Learning from Recurrence:

Internalizing Depth-Induced Preferences in Looped Language Models

Hui Chen
University of Wisconsin-Madison
hchen795@wisc.edu

Joshua Elsdon
Microsoft
joshua@elsdon.io

Can Goksen
Microsoft
can.goksen@microsoft.com

arXiv:2611.33468v1 [cs.CL] 24 Nov 2026

Abstract

A looped language model repeatedly applies the same block to its hidden state before predicting each token. The number of passes is its *recurrent depth*. Increasing this depth gives the model more computation per token with the same parameters. Can the shared block learn from the predictions it makes at greater depth? Following prior work, we retrofit pretrained OLMo-2-1B with a shared recurrent block and continue mid-training over varying recurrent depths. In the four-pass model, additional passes mainly redistribute probability among tokens that already rank highly after one pass. We use the continuations generated at greater depth to train only the shared block through a single pass, without reference answers or correctness filtering. We call this transfer *preference write-back*. Write-back further improves one-pass and multi-pass accuracy on most tasks. At their target depths, the updated models outperform a non-recurrent control matched for mid-training data and token budget on all six mathematical reasoning benchmarks, with a mean relative gain in accuracy of up to 46.2% across tasks. At four-pass inference, one-pass training outperforms full-depth training on the same targets in the model with input injection. This advantage disappears in the matched model without injection. On GSM8K, using the updated block only on passes two through four recovers nearly the entire four-pass gain, showing that an update learned in one pass can improve later applications of the same block.

1 Introduction

Looped Transformer language models scale inference-time computation by reapplying recurrent blocks before emitting each token, decoupling computation depth from layer depth (Dehghani et al., 2019; Saunshi et al., 2025; Geiping et al., 2025). This makes recurrent depth attractive both for serving—via parameter sharing, depth-wise batching, and KV-cache reuse (Lan et al., 2020; Bae et al., 2025a,b; Geiping et al., 2025)—and for latent test-time computation, where the model iterates in hidden-state space before emitting a token (Geiping et al., 2025; McLeish et al., 2026; Koishekenov et al., 2025; Zhu et al., 2025). Empirically, such models can improve with additional recurrent depth and, in some settings, rival language models several times larger (Saunshi et al., 2025; Geiping et al., 2025; Zhu et al., 2025).

These gains, however, do not follow from weight sharing alone. Although reapplying a shared block makes additional depth available, the model must learn to use it across repeated passes. Recent work has therefore developed recurrent architectures and depth-training strategies toward this goal (Geiping et al., 2025; McLeish et al., 2026; Koishekenov et al., 2025). Building on these approaches, we retrofit a pretrained language model with recurrence and adapt it through continued mid-training over varying execution depths. Yet developing a model that benefits from additional recurrence leaves a further question: what do these additional passes actually change in its predictions?

Recurrent models can learn genuine iterative computation in controlled settings, including repeated algorithmic operations and computing intermediate results in one recurrent pass that are used in subsequent passes (Schwarzschild et al., 2021; Bansal et al., 2022; Saunshi et al., 2025; Jolicoeur-Martineau, 2025; Kohli et al., 2026; Huang et al., 2026). But these results do not determine the role of additional passes in pretrained language models retrofitted with recurrence (Bae et al., 2025a; McLeish et al., 2026; Koishekenov et al., 2025): improvements with depth may reflect new intermediate computation, refinement of an already-formed predictive state, or a combination of the two.

We therefore examine how the next-token distribution changes with recurrent depth. Our intermediate readouts show that additional recurrence primarily redistributes probability among tokens that already rank highly after shallow execution. The leading token often gains probability, while changes in the relative ordering of nearby candidates can also alter which token the model favors. These changes in preference can shape the continuation the model produces

and the context of subsequent predictions. This observation connects to recent results on simple self-distillation, where changes in a model's own predictive distribution can provide useful training signal (Zhang et al., 2026). Here, deeper recurrent execution itself produces these preference changes, suggesting that the resulting behavior could guide further learning of the shared recurrent operator without an additional correctness signal.

This paper tests this possibility by using trajectories generated at deeper recurrent depth as supervision for the same model, while training only the shared recurrent block through a shallower execution. We refer to this process as *preference write-back*. We take the shared recurrent operator as the object of improvement, with its subsequent repeated use as a central consideration. Shallow execution provides a path for learning the update, while its value is assessed through subsequent applications of the same operator.

Specifically, we make the following **contributions**. We show that additional recurrence refines predictive preferences among plausible continuations (Section 2.2), and turn this depth-induced behavior into supervision without reference answers or filtering (Section 2.3). By separating teacher, training and inference depth, we find that learning from deeper behavior improves recurrent inference, with single-pass learning outperforming full-depth learning on the same targets in our input-injected models (Section 3.2). We further show that the shallow-learned update remains useful when the shared operator is applied at later recurrent passes (Section 3.3). Controlled comparisons reveal that the advantage of shallow learning depends on the recurrent recipe, connecting how recurrent computation is developed to how it can subsequently be improved (Section 3.4). We also examine the source of supervision and its transfer across question pools (Section 3.5), together with the retention of capabilities beyond mathematical reasoning (Section 3.6).

2 Learning from recurrence

2.1 Looped language models and repeated computation across depth

A shared computation across depth. Looped language models increase execution depth by repeatedly applying the same recurrent core to an evolving latent state. We develop this shared computation through recurrent mid-training, then seek to improve the same core using behavior revealed by deeper execution.

Input-conditioned recurrent computation. We retrofit the OLMo-2 1B stage-1 checkpoint (Team OLMo et al., 2025) into a non-recurrent input stack P, a shared recurrent core F_θ, and a non-recurrent output stack C, following prior recurrent architectures (Geiping et al., 2025) and prior retrofits of pretrained models (McLeish et al., 2026). The input and output stacks are termed the *prelude* and *coda* in prior work. We assign the first seven Transformer blocks to P, the middle four to F_θ, and the final five to C, following the partition used for this model family (Koishekenov et al., 2025). This partition places recurrence in a middle region where the angular changes between consecutive hidden representations have stabilized, retaining separate input and output transformations.

Let $x_{1:n} \in \mathcal{V}^n$ be a prefix of n tokens from vocabulary $\mathcal{V}$. The input stack, including token embeddings, produces $e = P(x_{1:n}) \in \mathbb{R}^{n\times d}$, where d is the model width. Starting from an initial state $s_0 \in \mathbb{R}^{n\times d}$, an execution of depth r computes

$$s_i = F_\theta(A_\psi([s_{i-1}; e])), \qquad i = 1, \dots, r.$$

Here, $s_i \in \mathbb{R}^{n\times d}$ is the state after pass i, and $[\cdot;\cdot]$ denotes concatenation along the hidden dimension. The recurrent core $F_\theta : \mathbb{R}^{n\times d} \to \mathbb{R}^{n\times d}$ and the position-wise linear projection $A_\psi : \mathbb{R}^{2d} \to \mathbb{R}^d$ reuse the same parameters θ and ψ at every pass. This *input injection* makes the same representation e available throughout the evolution of the recurrent state. After the final pass, the output stack $C : \mathbb{R}^{n\times d} \to \mathbb{R}^{n\times d}$ and language-model head $W \in \mathbb{R}^{|\mathcal{V}|\times d}$ produce the next-token distribution $p_r(\cdot \mid x_{1:n}) = \mathrm{softmax}(W[C(s_r)]_n)$, where $[\cdot]_n \in \mathbb{R}^d$ selects the final-position representation. The input and output stacks are each applied once.

Learning across recurrent depths. We adapt the pretrained parameters to recurrent execution through continued causal language-model training on the Dolmino mixture (Team OLMo et al., 2025). During this stage, the execution depth is sampled once per training step from a Poisson-lognormal depth distribution, following prior recurrent-depth training (Geiping et al., 2025). Let

$$r \sim \pi_t, \qquad \mathbb{E}_{r\sim\pi_t}[r] = \mu_t,$$

where r is the sampled number of recurrent passes and π_t denotes the Poisson-lognormal depth distribution at training step t. Rather than training immediately at the final depth distribution, we follow the recurrence curriculum of McLeish

et al. (2026) and gradually increase its mean μ_t toward a target value m. The language-model objective is therefore optimized over recurrent computations of varying depth, training the shared core to operate across the depth distribution.

We call m the model's *mean depth*, using target values of 4, 8, and 32. Mean depth characterizes the recurrent-depth distribution under which the model is learned, whereas the *execution depth* r specifies the number of recurrent passes used in an individual execution. Once mid-training is complete, r remains independently selectable for a fixed checkpoint, so shallow and deep executions are obtained by applying the same learned recurrent computation for different numbers of passes. Initialization and curriculum details are provided in Appendix C, and Appendix A shows that recurrence already improves accuracy during mid-training. The resulting checkpoint supplies both the model to be updated and the deeper executions from which we construct supervision.

2.2 Recurrence-induced preferences as supervision

Once mid-training is complete, execution depth can be varied while the learned parameters remain fixed. We examine how additional recurrence changes the next-token distribution of a fixed model, focusing on how repeated applications of the shared core reshape its preferences among candidate continuations. We first characterize this change, then use the preference of deeper execution as supervision.

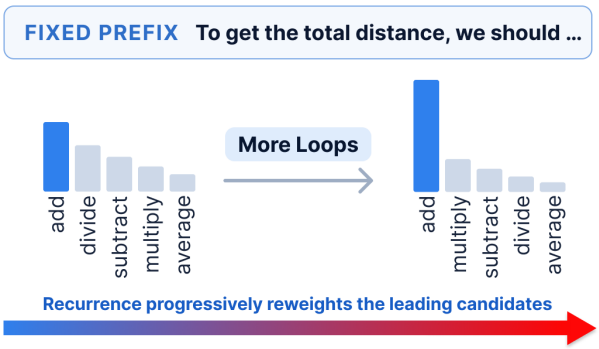

Figure 1: Recurrence revises predictive preference. At a fixed prefix, successive recurrent passes redistribute probability among leading candidates, occasionally changing which continuation is preferred.

Recurrence refines predictive preference. At the level of intermediate predictions, additional recurrence primarily refines preferences among already plausible continuations. We expose this change using intermediate readouts along a fixed deep-generation path, decoding each recurrent state with the same output stack while leaving the generation itself unchanged. In the mean-depth-4 model, the top-ranked token remains unchanged across four passes at 95.6% of generated positions, even as probability shifts toward the leading token while total top-10 mass changes little. When the final top-ranked token does change, 99.4% of those choices are already contained in the first-pass top five. The dominant effect of recurrence is therefore a redistribution of probability within the head of the predictive distribution, occasionally strong enough to change the selected token. These changes suggest using the choices made by deeper execution as behavioral supervision. Figure 1 illustrates this effect. Appendix B details the readout protocol and analysis.

From predictive preference to behavioral supervision. We next represent the preferences revealed by deeper recurrence as learnable targets. A decoding rule turns the deeper next-token distribution into a complete autoregressive continuation. This preserves the token preferred at each step and the contexts induced by those choices, through which subsequent preferences are expressed.

In our default construction, we freeze the pre-update checkpoint and choose a teacher depth r_{T}. For a prompt x, greedy decoding produces a continuation $\hat{y} = (\hat{y}_1, \ldots, \hat{y}_T)$ according to

$$\hat{y}_t = \arg\max_{v \in \mathcal{V}} p^{(0)}_{r_{\mathrm{T}}}(v \mid x, \hat{y}_{<t}), \qquad t = 1, \ldots, T,$$

where $p^{(0)}_{r_{\mathrm{T}}}$ is the predictive distribution of the frozen pre-update model after r_{T} recurrent passes. The resulting trajectory is a sequence-level realization of the preferences expressed by deeper execution. It retains the selected tokens and the contexts they create, without preserving the full predictive distribution.

