arXiv:2611.31094v2 [cs.CL] 24 Nov 2026

Unlocking the Power of Loops in Pretrained LLMs with Loop Residuals

Xuanlei Zhao [1] **Hailun Xu** **Kai Wang** [1] **Yang Luo** [1] **Yingxuan Ren** [1] **Zirui Zhu** [1] **Yuxuan Lou** [1] **Kun Xu**
Yan Wang [2] **Kanchan Sarkar** **Yang You** [1]

github.com/NUS-HPC-AI-Lab/LoopRes

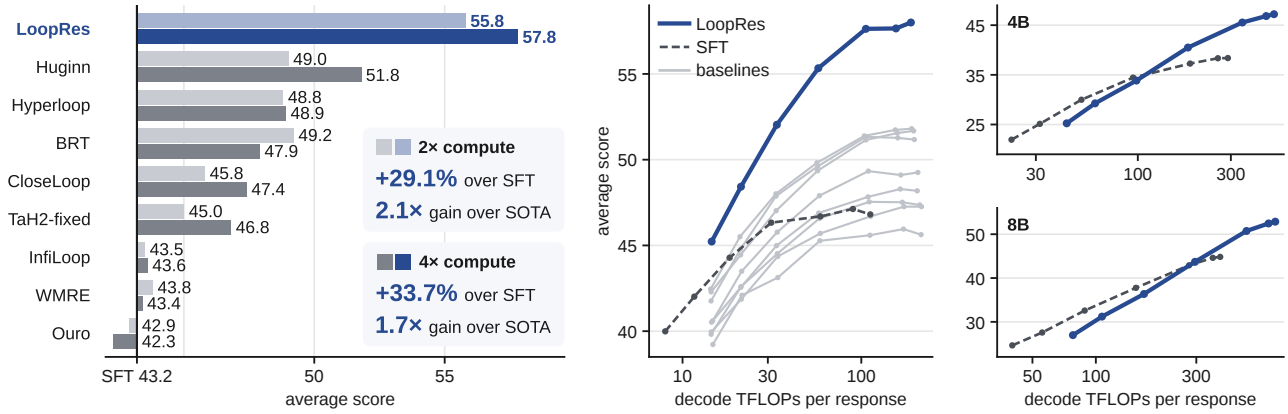

(a) 1.7B, 2× and 4× compute. *(b)* 1.7B scaling, 2× compute. *(c)* 4B/8B scaling, 2× compute.

Figure 1. **Average scores across compute budgets and backbones.** Average scores on seven math and code benchmarks, where SFT is the same model finetuned without looping. (a) Qwen3-1.7B-Base at 2× (light) and 4× (dark) compute, with SOTA the best baseline at the same compute. (b) Scores against decode compute as the thinking budget grows from 1K to 30K tokens. (c) The same sweep on Qwen3-4B-Base and Qwen3-8B-Base over five math and code benchmarks.

Abstract

Looped Transformers iterate a shared block of layers to increase computation without adding parameters, offering stronger performance and test-time scaling. However, this power has so far required training from scratch and yielded only modest gains for LLMs. Finetuning existing models into loops is far cheaper, yet also brings only marginal gains over SFT. We propose LoopRes, an effective, scalable, and simple approach that turns loops into significant and consistent gains for pretrained LLMs. Building on Attention Residuals, LoopRes lets each looped layer attend directly to the outputs of earlier loops. Since these outputs are iterates of the same layers and become harder to tell apart, we give each a trainable loop-aware key that encodes its structural identity. Attending to all earlier loops also lets layers skip intermediate ones, so we restrict each layer to a residual window of the latest W loops, forcing every loop to build incrementally on recent ones. Across 6 pretrained LLMs from 3 model families and 9 benchmarks, LoopRes delivers significant and consistent gains: at 2× compute, it improves Qwen3-1.7B-Base by 29.1% over SFT, more than doubling the gain of the state-of-the-art baseline. The gains hold on larger models up to 30B and across all three families. More importantly, from 1.7B to 8B, LoopRes scales more steeply than SFT as compute grows, reaching a higher ceiling while needing less compute on long outputs. These results establish LoopRes as an effective looping approach, delivering both better performance and better test-time scaling.

