited prefixes $s_{i,j} = (x, y_{i,<j})$; a teacher scores every visited prefix and the student is trained to match it there [1, 18, 19]. In self-distillation the teacher is the model itself with a privileged view; in OPSD [20] the frozen initial model additionally sees the reference solution. LoopOPD [16] applies this inside a looped model with an intermediate loop $k = 2$ as the student, the frozen base last loop as the teacher, and rollouts from $P_{\theta_{\mathrm{old}}}^{(k)}$:

$$\mathcal{L}_{\mathrm{LoopOPD}}(\theta) = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{|y_i|} \sum_{j=1}^{|y_i|} \mathrm{KL}\Big(P_\theta^{(k)}(\cdot \mid s_{i,j}) \,\Big\|\, P_{\theta_0}^{(T)}(\cdot \mid s_{i,j}) \Big). \tag{2}$$

The deployed last loop changes only through the parameters it shares with loop k.

3.2 A teacher beyond the last loop

LoopCD [12] applies contrastive decoding [2, 10] inside one looped model, with the last loop as the strong model and a reference loop r as the weak one; in log space it extrapolates past the last loop along the direction from loop r to loop T. Without training it raises the AIME24 pass@1 of Ouro-1.4B-Thinking from 50.83 to 59.17, mostly by re-ranking tokens rather than by a change of temperature [12]. Our teacher is this rule applied to the frozen base parameters, at temperature 1 and without truncation:

$$\log \pi_{\mathrm{T}}(\cdot \mid s) = \log \mathrm{softmax}\big((1 + w)\, \ell_T - w\, \ell_r\big) = \ell_T + w\, f - \log Z(s), \tag{3}$$

with $w \geq 0$, the base model's loop direction $f = \ell_T - \ell_r$, and $Z(s) = \sum_{v \in \mathcal{V}} P_{\theta_0}^{(T)}(v \mid s)^{1+w} P_{\theta_0}^{(r)}(v \mid s)^{-w}$. Both readouts come from one forward pass of the frozen model, and the teacher never changes during training.

Why the teacher must lie beyond the last loop. With $w = 0$ the teacher is the base model's last loop, which is exactly the student at initialisation: the loss below and its gradient are zero at every prefix, and training does not move. The last loop cannot learn from its own frozen copy. LoopOPD avoids this by making an intermediate loop the student, which leaves the deployed loop without a direct signal. ELPSD keeps the deployed loop as the student and moves the teacher beyond it; for $w > 0$ the two differ at initialisation by the gap in Proposition 1.

3.3 Objective

At every step $\pi_{\theta_{\text{old}}}$ samples G responses per prompt with the standard last-loop rollout engine, N in total. The loss is the exact reverse KL from the student to the teacher at every visited prefix, averaged over the tokens of each response and then over responses, the normalisation of LoopOPD:

$$\mathcal{L}_{\text{ELPSD}}(\theta) = \frac{1}{N} \sum_{i=1}^{N} \frac{1}{|y_i|} \sum_{j=1}^{|y_i|} \mathrm{KL}(\pi_\theta(\cdot \mid s_{i,j}) \,\|\, \pi_{\mathrm{T}}(\cdot \mid s_{i,j}))\,. \tag{4}$$

Each prefix contributes only its own KL, computed in closed form over the vocabulary, so the gradient is exact: prefixes are sampled, next tokens are not, and no importance ratio or clipping is needed (Appendix A). We take one optimiser update per step. The verifier is never called by the objective; we only log the accuracy of the training rollouts.

Reverse rather than forward KL. The reverse KL is mode-seeking: it penalises mass where the teacher has little, but does not force the student to cover every token the teacher supports. This suits an extrapolated teacher. Tokens that loop r strongly disprefers have large f, so the teacher can raise them far above their probability under the last loop, which rarely samples them. The forward KL $\mathrm{KL}(\pi_{\mathrm{T}} \,\|\, \pi_\theta)$ would require the student to cover all of them.

3.4 Analysis

All statements hold at a fixed prefix, which we drop; proofs are in Appendix A.