For N training prompts, we collect one such continuation per prompt to form

$$\mathcal{D}_{r_{\mathrm{T}}} = \left\{ \left(x^{(j)}, \hat{y}^{(j)} \right) \right\}_{j=1}^{N}.$$

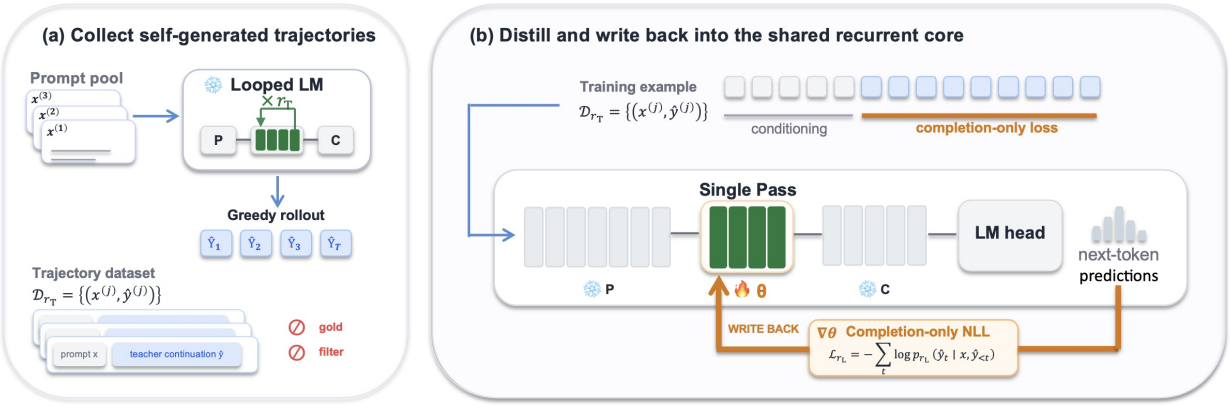

Figure 2: Preference write-back. (a) The frozen model generates one greedy continuation per prompt at teacher depth r_{T}, without reference answers or filtering. (b) These trajectories train only the shared recurrent core through a single pass (r_{L}=1). The updated core is reused at every pass.

Each target retains the complete generated continuation, including its reasoning and final answer. Reference answers to the target questions are neither provided during generation nor used to filter trajectories. Teacher depth determines which recurrent behavior is exposed, while the trajectory-construction rule determines which realization of that behavior is retained as supervision. How this supervision is written back into the shared computation is a separate choice, discussed next.

2.3 Writing preferences back into the shared recurrent core

Once deeper behavior has been converted into supervision, the remaining problem is to turn a transient effect of additional computation into a persistent change in the computation itself. In a weight-tied model, this means writing the behavior back into the shared operator that produced it and that subsequent recurrence will reuse. Crucially, the execution that reveals the behavior need not be the execution through which the operator learns it. We refer to this transfer from depth-induced behavior into reusable computation as *preference write-back* (Figure 2 and Algorithm 1 in Appendix C.3).

Decoupling behavior exposure from incorporation. Preference write-back separates the source of supervision from the computation through which that supervision is learned. The *teacher depth* r_{T} specifies the execution whose behavior is retained as supervision, while the *training depth* r_{L} specifies the recurrent computation through which that supervision acts on the shared operator. For a fixed target set $\mathcal{D}_{r_{\mathrm{T}}}$, changing r_{L} leaves both the behavioral target and the shared parameterization unchanged. Only the computational path connecting that target to the shared parameters changes. Behavior exposed at one depth can thus be learned through another execution of the same operator.

Writing behavior into the reusable operator. The supervision is behavioral, but its write-back destination is the shared recurrent operator. Let θ_0 denote the pre-update parameters of the recurrent core, and let ϕ collect the remaining model parameters, which are held fixed. ϕ includes the injection projection A_ψ, which acts in the forward computation but is not updated. Given $\mathcal{D}_{r_{\mathrm{T}}}$, we initialize $\theta = \theta_0$ and optimize the recurrent core through an execution of depth r_{L} using completion-only next-token likelihood.

$$\mathcal{L}_{r_{\mathrm{L}}}(\theta) = -\frac{1}{\sum_j T_j} \sum_{j=1}^{N} \sum_{t=1}^{T_j} \log p_{r_{\mathrm{L}}}\left(\hat{y}_t^{(j)} \mid x^{(j)}, \hat{y}_{<t}^{(j)}; \theta, \phi\right).$$

Here, $p_{r_{\mathrm{L}}}(\cdot\,; \theta, \phi)$ denotes the predictive distribution obtained after r_{L} applications of the recurrent core, in the notation of Section 2.1. Only the shared recurrent core is updated. All remaining model parameters stay fixed. The resulting parameters θ_1 store the deep-execution behavior in the same transformation that recurrence will invoke again. Training depth controls the path by which supervision reaches the operator, while the object being changed remains the reusable computation.

Reusing the written-back computation. After optimization, the updated core F_{θ_1} becomes the recurrent operator used at every subsequent pass, while the remaining parameters ϕ remain unchanged. At inference depth r_{E}, the model repeatedly applies F_{θ_1} as the recurrent state evolves. This reuse follows from weight sharing. Whether an update learned through shallow execution remains useful on states reached later in recurrence is a separate empirical question.

Preference write-back thus separates the depth that exposes behavior from the depth used to learn it and the depth at which the updated computation is reused. These roles are indexed by teacher depth r_{T}, training depth r_{L}, and inference depth r_{E}, respectively, while all three operate on the same shared recurrent computation. We therefore assess the update through its repeated use at inference, not only through the execution used to learn it.

3 Preference write-back in looped language models

Experimental setup. We study looped language models obtained by continuing mid-training an OLMo-2-1B base model toward mean recurrent depths of 4, 8, and 32. Unless otherwise stated, self-distillation follows the recipe introduced in Section 2. Each model generates greedy trajectories at its standard deeper execution depth, without correctness filtering, and only the shared recurrent block is updated through one recurrent pass. We compare the looped checkpoints with a non-looped control trained on the same data and 50B mid-training tokens. We evaluate each looped model before and after write-back at one pass and at its standard inference depth, focusing on mathematical reasoning under greedy decoding. Architecture, training, and evaluation details are given in Section 2 and Appendix C.

3.1 Learning from deeper recurrence improves recurrent inference

The looped checkpoints already benefit from additional recurrence before write-back. At their mean depths, all three outperform the non-looped control on each of the six mathematical reasoning benchmarks, with mean per-task relative gains of 26.5–39.9% (Table 1). On GSM8K, the mean-depth-4 checkpoint improves from 40.86% after one pass to 56.18% after four passes.

Learning from its own four-pass continuations raises one-pass accuracy to 48.52% and four-pass accuracy to 62.70%. The updated model continues to benefit from additional recurrence, gaining 14.18 points from one to four passes. Across the three models, write-back increases accuracy in 15 of the 18 model–benchmark comparisons at their mean depths, including all six for the mean-depth-4 model. GSM8K, GSM1K, GSM-Symbolic, and ASDiv improve at all three mean depths. The final models achieve mean per-task relative gains of 34.9–46.2% over the non-looped control.

Additional recurrence supplies supervision beyond one-pass self-training. On the mean-depth-4 model, we use the same question pool, one-pass learning, and four-pass inference, varying only the teacher depth (Table 7). Four-pass teacher continuations yield 62.70% on GSM8K and 54.36% on GSM1K, compared with 59.29% and 51.95% from one-pass teacher continuations.

Write-back further improves the looped checkpoint under a common learning budget. We compare gold-answer SFT, SSD-style self-training (Zhang et al., 2026), and distillation on a common random subset of 1,024 GSM8K training questions (Table 2). Both starting checkpoints share the pretrained initialization, mid-training data, data order, and 50B-token budget. The primary comparisons use a common cap of approximately 1.575×10^{16} estimated matrix-multiplication FLOPs for trajectory generation, teacher scoring, and student training. Sequence KD (Kim and Rush, 2016) learns from the same four-pass continuations as write-back and gives the best non-looped scores of 47.46% on GSM8K and 40.83% on GSM1K. The looped checkpoint already reaches 56.18% and 51.12% at four-pass inference.

On the looped checkpoint, DistiLLM-2 (Ko et al., 2025), GKD (Agarwal et al., 2024), and EOPD (Jin et al., 2026) use the same original four-pass teacher as write-back. Write-back achieves the highest four-pass accuracy on both tasks among the tested updates, at 60.50% and 51.54%. Its GSM8K gain of 4.32 points is more than twice the 2.12-point gain of DistiLLM-2, the strongest tested KD alternative on this task.

3.2 Learning depth shapes the update to the shared operator

The depth used for learning matters even when the teacher trajectories are fixed. We hold the targets, objective, trainable parameters, and optimization settings fixed, changing only the number of recurrent passes used during learning (Table 4). On the mean-depth-4 model, learning four-pass targets through one pass reaches 62.70% GSM8K accuracy at four-pass inference, compared with 56.86% after four-pass learning. The corresponding advantages are 3.32 points

Table 1: Mathematical reasoning before and after preference write-back across recurrent depths. Exact match (%) under greedy decoding at depth 1 and at the mean depth m. *Mean depth* rows report the looped checkpoints before write-back. The write-back rows update only the shared recurrent block using greedy depth-m trajectories (50,000 OpenMath questions for m=4, 10,000 otherwise). Parenthesized values and the last column report relative changes at depth m against the non-looped control.

Model	GSM8K		GSM1K		SVAMP		GSM-Symbolic		ASDiv		AQuA-RAT		Avg.
Depth	1	m	1	m	1	m	1	m	1	m	1	m	vs. control
Non-looped control	44.05	–	39.83	–	56.50	–	37.42	–	59.39	–	10.24	–	–
Mean depth 4	40.86	56.18 (+27.5%)	36.93	51.12 (+28.3%)	52.40	63.30 (+12.0%)	31.82	43.56 (+16.4%)	55.44	64.82 (+9.1%)	14.17	16.93 (+65.4%)	+26.5%
+ Self-distillation	48.52	62.70 (+42.3%)	44.65	54.36 (+36.5%)	56.40	64.70 (+14.5%)	36.92	47.64 (+27.3%)	59.09	66.55 (+12.1%)	12.20	18.11 (+76.9%)	+34.9%
Mean depth 8	22.59	56.33 (+27.9%)	20.50	52.53 (+31.9%)	39.90	66.00 (+16.8%)	15.76	45.06 (+20.4%)	46.33	66.20 (+11.5%)	10.24	23.62 (+130.8%)	+39.9%
+ Self-distillation	37.38	62.85 (+42.7%)	32.37	56.02 (+40.6%)	43.60	65.20 (+15.4%)	27.94	48.42 (+29.4%)	51.84	67.90 (+14.3%)	5.91	22.44 (+119.2%)	+43.6%
Mean depth 32	4.09	57.01 (+29.4%)	3.24	53.36 (+34.0%)	19.70	67.70 (+19.8%)	2.68	44.90 (+20.0%)	20.39	66.55 (+12.1%)	0.79	19.29 (+88.5%)	+34.0%
+ Self-distillation	31.84	62.09 (+41.0%)	26.39	54.36 (+36.5%)	36.20	65.70 (+16.3%)	20.66	48.08 (+28.5%)	44.95	67.07 (+12.9%)	2.76	24.80 (+142.3%)	+46.2%

Table 2: Write-back under a common learning budget. Exact match (%) at inference depths 1 and 4 after learning on the same 1,024 GSM8K training questions. B_{50} and L_0 denote the non-looped and mean-depth-4 checkpoints. L_0 paths and teacher refer to its four-pass continuations and predictive supervision, respectively. Gold denotes human solutions. Bold marks the highest score within each checkpoint block at each depth.