1. Introduction

Looped Transformers reuse a shared block of layers multiple times in a single forward pass, achieving deeper computation without adding parameters (Dehghani et al., 2019; Giannou et al., 2023). Depth scales with the number of loops rather than the number of layers, making loops a simple lever for both stronger performance and test-time

[1]National University of Singapore [2]Independent Researcher. Correspondence to: Xuanlei Zhao <xuanlei.zhao@u.nus.edu>, Yang You <yangyou@nus.edu.sg>.

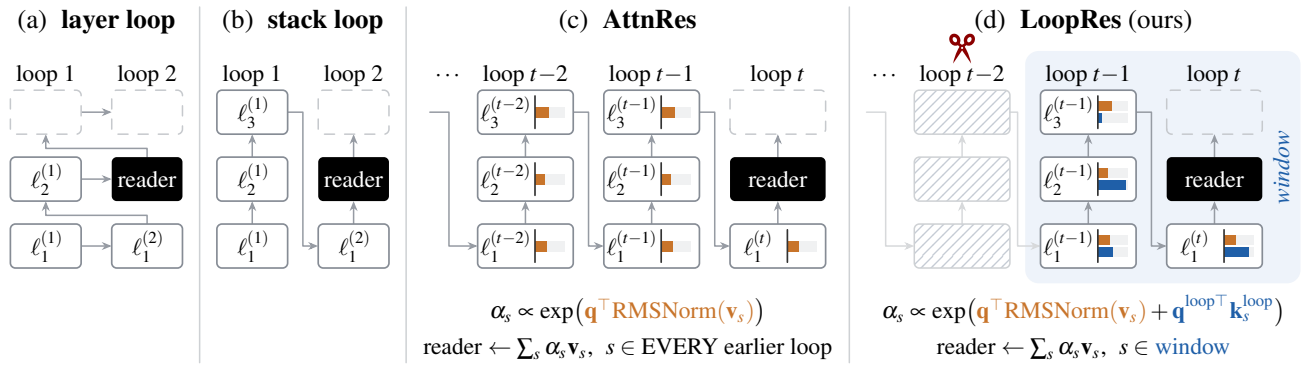

Figure 2. **How a looped layer takes the outputs of earlier layers as input.** Each panel diagrams three layers running for up to three loops, highlighting one layer's input, where $\ell_i^{(t)}$ denotes layer i of loop t and dashed layers have not yet run. Loops before $t-2$ are omitted in (c) and (d). (a) A layer loop iterates each layer through every loop, passing only the previous iteration forward. (b) A stack loop feeds only the last layer of each loop into the first of the next. (c) AttnRes attends to every earlier loop output with content-based weights. (d) LoopRes attends only to a window of the latest W loops (shaded, $W=1$ here) and adds a loop-aware key alongside the content key. Each bar marks the magnitude of a score term.

compute scaling. Yet turning this promise into significant and consistent gains for LLMs remains an open challenge.

One line of work pretrains looped LLMs from scratch (Geiping et al., 2025; Zhu et al., 2025). However, the gains remain modest or even negative at comparable cost (Zeitoun et al., 2026; Wang et al., 2026). Given the expense of pretraining, this direction offers limited returns.

A more promising approach finetunes pretrained LLMs into loops (Marchenko et al., 2026; You et al., 2026), reusing their learned capabilities at a fraction of the pretraining cost. As illustrated in Figure 2a,b, although these designs vary in their looping pattern (Zhu et al., 2025; You et al., 2026; Geiping et al., 2025; Zeitoun et al., 2026; Wang et al., 2026; Zhuang et al., 2026; Marques et al., 2026), they all compress the outputs of earlier loops into a narrow pathway. This limits what later loops can build on, yielding only marginal gains and leaving the potential of loops largely unrealized.

A natural remedy is to let each looped layer attend directly to the outputs of earlier loops. Attention Residuals (AttnRes) (Kimi Team et al., 2026) offer exactly this: each layer builds its input by attending to all earlier layer outputs. However, AttnRes is designed for standard Transformers, where the sources differ in content and are fixed in number. Our key insight is that loops break both assumptions: the sources are iterates of the same layers, so they become similar in content and grow in number.