Proposition 1 (Contrastive decomposition and initial gap)**.** *For every distribution π over $\mathcal{V}$,*

$$\mathrm{KL}(\pi \,\|\, \pi_{\mathrm{T}}) = (1+w)\,\mathrm{KL}\Big(\pi \,\Big\|\, P_{\theta_0}^{(T)}\Big) - w\,\mathrm{KL}\Big(\pi \,\Big\|\, P_{\theta_0}^{(r)}\Big) + \log Z, \tag{5}$$

with $\log Z = w\, D_{1+w}\Big(P_{\theta_0}^{(T)} \,\Big\|\, P_{\theta_0}^{(r)}\Big) \geq 0$, where $D_\alpha(p \,\|\, q) = \frac{1}{\alpha - 1} \log \sum_v p(v)^\alpha q(v)^{1-\alpha}$ is the Rényi divergence of order α. At initialisation, $\pi_\theta = P_{\theta_0}^{(T)}$ and

$$\mathrm{KL}\Big(P_{\theta_0}^{(T)} \,\Big\|\, \pi_{\mathrm{T}}\Big) = w\left[D_{1+w}\Big(P_{\theta_0}^{(T)} \,\Big\|\, P_{\theta_0}^{(r)}\Big) - \mathrm{KL}\Big(P_{\theta_0}^{(T)} \,\Big\|\, P_{\theta_0}^{(r)}\Big)\right] \geq 0,$$

which is non-decreasing in w and zero if and only if $w = 0$ or $P_{\theta_0}^{(T)} = P_{\theta_0}^{(r)}$.

By Equation (5), minimising $\mathcal{L}_{\mathrm{ELPSD}}$ pulls the last loop towards the base model's last loop with weight $1+w$ and pushes it away from the base model's loop r with weight w; $\log Z$ is constant in θ. The pull term has the larger weight, and the sum is a KL divergence, bounded below by zero. At $\theta = \theta_0$ the pull term has zero gradient, so the first updates move the last loop away from loop r. The initial gap is large where the two base loops disagree and zero where they agree.

Proposition 2 (A self-refreshed extrapolated teacher diverges)**.** *Idealise training as rounds. Let* $\pi^{(0)} = P_{\theta_0}^{(T)}$. *In round* k *the teacher is recomputed from the current last loop,* $\log \pi_{\mathrm{T}}^{(k)} = \log \mathrm{softmax}\big((1+w)\log \pi^{(k)} - w\,\ell_r\big)$, *with loop* r *held at its base value, and distillation reaches it exactly,* $\pi^{(k+1)} = \pi_{\mathrm{T}}^{(k)}$. *Then* $\log \pi^{(k)} = \ell_r + (1+w)^k f - \log Z_k$. *For* $w > 0$, $\pi^{(k)}$ *concentrates on* $\arg\max_v f(v)$, *and the only fixed point of the round map,* $P_{\theta_0}^{(r)}$, *is repelling. With the frozen teacher, the minimiser of* $\mathrm{KL}(\pi \,\|\, \pi_{\mathrm{T}})$ *is* π_{T} *itself, a fixed point of distillation.*

The limit concentrates on tokens where loop T most exceeds loop r, however unlikely they are under loop T, so it is a degenerate policy rather than a sharper teacher. A frozen teacher removes this feedback. The refreshed teacher of D-LoopOPD [16] is the current last loop itself, not an extrapolation, so its recursion has no factor larger than one. Our teacher also has the form of ExOPD's extrapolated optimum [19] with π_{ref} the base loop r, π_{teacher} the base loop T and $\lambda = 1 + w$. Summed along a response, the per-prefix reverse KLs give the sequence-level KL to LoopCD decoding in expectation (Appendix A).

3.5 Implementation

The teacher is a frozen float32 copy of the initial model on every training GPU. Its full-vocabulary distributions are too large to store with the batch, so for each micro-batch the teacher runs its loop-wise forward without gradients and passes its loop-1 and loop-4 hidden states and LM head into the student's forward, which computes the KL of Equation (4) in chunks of positions under activation checkpointing. Rollouts use the same vLLM [9] engine and settings as GRPO. We use $w = 0.5$ and $r = 1$, the LoopCD setting for Ouro-1.4B-Thinking on math [12].

4 Experiments

4.1 Setup

Training. All trained systems start from Ouro-1.4B-Thinking [21] (24 layers, $T = 4$) and update all parameters on 14,684 problems filtered from DeepMath [6]. Each step samples 64 prompts $\times$ $G = 8$ responses ($N = 512$) at temperature 1.0 and top-p 1.0 and performs one optimiser update; all runs take 100 steps with learning rate 10^{-6}, responses of at most 16,384 tokens, no KL term to a reference model, no entropy bonus, early exit disabled and seed 42. A binary outcome reward R enters only the objectives of GRPO and RLTT; for the other systems we compute it only to log training accuracy. Rollouts come from the last loop in every system. Appendix B lists all settings.