Model	Supervision	Trainable parameters	Learning depth	GSM8K		GSM1K	
				1	4	1	4
Non-looped (B_{50})	–	–	–	44.05	–	39.83	–
+ SSD self-training	Own samples	Full	1	42.23	–	37.43	–
+ Gold-answer SFT	Gold	Full	1	44.73	–	40.17	–
+ Sequence KD	L_0 paths	Full	1	**47.46**	–	**40.83**	–
+ EOPD	L_0 teacher	Full	1	45.72	–	39.50	–
Mean depth 4 (L_0)	–	–	–	40.86	56.18	36.93	51.12
+ Gold-answer SFT	Gold	Core	1	38.89	55.04	33.03	50.79
	Gold	Core	4	38.51	56.10	34.77	51.20
	Gold	Full	1	38.36	54.28	35.10	46.89
	Gold	Full	4	40.03	55.50	33.03	47.72
+ DistiLLM-2	L_0 teacher	Core	1	42.46	58.30	36.93	50.71
+ GKD (on-policy)	L_0 teacher	Core	1	41.93	56.71	36.76	49.05
+ EOPD	L_0 teacher	Core	1	41.70	56.86	36.51	50.46
	L_0 teacher	Full	1	40.71	55.80	34.11	46.56
+ Preference write-back	L_0 paths	Core	1	**42.76**	**60.50**	**39.50**	**51.54**

on GSM1K and 2.50 points on SVAMP. The same ordering holds on GSM8K and GSM1K at the mean depths of 8 and 32. The mean-depth-32 checkpoint scores only 4.09% on GSM8K after one pass, yet learning through that pass raises its 32-pass accuracy from 57.01% to 62.09%.

Training through only the final pass of a four-pass execution reaches 56.63% on GSM8K in the mean-depth-4 model. Both this setting and one-pass learning differentiate a single application of the core. The forward depth used during learning therefore matters even when backpropagation covers only one recurrent pass.

One-pass learning places the loss directly on the first-pass state. When the loss is read out after four passes, the gradient

Figure 3: Recurrent depth weakens and redirects the gradient reaching the first-pass state. (a) Cosine similarity to the one-pass gradient and relative norm at the first-pass state, with the loss placed after one to four passes. (b) Each pass's share of the sum of per-pass parameter-gradient norms under four-pass training. Setup in Appendix D.1.

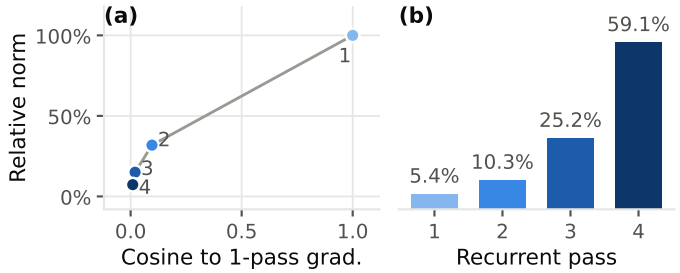

Table 3: Applying the shallow-learned block at later passes recovers nearly all of the four-pass gain. GSM8K exact match (%) at four-pass inference when each pass uses the original (✗) or the self-distilled (✓) shared block. Intervals and per-question counts are in Table 13.

Pass				
1	2	3	4	Accuracy (%)
✗	✗	✗	✗	56.18
✓	✗	✗	✗	56.25 (+0.1%)
✗	✓	✓	✓	62.47 (+11.2%)
✓	✓	✓	✓	62.70 (+11.6%)

Table 4: Single-pass learning outperforms full-depth learning on GSM8K and GSM1K at the three mean depths. Exact match (%) at depth 1 and at the mean depth m. Each block trains on the model's own greedy depth-m trajectories (*Teacher*) and varies the depth used during learning (*Training*). The $4^{\dagger}$ setting uses four passes with gradients through the last pass only. Changes are relative to the block's *Mean depth* row.

Model	Depth		GSM8K		GSM1K		SVAMP		Avg.
	Teacher	Training	Depth 1	Depth m	Depth 1	Depth m	Depth 1	Depth m	gain
Mean depth 4	–	–	40.86	56.18	36.93	51.12	52.40	63.30	–
	4	1	48.52	**62.70** (+11.6%)	44.65	**54.36** (+6.3%)	56.40	**64.70** (+2.2%)	+6.7%
+ Self-distillation	4	4	37.83	56.86 (+1.2%)	37.18	51.04 (−0.2%)	52.00	62.20 (−1.7%)	−0.2%
	4	$4^{\dagger}$	37.53	56.63 (+0.8%)	35.44	49.96 (−2.3%)	51.50	61.30 (−3.2%)	−1.5%
Mean depth 8	–	–	22.59	56.33	20.50	52.53	39.90	**66.00**	–
+ Self-distillation	8	1	37.38	**62.85** (+11.6%)	32.37	**56.02** (+6.6%)	43.60	65.20 (−1.2%)	+5.7%
	8	8	22.59	58.68 (+4.2%)	19.67	51.87 (−1.3%)	39.90	65.70 (−0.5%)	+0.8%
Mean depth 32	–	–	4.09	57.01	3.24	53.36	19.70	**67.70**	–
+ Self-distillation	32	1	31.84	**62.09** (+8.9%)	26.39	**54.36** (+1.9%)	36.20	65.70 (−3.0%)	+2.6%
	32	32	4.93	58.76 (+3.1%)	3.57	53.28 (−0.2%)	19.80	65.30 (−3.5%)	−0.2%

at that same state has 7.3% of the norm of the direct one-pass gradient and a cosine similarity of 0.009 (Figure 3a). At the parameter level, the first and fourth applications account for 5.4% and 59.1% of the sum of per-pass gradient norms under four-pass learning, respectively (Figure 3b). Appendix D.2 extends the performance comparison across inference depths and to a mean-depth-2 model.

3.3 Shallow-learned updates remain useful in later recurrence

Does an update learned through one pass improve later applications of the same shared operator? In the mean-depth-4 model, we choose at each of four recurrent passes whether to use the original or self-distilled shared block, while keeping the rest of the model fixed (Table 3). Using the updated block only on the first pass gives 56.25% GSM8K accuracy, close to the original model's 56.18%. Keeping the original block on pass 1 and using the updated block on passes 2–4 reaches 62.47%, compared with 62.70% when the update is used throughout. The later-pass replacement recovers 6.29 of the 6.52-point gain of the fully updated model. The update learned through one pass therefore remains useful on states reached later in recurrence.

Table 5: The advantage of shallow training depends on the recurrent recipe. Exact match (%) for matched mean-depth-4 models with and without input injection, self-distilled through one or four passes on their own greedy four-pass trajectories (*own*) or on those of the injected model (*injected*). Changes are relative to the *Mean depth 4* row with the same injection setting.

Model	Injection	Trajectories	Training depth	GSM8K		GSM1K		SVAMP		Avg.
				Depth 1	Depth m	Depth 1	Depth m	Depth 1	Depth m	gain
Mean depth 4	✓	–	–	40.86	56.18	36.93	51.12	52.40	63.30	–
+ Self-distillation	✓	own	1	48.52	**62.70** (+11.6%)	44.65	**54.36** (+6.3%)	56.40	**64.70** (+2.2%)	+6.7%
	✓	own	4	37.83	56.86 (+1.2%)	37.18	51.04 (−0.2%)	52.00	62.20 (−1.7%)	−0.2%
Mean depth 4	✗	–	–	43.90	50.80	38.42	45.31	54.90	58.70	–
+ Self-distillation	✗	own	1	47.16	52.39 (+3.1%)	42.90	45.31 (+0.0%)	54.60	56.50 (−3.7%)	−0.2%
	✗	own	4	44.58	50.72 (−0.1%)	40.41	46.72 (+3.1%)	53.90	59.00 (+0.5%)	+1.2%
	✗	injected	1	48.98	54.74 (+7.8%)	41.91	46.39 (+2.4%)	56.40	56.80 (−3.2%)	+2.3%
	✗	injected	4	45.26	54.81 (+7.9%)	39.75	47.63 (+5.1%)	56.20	59.60 (+1.5%)	+4.9%

Table 6: One-pass write-back across recurrent model families. Exact match (%) at depth 1 and the reported deeper depth m. Preference write-back uses 1,024 GSM8K training questions. Parentheses show relative changes from the original checkpoint at the same depth, computed from correct-answer counts. N/A marks a zero baseline. Nonzero baselines below 0.1% represent one correct answer. Bold marks the higher score in each comparison.

Model	m	Checkpoint	GSM8K				GSM1K			
			Depth 1		Depth m		Depth 1		Depth m	
Mean depth 4 (ours)	4	Original	40.86		56.18		36.93		51.12	
		Write-back	**42.76**	(+4.6%)	**60.50**	(+7.7%)	**39.50**	(+7.0%)	**51.54**	(+0.8%)
Retro-OLMo	16	Original	29.49		45.87		24.65		43.07	
		Write-back	**33.97**	(+15.2%)	**48.60**	(+6.0%)	**30.04**	(+21.9%)	**45.23**	(+5.0%)
Retro-LLaMA	8	Original	25.63		46.17		24.15		43.98	
		Write-back	**32.22**	(+25.7%)	**48.60**	(+5.3%)	**31.29**	(+29.6%)	**44.32**	(+0.8%)
Huginn (3.5B)	32	Original	0.00		32.37		0.08		28.46	
		Write-back	**1.59**	(N/A)	**32.68**	(+0.9%)	**1.58**	(+1,800.0%)	**28.80**	(+1.2%)
LoopUS	8	Original	0.08		**48.82**		0.08		48.22	
		Write-back	**55.65**	(+73,300.0%)	48.67	(−0.3%)	**53.53**	(+64,400.0%)	**48.46**	(+0.5%)
Ouro	4	Original	26.16		**79.68**		25.06		**77.43**	
		Write-back	**38.67**	(+47.8%)	77.18	(−3.1%)	**36.10**	(+44.0%)	74.69	(−3.5%)

3.4 Recurrent recipes shape subsequent learning

The advantage of shallow learning depends on the recurrent recipe. We compare mean-depth-4 models with and without input injection, starting from the same pretrained checkpoint and following matched mid-training procedures (Table 5). In the shared-target comparison, both models learn from the injected model's four-pass continuations.

With input injection, one-pass learning outperforms four-pass learning by 5.84 points on GSM8K and 3.32 points on GSM1K at four-pass inference. Without injection, four-pass learning gives comparable GSM8K accuracy (54.81% versus 54.74%) and higher accuracy on GSM1K (47.63% versus 46.39%) and SVAMP (59.60% versus 56.80%). How recurrent computation is developed therefore matters for how it can subsequently be improved.