This causes three problems. (i) **Similar iterates.** The outputs are iterates of the same layers and grow harder to tell apart as loops deepen. (ii) **Early convergence.** The AttnRes weights converge early in finetuning, and because the keys come from similar outputs, the model cannot learn to tell sources apart before the pattern is set. (iii) **Growing sources.** Each loop adds new sources to the full attention window, so layers can shortcut to earlier loops, skipping intermediate ones.

Based on these observations, we propose Loop Residuals (LoopRes), an effective, scalable, and simple approach that turns loops into significant and consistent gains for pretrained LLMs. LoopRes builds on AttnRes and introduces two improvements. (i) It assigns each source a trainable, loop-aware key alongside its content key, encoding the structural relationship of loops and layers. Because these keys are directly optimized, the model can learn to differentiate sources before the weights converge. (ii) It restricts each layer to a residual window of the latest W loops. This eliminates shortcuts to earlier loops, forces every loop to build incrementally on recent ones, and keeps the overhead constant as loops deepen.

Across 6 pretrained LLMs from 3 model families and 9 math and code benchmarks, LoopRes delivers significant and consistent gains. On Qwen3-1.7B-Base at 2× compute, it improves the average score by 29.1% over SFT, more than doubling the gain of the SOTA baseline, and its gain continues to grow as loops deepen. The gains are consistent: LoopRes improves by 10.0% to 37.7% on the larger Qwen3-{4B,8B,30B-A3B}-Base, as well as on Qwen3.5 and Llama-3.2 from different families. Notably, LoopRes on the 4B backbone even surpasses SFT on the 8B one, showing that loops can bridge the gap to larger models.

More importantly, LoopRes achieves better scaling at test time: from 1.7B to 8B, its score rises with a steeper slope than SFT as compute grows. It thus reaches a higher ceiling, significantly exceeding the best score of SFT. At the same time, it is more efficient on long outputs, reaching the same score at a lower compute cost. Together, these results establish loop depth as an axis of test-time scaling, benefiting both inference and post-training rollouts. We hope LoopRes serves as a simple and effective foundation

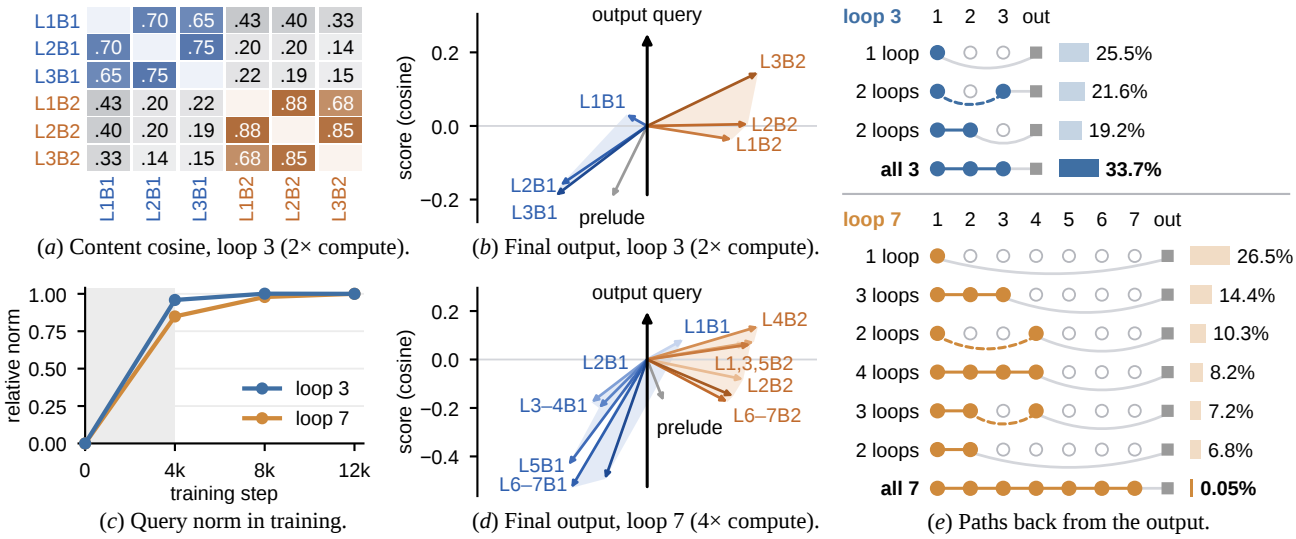

(*a*) Content cosine, loop 3 (2× compute). (*b*) Final output, loop 3 (2× compute). (*c*) Query norm in training. (*d*) Final output, loop 7 (4× compute). (*e*) Paths back from the output.