Compared systems. *Base* is Ouro-1.4B-Thinking without training. *GRPO* [15] trains the last loop with the clipped surrogate on group-normalised advantages. *RLTT* [17] runs the authors' released code in its original REINFORCE form, with exit-distribution weights over loops. *OPSD* [20] distills the last loop, on its own rollouts, towards the frozen initial model that additionally sees the reference solution of each problem, with the full-vocabulary generalised Jensen–Shannon divergence ($\beta = 0.5$, the paper's default) and no reward. ELPSD minimises Equation (4) with $w = 0.5$, $r = 1$.

Table 2: avg@8 on five mathematics benchmarks (Ouro-1.4B-Thinking, step 100, 4 loops, 8 samples, temperature 0.6). Avg. is the unweighted mean. Best per column in bold.

Method	AIME24	AIME25	AIME26	MATH-500	AMC23	Avg.
Base	45.42	40.83	42.08	91.78	85.63	61.15
GRPO	45.00	40.83	45.00	92.65	88.13	62.32
RLTT	32.50	27.92	29.17	85.58	75.63	50.16
OPSD	44.58	40.00	44.58	92.20	87.50	61.77
ELPSD	**47.50**	**41.25**	**46.67**	**93.05**	**90.00**	**63.69**

Evaluation. We evaluate the step-100 checkpoints on AIME24, AIME25, AIME26, MATH-500 [7, 11] and AMC23, five benchmarks with 630 problems in total. Every model runs 4 loops and decodes from its last loop without contrastive decoding. We draw 8 samples per problem (temperature 0.6, top-p 0.95, at most 16,384 tokens, seed 42) and report avg@8, the mean accuracy over all samples.

4.2 Main results

Table 2 reports avg@8. Reward-only training gives a small gain: GRPO raises the average avg@8 from 61.15 to 62.32 ($+1.17$); per benchmark, avg@8 changes by -0.42 (AIME24) to $+2.92$ (AIME26). Spreading the reward over loops falls below the base model: RLTT lowers avg@8 on all five benchmarks, by 6.20 (MATH-500) to 12.92 (AIME24) points, to an average of 50.16 (-10.99). During its training, loop-1 entropy falls from 1.684 to 0.382 nats, 23% of its start, and all four loops end between 0.36 and 0.41 nats, consistent with training intermediate loops on tokens sampled from the last loop. Reward-free self-distillation with privileged information helps only slightly: OPSD raises the average to 61.77 ($+0.62$) without a verifier, but it needs a reference solution for every training problem.

ELPSD gives the best avg@8 on all five benchmarks without any reward. It decodes exactly like Base, and raises the average from 61.15 to 63.69 ($+2.54$), by $+0.42$ (AIME25) to $+4.59$ (AIME26) per benchmark. It exceeds GRPO by 1.37 points on average and on every benchmark ($+0.40$ to $+2.50$), although GRPO calls the verifier on all 512 rollouts of every step and ELPSD never does. It also exceeds OPSD by 1.92 points on average and on every benchmark ($+0.85$ to $+2.92$) without any reference solution: in this setting the model's own extra loops are a better teacher than privileged context.

5 Conclusion

A looped language model computes, but does not sample from, a distribution beyond its last loop: the extrapolation from an early loop through the last one. ELPSD lets the model distill itself: its last loop learns, on its own rollouts, this LoopCD distribution of its own frozen initial weights, with no reward, verifier, external teacher or privileged context, and with unchanged rollouts and decoding. The objective pulls the last loop towards the base loop T and away from the base loop 1 (Proposition 1), and the teacher must stay frozen, since a self-refreshed one compounds the contrast (Proposition 2). Without any reward, ELPSD raises the average avg@8 of Ouro-1.4B-Thinking from 61.15 to 63.69 and is the best method on all five benchmarks, ahead of GRPO, RLTT and OPSD.

Future work. ELPSD is complementary to reward-based training: an outcome reward can be added on top of the self-distillation term, giving a reinforcement learning objective anchored at the extrapolated teacher. Schedules for w and r could extrapolate further as training proceeds, and because the self-distilled

model is again a looped model, test-time LoopCD can be applied to it in turn. Extending ELPSD to larger looped models and to domains beyond mathematics is a natural next step.

References

[1] Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. On-policy distillation of language models: Learning from self-generated mistakes. In *International Conference on Learning Representations*, 2024.

[2] Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R Glass, and Pengcheng He. DoLa: Decoding by contrasting layers improves factuality in large language models. In *International Conference on Learning Representations*, volume 2024, pages 54158–54183, 2024.