Shallow-learned updates also benefit deeper inference in other recurrent models. We apply write-back to five public looped checkpoints using the same 1,024 GSM8K training questions. Each model learns from its own deeper continuations by updating the shared core through one pass (Table 6). Retro-OLMo and Retro-LLaMA (McLeish et al., 2026) gain 2.73 and 2.43 points on GSM8K at 16 and 8 passes, respectively, and improve on GSM1K. Together with our model and the 3.5B Huginn (Geiping et al., 2025), all four tested models with fixed input injection improve on both tasks at their deeper inference settings.

After write-back, one-pass inference in LoopUS (Park et al., 2026), which uses gated recurrent updates, exceeds its original eight-pass accuracy on both tasks. Ouro (Zhu et al., 2025), which does not repeatedly inject a fixed input

Table 7: Useful self-distillation signals arise across teacher depths and trajectory construction choices. Exact match (%) with one-pass training. Panels vary the teacher depth, the question pool and the target decoding (*Setting*). The last two panels use the mean-depth-4 model, and all other settings follow the default recipe. Changes are relative to the matching *Mean depth* row.

Model	Setting	GSM8K		GSM1K		SVAMP		Avg.
		Depth 1	Depth m	Depth 1	Depth m	Depth 1	Depth m	gain
Mean depth 4	–	40.86	56.18	36.93	51.12	52.40	63.30	–
+ Self-distillation	teacher depth 4	48.52	**62.70** (+11.6%)	44.65	**54.36** (+6.3%)	56.40	**64.70** (+2.2%)	+6.7%
	teacher depth 1	38.06	59.29 (+5.5%)	36.18	51.95 (+1.6%)	53.10	63.60 (+0.5%)	+2.5%
Mean depth 32	–	4.09	57.01	3.24	53.36	19.70	**67.70**	–
+ Self-distillation	teacher depth 32	31.84	62.09 (+8.9%)	26.39	**54.36** (+1.9%)	36.20	65.70 (−3.0%)	+2.6%
	teacher depth 8	27.90	**62.85** (+10.2%)	23.90	53.69 (+0.6%)	34.70	65.10 (−3.8%)	+2.3%
+ Self-distillation	GSM8K train (7,473)	47.92	61.26 (+9.0%)	40.66	52.95 (+3.6%)	54.90	64.60 (+2.1%)	+4.9%
	OpenMath (10,000)	48.52	62.62 (+11.5%)	42.16	53.94 (+5.5%)	52.40	61.20 (−3.3%)	+4.6%
	OpenMath, curated (10,000)	46.78	61.56 (+9.6%)	41.83	54.27 (+6.2%)	54.10	63.90 (+0.9%)	+5.6%
	OpenMath (30,000)	50.11	**63.38** (+12.8%)	43.49	53.61 (+4.9%)	54.60	63.40 (+0.2%)	+5.9%
	OpenMath (50,000)	48.52	62.70 (+11.6%)	44.65	**54.36** (+6.3%)	56.40	**64.70** (+2.2%)	+6.7%
+ Self-distillation	greedy	48.52	**62.62** (+11.5%)	42.16	**53.94** (+5.5%)	52.40	**61.20** (−3.3%)	+4.6%
	sampled, T=0.7, k=10	45.11	60.50 (+7.7%)	41.24	53.03 (+3.7%)	52.10	59.80 (−5.5%)	+2.0%

Table 8: Average non-mathematical performance remains stable after write-back. OLMES averages at the mean depth over 17 tasks in six groups. Parenthesized percentages show the average per-task relative change from the non-looped control within each group. Per-task results are in Table 14.

Model	Knowledge (2)	Exams (2)	Commonsense (4)	Science QA (3)	Reading (3)	Open QA (3)	All (17)
Non-looped control	29.63	35.69	69.65	61.26	61.42	46.56	53.94
Mean-depth-4 Looped LM	32.39 (+9.0%)	39.97 (+12.0%)	70.04 (+0.6%)	66.01 (+8.6%)	67.20 (+10.5%)	47.45 (+1.5%)	56.88 (+6.3%)
+ Self-distillation	32.72 (+10.5%)	40.07 (+12.3%)	69.87 (+0.3%)	65.91 (+8.6%)	67.95 (+12.3%)	47.59 (+2.4%)	57.02 (+6.9%)
Mean-depth-8 Looped LM	34.34 (+17.5%)	41.65 (+16.7%)	69.91 (+0.4%)	67.27 (+10.8%)	64.94 (+6.8%)	46.32 (−2.6%)	56.90 (+6.8%)
+ Self-distillation	34.36 (+17.0%)	41.62 (+16.6%)	70.28 (+1.0%)	67.47 (+11.1%)	65.26 (+7.7%)	46.43 (−2.1%)	57.09 (+7.1%)
Mean-depth-32 Looped LM	34.04 (+16.5%)	41.86 (+17.3%)	70.14 (+0.8%)	67.02 (+10.5%)	63.56 (+3.1%)	45.89 (−3.3%)	56.58 (+6.0%)
+ Self-distillation	33.79 (+14.9%)	41.01 (+14.9%)	70.18 (+0.8%)	66.60 (+9.8%)	63.18 (+3.1%)	46.37 (−2.0%)	56.40 (+5.6%)

representation, gains 12.51 and 11.04 points at one pass on GSM8K and GSM1K, respectively, with lower four-pass accuracy on both tasks.

3.5 Sources of depth-induced supervision

Useful supervision can emerge before the model reaches its standard inference depth. We compare teacher trajectories generated at different depths while keeping one-pass learning and deeper inference fixed (Table 7). For the mean-depth-32 model, learning from eight-pass trajectories raises 32-pass GSM8K accuracy from 57.01% to 62.85%, compared with 62.09% from 32-pass trajectories. Across the three benchmarks, the corresponding mean per-task relative gains are 2.3% and 2.6%. Teacher depth can therefore be chosen separately from inference depth.

On the same 10,000 OpenMath questions, greedy continuations give higher accuracy than truncated sampling on all three evaluated tasks, including 62.62% versus 60.50% on GSM8K (Table 7). Learning from either GSM8K or OpenMath question pools also improves GSM8K and GSM1K in the mean-depth-4 model. At a fixed pool size of 10,000 questions, selecting questions for broader coverage raises SVAMP accuracy from 61.20% with random selection to 63.90%, with both pools improving GSM8K.

3.6 Retention of non-mathematical capabilities

Write-back retains the overall non-mathematical performance of the looped checkpoints. We evaluate 17 OLMES tasks covering knowledge, exams, commonsense reasoning, science QA, reading, and open-domain QA (Table 8). At their mean depths, the three looped checkpoints score 56.58–56.90% on the 17-task average, compared with 53.94% for the non-looped control. After learning from mathematical continuations, this average changes by less than 0.2 percentage points for each model. The final models achieve mean per-task relative gains of 5.6–7.1% over the non-looped control.

4 Related Work

Recurrent depth and looped language models. Weight sharing across depth dates back to Universal Transformers (Dehghani et al., 2019) and ALBERT (Lan et al., 2020). Looped Transformers can emulate iterative algorithms (Giannou et al., 2023) and approach the reasoning ability of deeper models (Saunshi et al., 2025). Recent work pre-trains recurrent-depth language models at scale (Geiping et al., 2025; Zhu et al., 2025), converts pretrained models into looped form (Bae et al., 2025a; McLeish et al., 2026; Koishekenov et al., 2025), and adapts the number of passes to each token (Bae et al., 2025b) or to a compute budget (Jeddi et al., 2026). Building on these approaches, we develop looped language models through recurrent mid-training and study how the predictive behavior revealed by additional passes can inform subsequent training.

Self-distillation and model-generated supervision. Extra inference-time computation can improve reasoning (Snell et al., 2025). Several lines of work use model predictions as supervision, including sequence-level distillation (Kim and Rush, 2016) and born-again self-distillation (Furlanello et al., 2018). Related work internalizes chain-of-thought into latent reasoning (Hao et al., 2025; Shen et al., 2025) and trains on a model's own generations, often filtered for correctness or consistency (Zelikman et al., 2022; Huang et al., 2023; Singh et al., 2024) and sometimes without any filter (Zhang et al., 2026; Agarwal et al., 2025). We use sequence-level self-distillation on the model's own deeper continuations to update only the shared recurrent core.

Supervision across depth and training through partial computation. Intermediate depths are often supervised directly through losses on hidden layers (Lee et al., 2015), deep-to-shallow self-distillation (Zhang et al., 2019) or early-exit losses (Elhoushi et al., 2024). Weight-tied iterative models are often trained through part of their computation. Implicit models use implicit or approximate gradients, including approximations based on a single operator application (Bai et al., 2019; Geng et al., 2021; Fung et al., 2022). Recurrent networks that extrapolate to more iterations (Schwarzschild et al., 2021) have been trained by backpropagating through a random suffix (Bansal et al., 2022). Recurrent-depth language models also truncate backpropagation (Geiping et al., 2025), as in truncated backpropagation through time (Williams and Peng, 1990). Preference write-back separates the forward depth used for learning from the depths used to generate supervision and to run inference. Our pass-wise intervention shows that an update learned through one pass remains useful at later recurrent passes. Controlled comparisons further show that the advantage of single-pass learning over full-depth learning depends on the recurrent recipe.

5 Conclusion

Once a model has learned to benefit from recurrent depth, its own deeper executions can provide supervision for further learning. Our analysis shows how additional recurrence refines predictive preferences among already plausible tokens, and preference write-back demonstrates how the model can learn from these changes to improve its reasoning. An important finding is that an update learned through a single application of the shared recurrent operator can improve its subsequent applications. Using the updated core only at later passes recovers nearly the full gain in our four-pass GSM8K experiment. The advantage of shallow learning over full-depth learning also varies with the recurrent recipe, linking the development of recurrent computation to its subsequent improvement. More broadly, these findings open up the question of how the behavior that emerges as recurrent computation develops can guide its continued learning.

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 *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/5be69a584901a26c521c2b51e40a4c20-Abstract-Conference.html.

Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, and Hao Peng. The unreasonable effectiveness of entropy minimization in LLM reasoning. In *Advances in Neural Information Processing Systems*, volume 38, pages 107150–107180, 2025. URL https://openreview.net/forum?id=UfFTBEsLgI.

Sangmin Bae, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Seungyeon Kim, and Tal Schuster. Relaxed recursive transformers: Effective parameter sharing with layer-wise LoRA. In *International Conference on Learning Representations*, 2025a. URL https://openreview.net/forum?id=WwpYSOkkCt.

Sangmin Bae, Yujin Kim, Reza Bayat, Sungnyun Kim, Jiyoun Ha, Tal Schuster, Adam Fisch, Hrayr Harutyunyan, Ziwei Ji, Aaron Courville, and Se-Young Yun. Mixture-of-recursions: Learning dynamic recursive depths for adaptive token-level computation. In *Advances in Neural Information Processing Systems*, volume 38, pages 96572–96617, 2025b. URL https://openreview.net/forum?id=QuqsEIVWIG.

Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Arpit Bansal, Avi Schwarzschild, Eitan Borgnia, Zeyad Emam, Furong Huang, Micah Goldblum, and Tom Goldstein. End-to-end algorithm synthesis with recurrent networks: Extrapolation without overthinking. In *Advances in Neural Information Processing Systems*, volume 35, pages 20232–20242, 2022.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=HyzdRiR9Y7.

Mostafa Elhoushi, Akshat Shrivastava, Diana Liskovich, Basil Hosmer, Bram Wasti, Liangzhen Lai, Anas Mahmoud, Bilge Acun, Saurabh Agarwal, Ahmed Roman, Ahmed Aly, Beidi Chen, and Carole-Jean Wu. LayerSkip: Enabling early exit inference and self-speculative decoding. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12622–12642, 2024. doi: 10.18653/v1/2024.acl-long.681.

Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley Osher, and Wotao Yin. JFB: Jacobian-free backpropagation for implicit networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6648–6656, 2022. doi: 10.1609/aaai.v36i6.20619.

Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1607–1616. PMLR, 2018. URL https://proceedings.mlr.press/v80/furlanello18a.html.

Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach. In *Advances in Neural Information Processing Systems*, volume 38, pages 41340–41391, 2025. URL https://openreview.net/forum?id=S3GhJooWIC.

Zhengyang Geng, Xin-Yu Zhang, Shaojie Bai, Yisen Wang, and Zhouchen Lin. On training implicit models. In *Advances in Neural Information Processing Systems*, volume 34, pages 24247–24260, 2021.

Mor Geva, Avi Caciularu, Kevin Wang, and Yoav Goldberg. Transformer feed-forward layers build predictions by promoting concepts in the vocabulary space. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 30–45, 2022. doi: 10.18653/v1/2022.emnlp-main.3.

Angeliki Giannou, Shashank Rajput, Jy-Yong Sohn, Kangwook Lee, Jason D. Lee, and Dimitris Papailiopoulos. Looped transformers as programmable computers. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 11398–11442. PMLR, 2023.

Yuling Gu, Oyvind Tafjord, Bailey Kuehl, Dany Haddad, Jesse Dodge, and Hannaneh Hajishirzi. OLMES: A standard for language model evaluations. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5020–5048, 2025.

Danny Halawi, Jean-Stanislas Denain, and Jacob Steinhardt. Overthinking the truth: Understanding how language models process false demonstrations. In *International Conference on Learning Representations*, 2024.

Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training large language models to reason in a continuous latent space. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=Itxz7S4Ip3.

Benhao Huang, Zhengyang Geng, and J Zico Kolter. Equilibrium reasoners: Learning attractors enables scalable reasoning. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 46586–46616, 2026. URL https://proceedings.mlr.press/v306/huang26l.html.

Jiaxin Huang, Shixiang Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. Large language models can self-improve. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1051–1068, 2023. doi: 10.18653/v1/2023.emnlp-main.67. URL https://aclanthology.org/2023.emnlp-main.67/.

Stanisław Jastrzębski, Devansh Arpit, Nicolas Ballas, Vikas Verma, Tong Che, and Yoshua Bengio. Residual connections encourage iterative inference. In *International Conference on Learning Representations*, 2018.

Ahmadreza Jeddi, Marco Ciccone, and Babak Taati. LoopFormer: Elastic-depth looped transformers for latent reasoning via shortcut modulation. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=RzYXb5YWBs.

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 *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, pages 53926–53943, 2026. URL https://proceedings.mlr.press/v306/jin26e.html.

Alexia Jolicoeur-Martineau. Less is more: Recursive reasoning with tiny networks. *arXiv preprint arXiv:2510.04871*, 2025. URL https://arxiv.org/abs/2510.04871.

Yoon Kim and Alexander M. Rush. Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1317–1327, 2016. doi: 10.18653/v1/D16-1139. URL https://aclanthology.org/D16-1139/.

Jongwoo Ko, Tianyi Chen, Sungnyun Kim, Tianyu Ding, Luming Liang, Ilya Zharkov, and Se-Young Yun. DistiLLM-2: A contrastive approach boosts the distillation of LLMs. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 31044–31062, 2025. URL https://proceedings.mlr.press/v267/ko25a.html.

Harsh Kohli, Srinivasan Parthasarathy, Huan Sun, and Yuekun Yao. Loop, think, & generalize: Implicit reasoning in recurrent-depth transformers. In *Third Conference on Language Modeling*, 2026. URL https://proceedings.colm.cc/paper_files/paper/2026/hash/37050ebbbd7096719ab96cec19a4c69f-Abstract-Conference.html.

Yeskendir Koishekenov, Aldo Lipani, and Nicola Cancedda. Encode, think, decode: Scaling test-time reasoning with recursive latent thoughts. *arXiv preprint arXiv:2510.07358*, 2025.

Vedang Lad, Jin Hwa Lee, Wes Gurnee, and Max Tegmark. Remarkable robustness of LLMs: Stages of inference? In *Advances in Neural Information Processing Systems*, volume 38, pages 130050–130083, 2025.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite BERT for self-supervised learning of language representations. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=H1eA7AEtvS.

Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38 of *Proceedings of Machine Learning Research*, pages 562–570. PMLR, 2015.

Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning

to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167, 2017.

Sean McLeish, Ang Li, John Kirchenbauer, Dayal Singh Kalra, Brian R. Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Jonas Geiping, Tom Goldstein, and Micah Goldblum. Teaching pretrained language models to think deeper with retrofitted recurrence. In *Third Conference on Language Modeling*, 2026. URL https://proceedings.colm.cc/paper_files/paper/2026/hash/50ee6db59fca8643dc625829d4a0eab9-Abstract-Conference.html.

Shen-yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing English math word problem solvers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 975–984, 2020.

Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. GSM-Symbolic: Understanding the limitations of mathematical reasoning in large language models. In *The Thirteenth International Conference on Learning Representations*, 2025.

Taekhyun Park, Yongjae Lee, Dohee Kim, and Hyerim Bae. LoopUS: Recasting pretrained LLMs into looped latent refinement models. *arXiv preprint arXiv:2605.11011*, 2026. URL https://arxiv.org/abs/2605.11011.

Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP models really able to solve simple math word problems? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094, 2021.

Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, and Sashank J. Reddi. Reasoning with latent thoughts: On the power of looped transformers. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=din0lGfZFd.

Avi Schwarzschild, Eitan Borgnia, Arjun Gupta, Furong Huang, Uzi Vishkin, Micah Goldblum, and Tom Goldstein. Can you learn an algorithm? Generalizing from easy to hard problems with recurrent networks. In *Advances in Neural Information Processing Systems*, volume 34, pages 6695–6706, 2021.

Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. CODI: Compressing chain-of-thought into continuous space via self-distillation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 677–693, 2025. doi: 10.18653/v1/2025.emnlp-main.36. URL https://aclanthology.org/2025.emnlp-main.36/.

Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron T Parisi, Abhishek Kumar, Alexander A Alemi, Alex Rizkowsky, Azade Nova, Ben Adlam, Bernd Bohnet, Gamaleldin Fathy Elsayed, Hanie Sedghi, Igor Mordatch, Isabelle Simpson, Izzeddin Gur, Jasper Snoek, Jeffrey Pennington, Jiri Hron, Kathleen Kenealy, Kevin Swersky, Kshiteej Mahajan, Laura A Culp, Lechao Xiao, Maxwell Bileschi, Noah Constant, Roman Novak, Rosanne Liu, Tris Warkentin, Yamini Bansal, Ethan Dyer, Behnam Neyshabur, Jascha Sohl-Dickstein, and Noah Fiedel. Beyond human data: Scaling self-training for problem-solving with language models. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=lNAyUngGFK. Expert Certification.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=4FWAwZtd2n.

Team OLMo, Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Allyson Ettinger, Michal Guerquin, David Heineman, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Jake Poznanski, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2 OLMo 2 furious, 2025. URL https://arxiv.org/abs/2501.00656.

Shubham Toshniwal, Wei Du, Ivan Moshkov, Branislav Kisacanin, Alexan Ayrapetyan, and Igor Gitman. OpenMathInstruct-2: Accelerating AI for math with massive open-source instruction data. In *The Thirteenth International Conference on Learning Representations*, 2025.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny

Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837, 2022.

Ronald J. Williams and Jing Peng. An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural Computation*, 2(4):490–501, 1990. doi: 10.1162/neco.1990.2.4.490.

Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. STaR: Bootstrapping reasoning with reasoning. In *Advances in Neural Information Processing Systems*, volume 35, pages 15476–15488, 2022.

Hugh Zhang, Jeff Da, Dean Lee, Vaughn Robinson, Catherine Wu, Will Song, Tiffany Zhao, Pranav Raja, Charlotte Zhuang, Dylan Slack, Qin Lyu, Sean Hendryx, Russell Kaplan, Michele Lunati, and Summer Yue. A careful examination of large language model performance on grade school arithmetic. In *Advances in Neural Information Processing Systems, Datasets and Benchmarks Track*, volume 37, pages 46819–46836, 2024.

Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3713–3722, 2019.

Ruixiang Zhang, Richard He Bai, Huangjie Zheng, Navdeep Jaitly, Ronan Collobert, and Yizhe Zhang. Embarrassingly simple self-distillation improves code generation. *arXiv preprint arXiv:2604.01193*, 2026.

Rui-Jie Zhu, Zixuan Wang, Kai Hua, Tianyu Zhang, Ziniu Li, Haoran Que, Boyi Wei, Zixin Wen, Fan Yin, He Xing, Lu Li, Jiajun Shi, Kaijing Ma, Shanda Li, Taylor Kergan, Andrew Smith, Xingwei Qu, Mude Hui, Bohong Wu, Qiyang Min, Hongzhi Huang, Xun Zhou, Wei Ye, Jiaheng Liu, Jian Yang, Yunfeng Shi, Chenghua Lin, Enduo Zhao, Tianle Cai, Ge Zhang, Wenhao Huang, Yoshua Bengio, and Jason Eshraghian. Scaling latent reasoning via looped language models. *arXiv preprint arXiv:2510.25741*, 2025. URL `https://arxiv.org/abs/2510.25741`.

A Recurrence during mid-training

Before asking what additional recurrence contributes to a trained model, we check that recurrence pays off during mid-training itself. Figure 4 follows the Looped LMs of this paper on GSM8K and GSM-Symbolic as mid-training proceeds from the stage-1 checkpoint (Appendix C.2), as their gain over the non-looped control trained on the same 50B tokens. Every Looped LM finishes with a positive gain on both tasks, and on GSM8K each one crosses zero within 24B tokens, less than half of the control's training. Recurrent computation is therefore already an effective part of these models before any self-distillation; the remaining appendices and Sections 2 and 3 study what it adds and how to write it back into the shared block.

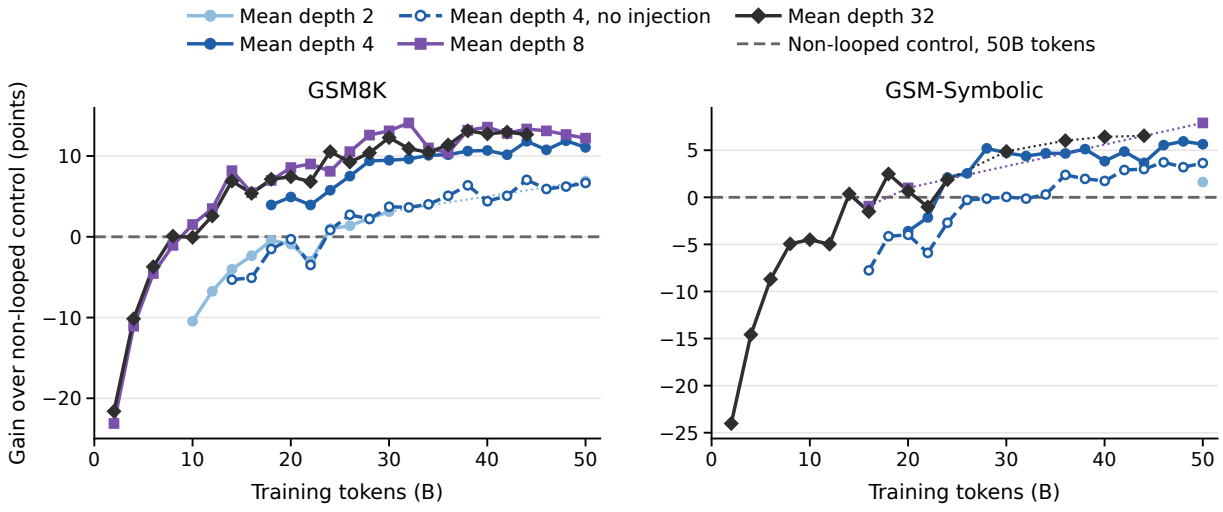

Figure 4: Gain in exact match (points) over the non-looped control after 50B tokens (dashed line at zero) during mid-training, for the Looped LMs of this paper, each evaluated at a fixed depth (its mean depth; eight passes for the mean-depth-32 model). Dotted segments join checkpoints more than 4B tokens apart.