Figure 3. **Visualization of AttnRes in a looped model.** Huginn + AttnRes on Qwen3-1.7B-Base at 2× compute (L=3) and 4× compute (L=7). LtBj denotes the output of block j in loop t. (a) Cosine similarity between block outputs across all loops. (b,d) Attention scores from the output query to each source at L=3 and L=7. Arrow height marks the score. (c) Output query norm over finetuning steps, normalized by its converged value. (e) Weight share on each path from the output back to the first loop, obtained by multiplying the attention weights along the path.

for scaling LLMs along this axis.

2. Realizing the Promise of Looping

2.1. Preliminaries

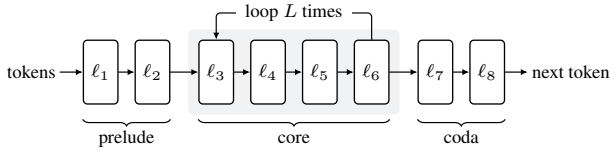

Figure 4. **The looped architecture of Huginn.** The layers of an LLM are split into a prelude, a core, and a coda. The prelude and the coda run once, and the core runs L times.

Looped language models. We adopt the looped architecture of Huginn (Geiping et al., 2025) as shown in Figure 4 and build it from a pretrained LLM by splitting its layers into three consecutive parts: a prelude, a core, and a coda. The prelude runs once on the token embeddings and produces $\boldsymbol{p}$. The core then runs L times, and we call each run a loop, indexed by $t \in \{1, \ldots, L\}$. The coda runs once on the output of loop L to predict the next token, and the model is finetuned with the cross-entropy loss on this prediction.

Attention Residuals. A standard Transformer feeds each sub-layer the residual stream. Attention Residuals (AttnRes) (Kimi Team et al., 2026) replace the stream with a learned input. In its block form, the layers are grouped into blocks, and each sub-layer takes as input a set of sources $\mathcal{S}$: the embedding, the summed output of every completed block, and the partial sum of the current block. AttnRes weighs these sources by a softmax attention, and the input of the sub-layer is formulated as

$$\boldsymbol{x} = \sum_{s \in \mathcal{S}} \alpha_s \boldsymbol{v}_s, \qquad \alpha_s = \frac{\exp(\boldsymbol{q}^\top \boldsymbol{k}_s)}{\sum_{s' \in \mathcal{S}} \exp(\boldsymbol{q}^\top \boldsymbol{k}_{s'})}, \quad (1)$$

where $\boldsymbol{q}$ is a learnable query of the sub-layer initialized to zero, $\boldsymbol{v}_s$ is a source, and $\boldsymbol{k}_s = \mathrm{RMSNorm}(\boldsymbol{v}_s)$ is its key. We call $\boldsymbol{k}_s$ a content key, since it is derived from the source and makes the weights content-dependent.

2.2. Applying Attention Residuals to Loops

We apply AttnRes only to the core of the looped model. The prelude and the coda keep the standard residual stream. Within the core, we group its layers into B blocks. We denote the output of block i in loop t by $\boldsymbol{b}_i^{(t)}$, and the running partial sum within block i by $\boldsymbol{u}_i^{(t)}$. A sub-layer in block i of loop t then takes as input

$$\mathcal{S}_i^{(t)} = \{\boldsymbol{p}, \underbrace{\boldsymbol{b}_1^{(1)}, \ldots, \boldsymbol{b}_B^{(t-1)}}_{\text{previous loops}}, \underbrace{\boldsymbol{b}_1^{(t)}, \ldots, \boldsymbol{b}_{i-1}^{(t)}, \boldsymbol{u}_i^{(t)}}_{\text{loop } t}\}, \quad (2)$$

where $\boldsymbol{p}$ is the prelude output. Each sub-layer forms its input from these sources by Equation 1, using its own query in each loop. The output of loop t is formed the same way, with its own query.

2.3. Challenges of Looping with Attention Residuals

Similar iterates. As shown in Figure 3a, the iterates of the same block are far more similar to each other than to other blocks. Because the content keys are derived from these sources, similar iterates receive similar scores. With three

loops, the output query in Figure 3b clearly separates the two blocks, but the iterates within each block are hard to tell apart. With seven loops in Figure 3d, this worsens, as many iterates collapse onto nearly the same score.