[3] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal Transformers. *arXiv preprint arXiv:1807.03819*, 2018.

[4] Bowen Fang, Dingmin Wang, and Hengzhi Pei. D^2-OPD: Separating direction and distance in extrapolative on-policy distillation. *arXiv preprint arXiv:2611.07241*, 2026.

[5] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, et al. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

[6] Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, et al. DeepMath-103K: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning. In *International Conference on Learning Representations*, volume 2026, pages 138306–138322, 2026.

[7] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Advances in Neural Information Processing Systems Datasets and Benchmarks Track*, 2021.

[8] 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. *arXiv preprint arXiv:2603.07079*, 2026.

[9] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. 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*, 2023.

[10] Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori B Hashimoto, Luke Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization. In *Proceedings of the 61st annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 12286–12312, 2023.

[11] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations*, 2024.

[12] Weihao Liu, Huangjie Zheng, Tianrong Chen, Rohit Dilip, Richard He Bai, Yizhu Jiao, Yuyang Wang, and Ruixiang Zhang. Decoding looped Transformers better for (almost) free. *arXiv preprint arXiv:2610.02185*, 2026.

[13] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[14] Hangliang Ren. LSRL: Process-supervised GRPO on latent recurrent states improves mathematical reasoning. In *EMNLP (Findings)*, pages 12534–12545, 2025.

[15] 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. *arXiv preprint arXiv:2402.03300*, 2024.

[16] Yi Wang, Rui Qian, Yu Li, Haoyang Yao, and Wenjie Wang. Recurrent self-improvement: Dynamic cross-loop on-policy distillation for looped language models. *arXiv preprint arXiv:2610.10623*, 2026.

[17] Jonathan Williams and Esin Tureci. Prioritize the process, not just the outcome: Rewarding latent thought trajectories improves reasoning in looped language models. *arXiv preprint arXiv:2602.10520*, 2026.

[18] Hongling Xu, Qi Zhu, Heyuan Deng, Jinpeng Li, Lu Hou, Yasheng Wang, Lifeng Shang, Ruifeng Xu, and Fei Mi. KDRL: Post-training reasoning LLMs via unified knowledge distillation and reinforcement learning. *arXiv preprint arXiv:2506.02208*, 2025.

[19] Wenkai Yang, Weijie Liu, Ruobing Xie, Kai Yang, Saiyong Yang, and Yankai Lin. Learning beyond teacher: Generalized on-policy distillation with reward extrapolation. *arXiv preprint arXiv:2602.12125*, 2026.

[20] 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. *arXiv preprint arXiv:2601.18734*, 2026.

[21] Rui-Jie Zhu, Zixuan Wang, Kai Hua, Tianyu Zhang, Ziniu Li, Haoran Que, Boyi Wei, Zixin Wen, Fan Yin, He Xing, et al. Scaling latent reasoning via looped language models. *arXiv preprint arXiv:2510.25741*, 2025.

A Proofs

Proof of Proposition 1. By Equation (3), $\log \pi_{\mathrm{T}} = (1+w)\ell_T - w\ell_r - \log Z$. Splitting $\log \pi = (1+w)\log \pi - w \log \pi$ in $\mathbb{E}_\pi[\log \pi - \log \pi_{\mathrm{T}}]$ gives Equation (5). Next,

$$Z = \sum_v P_{\theta_0}^{(T)}(v)^{1+w} P_{\theta_0}^{(r)}(v)^{-w} = \mathbb{E}_{P_{\theta_0}^{(T)}}\left[(P_{\theta_0}^{(T)}/P_{\theta_0}^{(r)})^w\right] = \exp\left(w\, D_{1+w}\left(P_{\theta_0}^{(T)} \,\middle\|\, P_{\theta_0}^{(r)}\right)\right)$$

by the definition with $\alpha = 1+w$, and Rényi divergences are non-negative. Setting $\pi = P_{\theta_0}^{(T)}$ in Equation (5) removes the first term. The Rényi divergence is non-decreasing in its order and equals the KL divergence at order 1, so the bracket is non-negative and non-decreasing in w, and so is its product with w. The gap is zero if and only if $\pi_{\mathrm{T}} = P_{\theta_0}^{(T)}$, i.e. $wf(v) = \log Z$ for all v. For $w > 0$ this makes f constant; both loops are normalised softmax outputs with full support, so $f = 0$. □