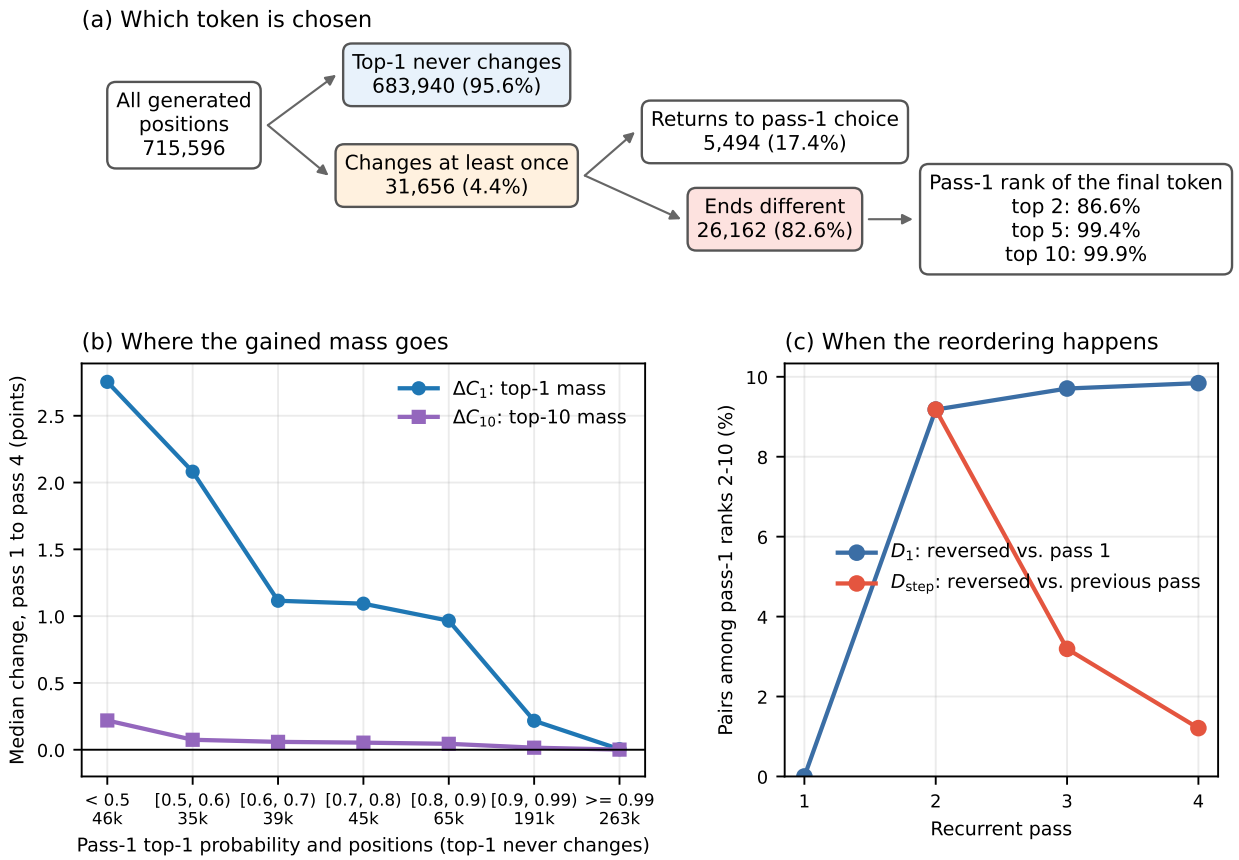

Figure 5: Effect of passes 2–4 on the readout distribution over 715,596 generated positions of GSM8K training problems. (a) Stability of the top-1 token across passes, and the pass-1 rank of the final token when it differs; each percentage is relative to the box it branches from. (b) Median change from q_1 to q_4 in top-1 probability (ΔC_1) and top-10 mass (ΔC_{10}), binned by the pass-1 top-1 probability, at positions whose top-1 never changes. (c) Pairwise reordering among the pass-1 ranks 2–10, relative to pass 1 (D_1) and to the preceding pass (D_{step}).

B Recurrence primarily reweights the top of the distribution

Reading out along a depth-4 trajectory. We analyze the mean-depth-4 Looped LM of Table 1. During an ordinary greedy generation at depth 4, we additionally decode the state s_k after each pass $k = 1, \ldots, 4$ with the same output stack and language-model head, giving the intermediate readouts

$$q_k(\cdot \mid x, \hat{y}_{<t}) = \text{softmax}(W[C(s_k)]_n)\,, \qquad k = 1, \ldots, 4.$$

The readouts are diagnostic only: each decodes from a copy of the output-stack cache, none feeds back into the recurrence, and tokens are selected from q_4 alone, so the generation is exactly the ordinary depth-4 generation. Because every q_k is conditioned on the depth-4 trajectory, and its output stack attends to the depth-4 states of earlier positions, it describes an intermediate state of that execution and differs in general from p_k, the distribution of an independent execution of depth k (Section 2.1). Figure 5 summarizes how the readouts change from pass to pass.

Recurrence primarily reshapes the head of the distribution. Additional passes mainly reshape the relative probabilities of candidates already near the top after the first pass, rather than promote previously low-ranked candidates. The effect varies in magnitude: at many positions the probabilities change while the ordering is preserved, at others nearby candidates are reordered, and only occasionally does the top-1 prediction change.

We apply this diagnostic to all 7,473 GSM8K training problems under greedy depth-4 generation, yielding 715,596 generated-token positions. (a) The top-1 prediction is stable at most positions: it is the same at all four passes at 95.6%

of positions, and when the final top-1 token differs from the pass-1 choice, it was already among the pass-1 top five 99.4% of the time (Figure 5a); (b) Where the top-1 remains fixed, recurrence mainly redistributes probability within the head of the distribution: the top-1 probability increases while the total top-10 mass changes little, with the largest shifts where the first pass is uncertain and the effect vanishing as it becomes confident (Figure 5b); (c) When recurrence does reorder nearby candidates, most of that change occurs early: pass 2 already accounts for 93% of the pairwise reordering among the pass-1 ranks 2–10 observed by pass 4, while later passes make progressively smaller changes (Figure 5c).

From readout to supervision. Within the four-pass generation, additional passes mainly redistribute probability among tokens that already rank highly in the first-pass readout of that same trajectory. Greedy decoding at the deeper execution turns this preference into a complete continuation, and these continuations are the self-distillation targets of Section 2.2.

Relation to layer-wise readouts. Decoding intermediate states with the model's own output head is a common probe of residual networks, whose layers refine a prediction iteratively (Jastrzębski et al., 2018): feed-forward updates promote candidate tokens in vocabulary space (Geva et al., 2022), late layers sharpen the next-token distribution (Lad et al., 2025), and intermediate predictions can be more accurate than final ones (Halawi et al., 2024). In a looped model every pass shares one state space and one operator, so the same output stack reads out each pass directly, without learned per-layer probes.

C Implementation details

This appendix gives the implementation behind Sections 2 and 3: the looped architecture (Appendix C.1), the mid-training that produces each Looped LM (Appendix C.2), the self-distillation recipe (Appendix C.3), and the evaluation protocol (Appendix C.4).

Notation in the tables. The tables name the three depths of Section 2.3 directly: *Teacher* is the teacher depth r_{T}, *Training* is the training depth r_{L}, and the *Depth 1* and *Mean depth* columns give the inference depth r_{E}. The mean depth m is the target of the depth curriculum during mid-training (Appendix C.2), and it is also the inference depth at which each Looped LM is used.

C.1 Architecture

Partition of OLMo-2 1B. OLMo-2 1B has 16 Transformer blocks of width $d = 2048$ with 16 attention heads (Team OLMo et al., 2025); its released configuration specifies a SwiGLU feed-forward layer of width 8,192 and a vocabulary of 100,352 tokens. The input stack P consists of blocks 1–7, the shared recurrent core F_θ, called the shared recurrent block in Section 3, of blocks 8–11, and the output stack C of blocks 12–16. F_θ is a single module invoked once per recurrent pass, so an execution of depth r applies the same 268,468,224 parameters r times; the complete model has 1.49B parameters.

Input injection and the initial state. The injection projection A_ψ is one bias-free linear map from $\mathbb{R}^{2d}$ to $\mathbb{R}^d$, applied at each position, with no gate, normalization or activation. The same input-stack output e is concatenated with the recurrent state at every pass. During training, the initial state s_0 is drawn from a truncated normal distribution with standard deviation 0.0087. At evaluation, s_0 is drawn with a seed derived from the prompt tokens, so executions of different depth on the same prompt start from the same state.

Model without input injection. The model without input injection in Table 5 removes both the concatenation and the separate initial state:

$$s_0 = e, \qquad s_i = F_\theta(s_{i-1}), \qquad i = 1, \ldots, r.$$

It is therefore a different recurrence rather than injection with a zeroed projection. Its partition, output stack and mid-training procedure match those of the injected mean-depth-4 model.

Table 9: Mid-training of the Looped LMs and the non-looped control. *Backpropagation* gives the passes that carry gradient at each training step.

Model	Input injection	Optimizer	Backpropagation
Non-looped control	✗	AdamW	single fixed pass
Mean-depth-2 Looped LM	✓	Muon	every pass
Mean-depth-4 Looped LM	✓	Muon	every pass
Mean-depth-4 Looped LM, no injection	✗	Muon	every pass
Mean-depth-8 Looped LM	✓	AdamW	last eight passes
Mean-depth-32 Looped LM	✓	AdamW	last eight passes

C.2 Mid-training of the Looped LMs

Common setup. Every Looped LM and the non-looped control start from the OLMo-2 1B stage-1 checkpoint and continue causal language-model training on the Dolmino mixture (Team OLMo et al., 2025) for up to 50B tokens, with context length 4,096, a global batch of 2,097,152 tokens (512 sequences), the same data order, BF16 precision and fully sharded data parallelism. Table 9 lists the optimizer and the gradient path of each model. All models use a peak learning rate of 7.4×10^{-5}, decayed linearly to zero over 23,842 steps without warmup, $\beta = (0.9, 0.95)$, weight decay 0.033, gradient clipping at norm 1.0 and a z-loss coefficient of 10^{-5}. Muon updates the matrix parameters with its learning rate rescaled to the update scale of AdamW, and AdamW updates the embeddings, normalization parameters and language-model head.