Early convergence. As shown in Figure 3c, the query norm reaches most of its final value within the first third of finetuning and then levels off. We hypothesize that this commits the model to its early connection pattern, since content keys change only indirectly, through the hidden states.

Growing sources. The source set grows with every loop, and AttnRes attends to all of them, so layers can skip intermediate loops. As shown in Figure 3e, the weight on the path through every loop falls from 33.7% with three loops to 0.05% with seven loops, and the most likely paths never reach the later loops, suggesting that deeper loops add little effective depth.

2.4. LoopRes

Building on the analysis above, we propose LoopRes, an effective, scalable, and simple approach that turns loops into significant and consistent gains for pretrained LLMs. As illustrated in Figure 2d, LoopRes keeps the input-dependent mechanism of AttnRes and introduces two improvements: (i) It assigns each source a loop-aware key alongside its content key, which resolves the identity of each iterate. (ii) It restricts each layer to a residual window of the previous W loops, which bounds the number of sources as the loop deepens. Figure 9 provides PyTorch-style pseudo code.

Loop-aware keys. We assign each source a loop-aware key $\boldsymbol{k}_s^{\text{loop}}$ alongside its content key $\boldsymbol{k}_s = \text{RMSNorm}(\boldsymbol{v}_s)$: a trainable vector per block per loop, initialized from $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$. Each sub-layer uses the two keys with two queries, and the input weights in Equation 1 become

$$\alpha_s = \frac{\exp(\boldsymbol{q}^{\text{loop}\top}\boldsymbol{k}_s^{\text{loop}} + \boldsymbol{q}^\top \boldsymbol{k}_s)}{\sum_{s' \in \mathcal{S}} \exp(\boldsymbol{q}^{\text{loop}\top}\boldsymbol{k}_{s'}^{\text{loop}} + \boldsymbol{q}^\top \boldsymbol{k}_{s'})}, \tag{3}$$

where $\boldsymbol{q}^{\text{loop}}$ is the loop-aware query and $\boldsymbol{q}$ the content query, both initialized to zero, so that each sub-layer starts from uniform weights over its sources.

While the content key determines how relevant each source is to the current token, the loop-aware key brings two benefits: (i) It encodes the structure of the loop, distinguishing the iterates of the same layers by their coordinates. (ii) It is optimized directly, so the connection pattern no longer adapts only through the hidden states.

Residual window. LoopRes restricts the residual sources to a window of the previous W loops. This brings three benefits: (i) It prevents layers from skipping to distant loops, forcing each loop to build on the latest loops. (ii) It keeps the softmax from shrinking the weight of the current loop as sources accumulate. (iii) It bounds each input to $(W+1)B+1$ sources, thus limiting compute, memory, and communication overhead to W loops.

Efficiency. A naive implementation of Equation 3 loads all queries and keys, then writes the weighted sum for every sub-layer. Following the design of FlashAttention (Dao, 2024), we fuse these steps into a single Triton (Tillet et al., 2019) kernel that loads each source once from HBM and writes only the result back, significantly reducing IO cost and peak memory.

3. Experiments

We validate the performance of LoopRes and analyze its design through extensive experiments. In particular, we investigate the following questions:

- Can LoopRes consistently outperform existing looped designs? (Table 1, Figure 5)
- Does LoopRes scale to deeper loops, larger backbones, and other model families? (Figure 6, Table 3, Table 4)
- How much does each component contribute? (Table 5)
- What patterns does LoopRes learn? (Figure 7, Figure 8)

3.1. Setups

Models and baselines. We use Qwen3-{1.7B, 4B, 8B, 30B-A3B}-Base (Yang et al., 2025), Qwen3.5-2B-Base (Qwen Team, 2026), and Llama-3.2-1B (Grattafiori et al., 2024) as backbones, with Qwen3-1.7B-Base unless otherwise stated. We compare LoopRes with SFT, the backbone finetuned without looping, and with 8 looped baselines. Five of them loop all layers: Ouro (Zhu et al., 2025), TaH2-fixed (You et al., 2026), CloseLoop (Marchenko et al., 2026), InfiLoop (Li et al., 2026), and BRT (Marques et al., 2026). The other 3 loop the middle half of the layers as LoopRes does: Huginn (Geiping et al., 2025), Hyperloop (Zeitoun et al., 2026), and WMRE (Zhuang et al., 2026). All baselines are compared at $2\times$ and $4\times$ compute, corresponding to $L = 2, 4$ for all-layer baselines and $L = 3, 7$ for half-layer ones. LoopRes uses $W = 1$ at both compute budgets.