Proof of Proposition 2. For $k = 0$, $\log \pi^{(0)} = \ell_T = \ell_r + f$. If the claim holds for k, then up to a constant $\log \pi^{(k+1)} = (1+w)\big(\ell_r + (1+w)^k f\big) - w\,\ell_r = \ell_r + (1+w)^{k+1} f$. Hence for $v^\star$ maximising f and $f(v) < f(v^\star)$, the ratio $\pi^{(k)}(v)/\pi^{(k)}(v^\star)$ equals $P_{\theta_0}^{(r)}(v)/P_{\theta_0}^{(r)}(v^\star)$ times $\exp\big((1+w)^k(f(v) - f(v^\star))\big) \to 0$. A fixed point has $w(\log \pi - \ell_r)$ constant, so $\pi = P_{\theta_0}^{(r)}$, and a perturbation g of $\log \pi - \ell_r$ is mapped to $(1+w)g$ up to a constant, so the fixed point is repelling. □

The tokens in $\arg\max_v f(v)$ are those where loop T most exceeds loop r, however unlikely they are under loop T, so the limit is a degenerate policy rather than a sharper teacher.

Exact gradient of Equation (4). At a prefix s, using $\sum_v \pi_\theta(v)\nabla_\theta \log \pi_\theta(v) = 0$,

$$\nabla_\theta \mathrm{KL}(\pi_\theta \,\|\, \pi_{\mathrm{T}}) = \sum_{v \in \mathcal{V}} \pi_\theta(v)\big[\log \pi_\theta(v) - \log \pi_{\mathrm{T}}(v)\big]\nabla_\theta \log \pi_\theta(v).$$

Prefixes are sampled, next tokens are not, so no importance ratio or clipping is needed.

Sequence-level view. With $\pi_\theta(y \mid x)$ and $\pi_{\mathrm{T}}(y \mid x)$ the products of per-token distributions (the latter is LoopCD decoding of the base model at temperature 1), the expected sum of per-prefix reverse KLs along a response from π_θ equals $\mathrm{KL}(\pi_\theta(y \mid x) \,\|\, \pi_{\mathrm{T}}(y \mid x))$. As in LoopOPD and KDRL [18], each prefix contributes its own KL term.

Relation to ExOPD. $P_{\theta_0}^{(r)}(P_{\theta_0}^{(T)}/P_{\theta_0}^{(r)})^{1+w} = (P_{\theta_0}^{(T)})^{1+w}(P_{\theta_0}^{(r)})^{-w}$, so the teacher has the form of ExOPD's extrapolated optimum [19] at every prefix, with $\pi_{\mathrm{ref}} = P_{\theta_0}^{(r)}$, $\pi_{\mathrm{teacher}} = P_{\theta_0}^{(T)}$ and $\lambda = 1 + w$.

B Experimental details

Method-specific settings. GRPO: clipped surrogate (clip 0.2, dual clip 3) on group-normalised advantages, last loop. RLTT: authors' code, REINFORCE on $\hat{A}_i \sum_t q_t \log P_\theta^{(t)}$ with the exit distribution q_t, exit gate trained. OPSD: frozen initial model with the reference solution in its context as teacher, last loop as student on its own rollouts, full-vocabulary generalised JSD with $\beta = 0.5$ (the defaults of Zhao et al. [20]). ELPSD: $w = 0.5$, $r = 1$, exact reverse KL, per-response token mean; the frozen float32 teacher is held unsharded on every GPU (about 5.6 GB), runs its loop-wise forward per micro-batch without gradients, and passes its loop-1 and loop-4 hidden states and LM head into the student's forward, where the KL is computed in chunks of positions under activation checkpointing.

Table 3: Settings shared by all trained systems.

Setting	Value
Initial model	Ouro-1.4B-Thinking, 24 layers, $T = 4$ loops, full-parameter training
Training data	14,684 filtered DeepMath problems, binary outcome reward
Prompts $\times$ responses per step	$64 \times 8 = 512$
Optimiser updates per step	1 (mini-batch = 64 prompts, 1 epoch)
Training steps	100
Learning rate	10^{-6}, 10 warm-up steps, then constant
Optimiser	AdamW [13], betas $(0.9, 0.99)$, weight decay 0.1
Gradient clipping	global norm 0.1
KL term to a reference model	none
Entropy bonus	none
Max prompt / response length	2,048 / 16,384 tokens
Rollout sampling	temperature 1.0, top-p 1.0, vLLM in bf16
Early exit	disabled (log-probabilities always from loop 4)
Seed (training and data order)	42
Hardware	one node, 8 GPUs