Depth sampling and curriculum. At each training step, one depth r is drawn for the whole batch from the Poisson-lognormal distribution of Geiping et al. (2025), with log-normal scale $\sigma = 0.5$ and parameterized by the current curriculum mean μ_t of Section 2.1. Over the first 75% of the T training steps, μ_t rises from 1 to the target mean depth m along a one-minus-square-root schedule, and it is held at m afterwards:

$$\mu_t = \max\Big(1,\ \Big\lceil m\left(1 - \sqrt{1 - \min(t/t_{\mathrm{w}},\, 1)}\right)\Big\rceil\Big), \qquad t_{\mathrm{w}} = 0.75\,T.$$

Gradient path through the recurrence. For the mean-depth-8 and mean-depth-32 models, all passes before the last eight of a sampled depth run without gradient, so backpropagation covers at most eight passes, following Geiping et al. (2025); this bounds the length of the gradient path through the recurrence and keeps optimization stable at large depth. Because e is injected at every pass, the input stack still receives gradient through the differentiated passes.

Non-looped control. The non-looped control of Tables 1 and 8 is trained with the same data and setup, but with a single fixed recurrent pass and no input injection. With $s_0 = e$ and one pass, it computes $C(F_\theta(P(x)))$, which is the original 16-block OLMo-2 architecture. Appendix A follows the accuracy of the Looped LMs through this stage.

C.3 Self-distillation

Teacher trajectories. The teacher is the frozen Looped LM before the update. For each training question it generates one continuation at teacher depth r_{T} by greedy decoding, using the fixed eight-shot chain-of-thought prompt of OLMES for GSM8K (Gu et al., 2025; Team OLMo et al., 2025) and at most 512 new tokens. Generation ends at the first stop string or end-of-sequence token. The generated token ids, including the stop token, are the target without further editing; a continuation that reaches the token limit is kept as generated. Targets are generated once and reused in every epoch, and the reference answer of the training question is never used. The sampled targets in Table 7 apply temperature $T = 0.7$ and then top-k truncation with $k = 10$ at each step.

Objective and updated parameters. A training sequence is the eight-shot prompt followed by the target continuation. The loss is the completion-only next-token cross-entropy of Section 2.3, normalized by the total number of target tokens in the batch; no distillation term on the teacher probabilities is added. Only the 32 parameter tensors of F_θ are updated. The embeddings, input stack, injection projection A_ψ, output stack and language-model head are frozen, leaving 18% of the parameters trainable.

Algorithm 1 Preference write-back by self-distillation

Input: Looped LM with shared core θ_0 and remaining parameters ϕ; prompts $\{x^{(j)}\}_{j=1}^{N}$;
teacher depth r_{T}; training depth r_{L}; number of epochs E

Expose: the frozen model generates targets at teacher depth
for $j = 1, \ldots, N$ **do**
$\hat{y}^{(j)} \leftarrow$ greedy continuation of $x^{(j)}$ under $p_{r_{\mathrm{T}}}^{(0)}$, at most 512 tokens, kept unfiltered
$\mathcal{D}_{r_{\mathrm{T}}} \leftarrow \{(x^{(j)}, \hat{y}^{(j)})\}_{j=1}^{N}$
Incorporate: update only the shared core through r_{L} passes
$\theta \leftarrow \theta_0$
for epoch $= 1, \ldots, E$ **do**
for each batch $B \subset \mathcal{D}_{r_{\mathrm{T}}}$ of 16 sequences **do**
$\theta \leftarrow \mathrm{AdamW}\left(\theta, \nabla_\theta \mathcal{L}_{r_{\mathrm{L}}}(\theta; B)\right)$ *completion-only cross-entropy; ϕ frozen*
Reuse: the updated core serves every pass at any inference depth
return $\theta_1 \leftarrow \theta$, used with ϕ at inference depth r_{E}

Table 10: Self-distillation settings shared by all runs, and the size of the OpenMath question pools.

Setting	Value
Updated parameters	shared recurrent block F_θ (268M of 1.49B)
Objective	completion-only next-token cross-entropy
Optimizer	AdamW, $\beta = (0.9, 0.95)$, no weight decay
Learning rate	3×10^{-6}; 10 linear warmup steps, then cosine decay
Batch	16 sequences
Gradient clipping	global norm 1.0
Epochs	3 (OpenMath pools), 6 (GSM8K training questions)
Precision	BF16 computation, FP32 parameters
Targets	one teacher continuation per question, at most 512 tokens
OpenMath question pools (targets of the mean-depth-4 model)	
10,000 questions	992,459 target tokens per epoch
10,000 questions, curated	1,019,832 target tokens per epoch
30,000 questions	2,973,623 target tokens per epoch
50,000 questions	4,953,856 target tokens per epoch

Algorithm 1 summarizes the procedure.

Training depth. Training at depth r_{L} runs r_{L} recurrent passes forward, reads the loss out after the last pass, and backpropagates through all r_{L} passes; because F_θ is shared, the gradient accumulates contributions from each of its uses. The $4^\dagger$ row of Table 4 runs the first three passes without gradient and backpropagates through the fourth only. For the depth-8 and depth-32 training rows of Table 4, activations are recomputed pass by pass to fit in memory, which leaves the gradient unchanged.

Optimization. Table 10 lists the optimization settings, which are shared by every self-distillation run in the paper. Runs differ only in the teacher depth, training depth, question pool and target decoding stated in the corresponding table.

Question pools. The default pool contains question text from the `augmented_gsm8k` portion of OpenMathInstruct-2 (Toshniwal et al., 2025). Only the problem statements are read; the released solutions and answers are not used. Questions are normalized for Unicode and whitespace, deduplicated, and screened against the GSM8K test, GSM1K and SVAMP questions and the development questions of Appendix C.4. A lexical audit against these 4,214 questions (word five-gram Jaccard similarity of at least 0.8, or numeric-template similarity of at least 0.9) finds no near match in the final pools. The 10,000-, 30,000- and 50,000-question pools are nested. The curated 10,000-question pool of Table 7 is selected from the 50,000-question pool using the question text alone: it matches the length distribution of the random

Table 11: Mathematical evaluation sets. All prompts are fixed chain-of-thought demonstrations shared by every question of a set.

Benchmark	Questions	Shots	Token limit
GSM8K test (Cobbe et al., 2021)	1,319	8	512
GSM1K (Zhang et al., 2024)	1,205	8	512
SVAMP (Patel et al., 2021)	1,000	8	512
GSM-Symbolic, main (Mirzadeh et al., 2025)	5,000	8	1,024
ASDiv (Miao et al., 2020)	2,305	8	512
AQuA-RAT (Ling et al., 2017)	254	4	512

10,000-question pool, favors rare problem types, and keeps at most one question from each cluster of near-duplicates, and it reuses the targets already generated for those questions. The GSM8K pool consists of all 7,473 training questions (Cobbe et al., 2021).

Checkpoint selection. For the OpenMath runs in Table 1, we selected checkpoints by development-set accuracy at each model's mean depth. For the full GSM8K training run in Table 7, we selected epoch 2 by mean relative gain in test accuracy at four passes.

C.4 Evaluation protocol

Mathematical reasoning. Table 11 lists the mathematical evaluation sets. Every model is evaluated with greedy decoding in BF16 at depth 1 and at its mean depth. Generation stops at the first stop string or end-of-sequence token, and a response that reaches the token limit remains in the denominator. For GSM8K, GSM1K and SVAMP the answer is the number after the last `####` if present and otherwise the last number in the response, normalized for commas, trailing zeros and sign; GSM-Symbolic uses the answer extraction of the OLMES GSM8K task; for ASDiv it is the number after "The final answer is", falling back to the last number; for AQuA-RAT it is the last explicitly labelled option letter. Accuracy is exact match of the extracted answer with the reference answer.

GSM8K and GSM-Symbolic use the OLMES eight-shot prompt (Gu et al., 2025; Team OLMo et al., 2025); GSM1K and SVAMP use the eight arithmetic chain-of-thought demonstrations of Wei et al. (2022), and ASDiv uses the same eight demonstrations in a problem–solution format. AQuA-RAT uses four fixed demonstrations with their rationales, drawn from its development split. The non-mathematical tasks of Tables 8 and 14 are evaluated with the OLMES harness at each model's mean depth (depth 1 for the non-looped control), with the shots and metrics listed in Table 14.

Summary columns. Relative changes in the tables are per-task changes divided by the reference value, the Looped LM before self-distillation or, where stated, the non-looped control. The *Avg.* columns of Tables 1 and 4–7 are unweighted means of these per-task relative changes. In Table 8, each group score is the mean of its task scores and its percentage the mean of the per-task relative changes.

Development set and intervals. We used a fixed development set of 690 questions from the full GSM8K training pool and excluded those questions from the OpenMath pools. Intervals are 95% question-paired bootstrap intervals over 10,000 resamples of the evaluation questions. Table 12 gives them for the comparisons of Table 1: every Looped LM is above the non-looped control on all six tasks, and the gains of self-distillation on GSM8K and GSM-Symbolic exclude zero for all three models.

C.5 Compute

Mid-training runs on NVIDIA A100 80GB GPUs, and self-distillation, including the generation of teacher trajectories, on NVIDIA A800 80GB GPUs. Self-distillation is small next to mid-training. Even at teacher depth 32, generating the trajectories for 10,000 questions takes 18 GPU-hours, and three epochs of training on them at depth 1 take 2 GPU-hours.

Table 12: Paired differences in exact match (points) behind Table 1, at each model's mean depth, with 95% question-paired bootstrap intervals. The non-looped control is evaluated at depth 1.

Comparison	GSM8K	GSM1K	SVAMP	GSM-Symbolic	ASDiv	AQuA-RAT
+ Self-distillation minus Looped LM						
Mean depth 4	+6.5 [4.2, 8.8]	+3.2 [1.1, 5.5]	+1.4 [−1.1, 3.9]	+4.1 [3.0, 5.3]	+1.7 [0.5, 3.0]	+1.2 [−3.9, 6.3]
Mean depth 8	+6.5 [4.4, 8.6]	+3.5 [1.2, 5.8]	−0.8 [−3.0, 1.4]	+3.4 [2.3, 4.5]	+1.7 [0.5, 3.0]	−1.2 [−6.7, 4.3]
Mean depth 32	+5.1 [2.7, 7.6]	+1.0 [−1.6, 3.6]	−2.0 [−4.6, 0.6]	+3.2 [1.9, 4.4]	–	+5.5 [0.0, 11.0]
Looped LM minus non-looped control						
Mean depth 4	+12.1 [9.6, 14.6]	+11.3 [8.7, 13.8]	+6.8 [4.0, 9.7]	+6.1 [4.9, 7.3]	+5.4 [3.9, 6.9]	+6.7 [2.0, 11.8]
Mean depth 8	+12.3 [9.8, 14.8]	+12.7 [10.2, 15.2]	+9.5 [6.7, 12.3]	+7.6 [6.4, 8.9]	+6.8 [5.3, 8.3]	+13.4 [7.9, 19.3]
Mean depth 32	+13.0 [10.4, 15.5]	+13.5 [11.0, 16.2]	+11.2 [8.3, 14.2]	+7.5 [6.2, 8.7]	+7.2 [5.7, 8.7]	+9.1 [3.5, 15.0]

D Further evidence on training depth

This appendix supports Sections 3.2 and 3.3. It details the gradient measurements of Figure 3 (Appendix D.1), follows the two training depths across inference depths and model families (Appendix D.2), and gives the full record of the pass-wise block replacement of Table 3 (Appendix D.3).