Training setup. We finetune on math and code questions from AM-Qwen3-Distilled (Tian et al., 2025) and UltraData-SFT-2605 (OpenBMB, 2026). Qwen3-1.7B-Base trains on replies regenerated by Qwen3-8B, totalling 572K conversations and 3.1B tokens; all other backbones train on the original replies of Qwen3-235B-A22B. All models share the same hyperparameters: 12,000 steps with a learning rate of 3×10^{-5}, as detailed in Section A.4.

Evaluation setup. We report the loss on 1,493 held-out conversations and evaluate on seven benchmarks for backbones below 4B: MATH500 (Lightman et al., 2024), OlympiadBench (He et al., 2024), AMC23, AIME25, and AIME26

Table 1. **Comparison with baselines at 2× and 4× compute on Qwen3-1.7B-Base.** Compute is measured relative to SFT. HE+ denotes HumanEval+, and both code benchmarks report the plus tests. *avg.* averages the seven math and code scores. Each Δ column gives a method's change relative to SFT. Within each compute budget, **bold** marks the best and underline the second best.

method	math					code		avg.	Δ	loss	Δ
	AIME25	AIME26	AMC23	Olympiad	MATH500	HE+	MBPP+				
SFT	14.6	11.7	48.6	40.7	73.4	57.3	56.2	43.2	–	0.312	–
2× compute											
Huginn	19.6	14.4	54.1	46.5	80.4	64.2	64.0	49.0	13.5%↑	0.303	2.8%↓
Ouro	12.5	9.2	47.8	40.4	74.6	59.1	56.6	42.9	0.8%↓	0.315	1.1%↑
Hyperloop	18.3	16.7	54.4	47.1	77.6	66.9	60.4	48.8	12.9%↑	0.303	2.9%↓
TaH2-fixed	15.8	12.5	50.0	41.3	75.8	62.3	57.0	45.0	4.0%↑	0.306	1.7%↓
WMRE	16.2	10.6	51.2	40.7	74.0	60.1	54.1	43.8	1.5%↑	0.311	0.3%↓
CloseLoop	16.5	13.3	50.0	44.1	78.8	61.4	56.7	45.8	6.0%↑	0.309	1.0%↓
InfiLoop	14.4	10.8	49.4	42.4	73.0	58.4	56.3	43.5	0.7%↑	0.312	0.0%↑
BRT	21.7	15.4	53.6	45.0	79.0	66.6	63.2	49.2	13.9%↑	0.302	3.0%↓
LoopRes	**27.9**	**19.4**	**64.7**	**53.8**	**84.0**	**71.8**	**68.9**	**55.8**	**29.1%↑**	**0.292**	**6.3%↓**
4× compute											
Huginn	24.4	18.3	58.8	50.4	80.0	66.2	64.3	51.8	19.8%↑	0.299	4.1%↓
Ouro	13.5	11.2	45.6	40.4	75.4	56.9	52.8	42.3	2.2%↓	0.321	3.1%↑
Hyperloop	18.5	15.8	54.7	46.7	79.6	63.4	63.6	48.9	13.2%↑	0.303	2.8%↓
TaH2-fixed	18.8	12.7	52.2	43.1	77.0	64.0	59.8	46.8	8.3%↑	0.303	2.6%↓
WMRE	12.7	11.2	49.5	41.9	73.4	60.1	55.2	43.4	0.5%↑	0.312	0.1%↑
CloseLoop	18.3	13.8	55.5	44.1	77.0	64.8	58.3	47.4	9.7%↑	0.307	1.4%↓
InfiLoop	13.5	10.8	49.8	40.9	76.4	57.9	55.8	43.6	0.9%↑	0.312	0.0%↑
BRT	21.2	13.8	54.7	44.7	77.0	64.2	59.4	47.9	10.7%↑	0.305	2.0%↓
LoopRes	**30.2**	**24.6**	**67.5**	**52.9**	**84.2**	**75.3**	**69.8**	**57.8**	**33.7%↑**	**0.291**	**6.8%↓**