D.1 Recurrent pass-gradient diagnostic

Setup. Both panels of Figure 3 use the mean-depth-4 Looped LM before self-distillation, the point from which training at depth 1 and training at depth 4 start with the same parameters and the same targets. We take 64 fixed teacher trajectories from the OpenMath question pool (6,892 target tokens) and evaluate the completion-only loss of Section 2.3 in FP32.

Panel (a): gradient at the same first-pass state. For each trajectory, let s_1 be the state after the first recurrent pass and L_k the loss read out after k passes. We compute $g_k = \partial L_k / \partial s_1$ over the full state tensor for $k = 1, \ldots, 4$; the state s_1 is verified to be identical across the four depths. The panel plots the cosine between g_k and the one-pass gradient g_1, and the ratio $\|g_k\|/\|g_1\|$, each computed per trajectory and averaged over the 64 trajectories; error bars are 95% bootstrap intervals over trajectories. Most of the change occurs at the first additional pass: with two passes, the cosine is already 0.10 and the norm 32% of the one-pass gradient.

Panel (b): parameter gradient by pass. To attribute the gradient of four-pass training to individual uses of the shared block, we untie the block into four copies with identical weights, one per pass. Their gradients $g^{(1)}, \ldots, g^{(4)}$ sum to the gradient of the shared parameters, with a maximum relative error of 1.3×10^{-6}. The panel shows the share $\|g^{(i)}\| / \sum_j \|g^{(j)}\|$ for each pass, computed on four groups of 16 trajectories, the batch size of self-distillation, and averaged over the groups.

Together, the two panels suggest that training at depth 1 applies the direct signal to the first use of the shared block, whereas training at depth 4 reaches it only as a small, rotated remnant and updates the block mainly through its uses on late recurrent states.

D.2 Training depth across inference depths and models

Figure 6 follows each injected Looped LM and its two self-distilled models across the inference depths we evaluated, for mean depths 2, 4, 8 and 32. On GSM8K and GSM1K, training at depth 1 improves the model at every inference depth in all four families, and the gain from recurrence is retained: the mean-depth-4 model self-distilled at depth 1 rises from 48.52% on GSM8K at one pass to 59.06% at two and 62.70% at four. The gain at one pass is largest where that execution starts weakest: the mean-depth-32 model moves from 4.09% to 31.84% at one pass while also improving at 8, 16 and 32 passes. Training on the same targets through the full mean depth ends below training at depth 1 at the mean depth on both sets in every family.

The mean-depth-2 model repeats the controlled comparison of Table 4 on a second injected checkpoint, mid-trained toward mean depth 2: the teacher generates one greedy trajectory at depth 2 for each of the 10,000 OpenMath questions,

and the same targets train the shared block through one pass or through two, with all other settings as in Appendix C.3. At the deployed depth, training at depth 1 is ahead by 4.2, 2.3 and 3.8 points on GSM8K, GSM1K and SVAMP, and every 95% CI of these paired differences excludes zero. Training at depth 1 is also the cheaper update: on the same trajectories and hardware, three epochs of the mean-depth-4 model take 0.72 GPU-hours at depth 1 against 1.14 at depth 4, with 14% less peak memory.

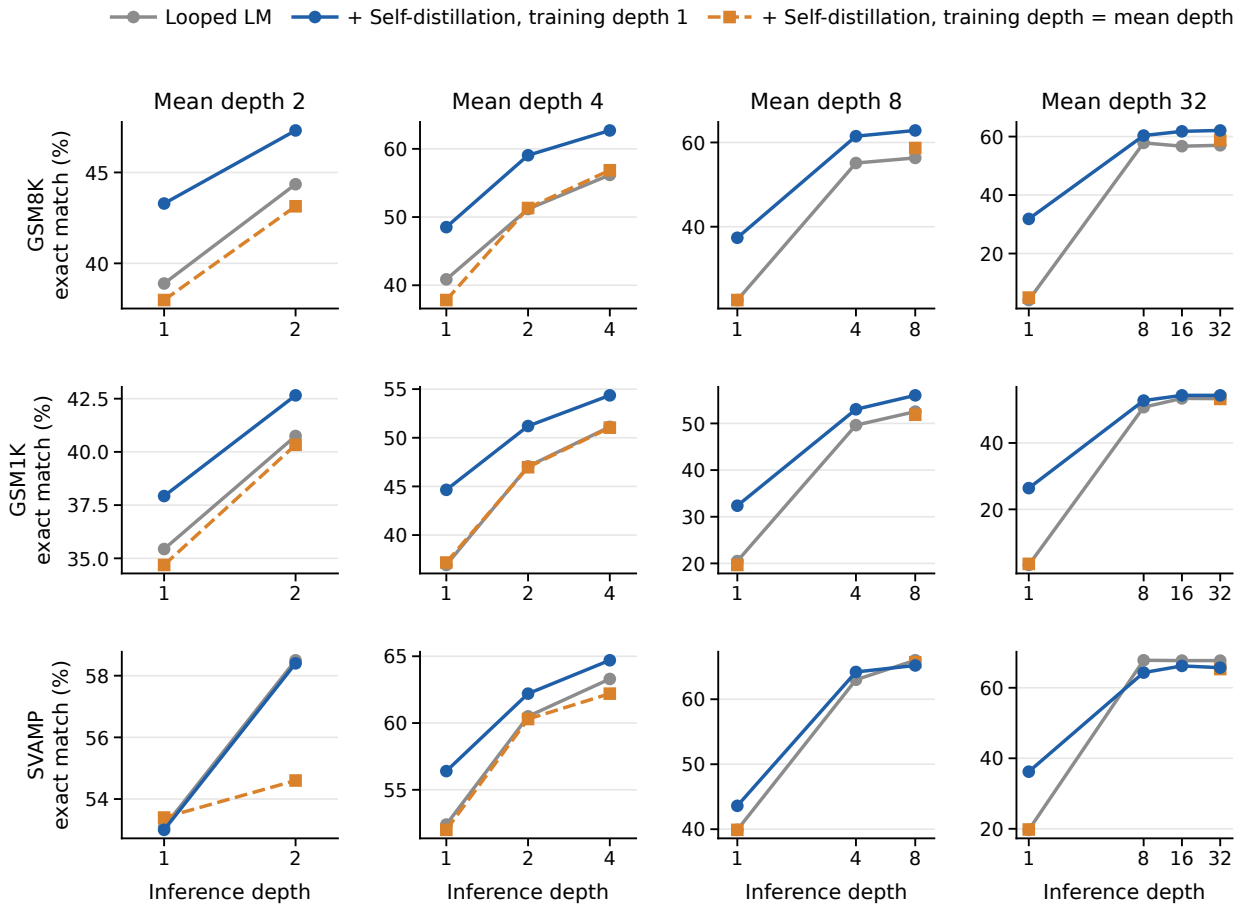

Figure 6: Exact match (%) across inference depths for four injected Looped LMs (columns) on GSM8K, GSM1K and SVAMP (rows). Each Looped LM is self-distilled from its own greedy mean-depth trajectories through one pass or through its full mean depth; lines connect adjacent evaluated depths. Inference depth is on a logarithmic axis, and the vertical range differs between panels. The values at depth 1 and at the mean depth of the mean-depth-4, 8 and 32 models are those of Table 4.

D.3 Pass-wise block replacement

The intervention of Table 3 executes the mean-depth-4 Looped LM for four passes and chooses, for each pass, either the original shared block or the block of the self-distilled model of Table 1. All other parameters are shared by the two models; only the 32 tensors of the recurrent block differ. Each pattern is a complete execution: generation and the cache are rebuilt under that pattern, so the comparison locates where the updated block is needed rather than splitting the gain into additive per-pass contributions. On eight check examples, the two pure patterns reproduce ordinary inference of the two models exactly, in generated tokens, next-token logits and cache contents.

Table 13 gives the record of the four patterns on GSM8K. Using the updated block on passes 2–4 recovers 97% of the gain obtained by using it throughout. Using it only on the first pass corrected 39 answers and introduced 38 errors, for a net gain of one correct answer out of 1,319.

Table 13: Pass-wise replacement of the shared block on GSM8K (1,319 questions) at four-pass inference. ✓: the self-distilled block is used at that pass; ✗: the original block. Change is in points relative to the original model; repaired / broken count the questions that become correct / incorrect.

Pass							
1	2	3	4	Accuracy (%)	Change	95% CI	Repaired / broken
✗	✗	✗	✗	56.18	–	–	–
✓	✗	✗	✗	56.25	+0.1	[−1.2, 1.4]	39 / 38
✗	✓	✓	✓	62.47	+6.3	[4.0, 8.6]	164 / 81
✓	✓	✓	✓	62.70	+6.5	[4.2, 8.8]	165 / 79

E Retention beyond mathematics

Table 14 gives the per-task scores behind Table 8. Self-distillation leaves individual tasks close to where they were: 47 of the 51 task scores of the three self-distilled models move by less than one point, and the mean-depth-8 model stays within one point on every task. The advantage that looping gives over the non-looped control is preserved as well: after self-distillation, each Looped LM remains above the control on at least 13 of the 17 tasks.

Table 14: Scores (%) on the 17 non-mathematical OLMES tasks, grouped as in Table 8: the non-looped control at depth 1 and, for each Looped LM, the model before and after self-distillation at its mean depth. The self-distilled models are those of Table 1; shots and metrics follow the OLMES task definitions (Appendix C.4), and AGIEval is evaluated in both its 3-shot and its 1-shot configuration.

			Control	Mean-depth-4		Mean-depth-8		Mean-depth-32	
Task	Shots	Metric	Depth 1	Looped LM	+ Self-dist.	Looped LM	+ Self-dist.	Looped LM	+ Self-dist.
Knowledge									
MMLU	5	accuracy	43.92	48.17	48.44	50.14	50.43	49.69	49.66
MMLU-Pro	5	accuracy	15.33	16.61	16.99	18.54	18.28	18.38	17.91
Exams									
AGIEval	3	accuracy	36.10	40.02	40.07	42.67	42.69	42.58	41.83
AGIEval	1	accuracy	35.28	39.92	40.07	40.63	40.54	41.14	40.20
Commonsense									
CommonsenseQA	5	accuracy	66.50	66.42	66.18	67.08	67.49	66.67	67.08
HellaSwag	5	accuracy	70.00	70.60	70.70	70.00	70.10	70.10	70.10
PIQA	5	accuracy	75.50	75.20	74.90	74.30	75.20	74.90	74.80
WinoGrande	5	accuracy	66.61	67.96	67.72	68.27	68.35	68.90	68.75
Science QA									
ARC-Challenge	5	accuracy	51.79	59.13	59.13	60.32	59.81	61.26	59.81
ARC-Easy	5	accuracy	78.60	80.50	79.60	81.70	82.00	80.80	80.40
OpenBookQA	5	accuracy	53.40	58.40	59.00	59.80	60.60	59.00	59.60
Reading									
CoQA	0	F1	67.14	74.50	74.10	68.05	68.28	66.59	64.44
SQuAD	5	F1	82.41	87.31	88.28	87.56	87.45	88.40	88.16
DROP	5	F1	34.72	39.80	41.46	39.21	40.04	35.69	36.94
Open QA									
Jeopardy	5	F1	63.81	65.42	65.54	64.83	64.88	64.59	65.50
NaturalQuestions	5	F1	20.84	20.83	21.48	18.77	19.02	18.77	19.28
TriviaQA	5	F1	55.02	56.11	55.75	55.36	55.39	54.32	54.34