Table 2. **Training overhead of LoopRes on Qwen3-1.7B-Base.** We report the training step time and the peak activation memory, measured at 16k sequence length.

variant	step time (ms)		peak mem. (GiB)	
plain residual	1886.3	–	8.56	–
LoopRes	1976.6	+4.8%	8.62	+0.7%
w/o loop-aware key	1968.9	+4.4%	8.62	+0.7%
w/o fused kernel	2922.3	+54.9%	12.63	+47.5%

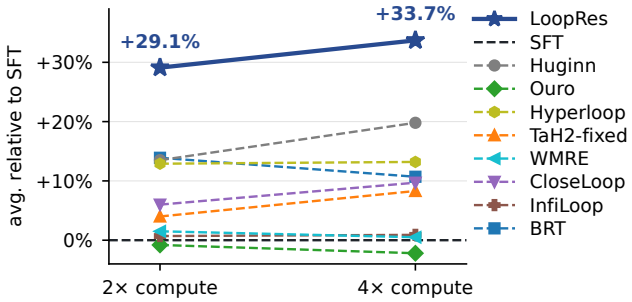

Figure 5. **Average score relative to SFT at 2× and 4× compute.** Each line connects a design's change in the average of the seven math and code scores in Table 1, on Qwen3-1.7B-Base.

in math, and HumanEval+ and MBPP+ (Liu et al., 2023) in code. Since the 4B and larger backbones saturate the easier benchmarks, we evaluate them on five harder ones instead: AIME25, AIME26, OlympiadBench, and IMO-AnswerBench (Luong et al., 2025) in math, and LiveCodeBench v6 (Jain et al., 2025) in code. We sample with temperature 0.6, top-p 0.95, and top-k 20 up to 32K tokens, and report pass@1 averaged over approximately 1,000 samples per benchmark.

3.2. Main Results

Performance results. Table 1 compares LoopRes, SFT, and 8 baselines across 7 benchmarks and the loss at 2× and 4× compute. To ensure a fair comparison, we reimplement every baseline on the same backbone and finetune it on the same data with the same hyperparameters, as detailed in Section A.2.

Based on the results, we make the following observations: (i) LoopRes achieves the best score on every benchmark at both compute budgets, more than doubling the gain of the strongest baseline at 2× compute. (ii) The relative gains are largest on the hardest benchmarks (e.g., AIME), likely because harder questions benefit most from the extra computation per token that deeper loops provide.

Loop gains. Figure 5 illustrates how the gain over SFT grows as compute doubles. Most baselines gain little or even regress, and the few that continue to improve remain

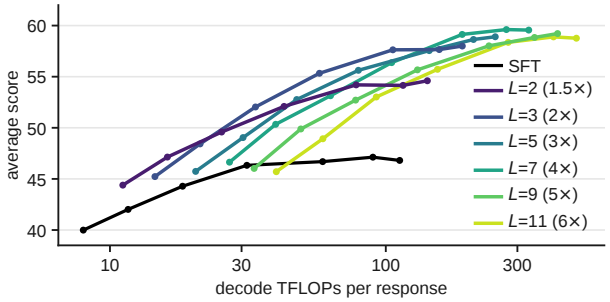

Figure 6. **Test-time scaling across loop depths.** Average score on seven benchmarks against decode compute as the thinking budget grows from 1K to 30K tokens, on Qwen3-1.7B-Base.

Table 3. **Results on larger backbones at 2× compute.** On AIME25 and AIME26, OlympiadBench, IMO-AnswerBench, and LiveCodeBenchv6.

method	A25	A26	Oly	IMO	LCB	avg.	Δ
Qwen3-4B-Base							
SFT	35.4	38.3	61.6	20.5	27.4	36.6	–
LoopRes	**47.1**	**51.0**	**68.3**	**29.6**	**32.9**	**45.8**	+9.2
Qwen3-8B-Base							
SFT	39.6	46.0	65.5	30.9	32.6	42.9	–
LoopRes	**51.5**	**62.9**	**72.3**	**35.5**	**35.1**	**51.5**	+8.6
Qwen3-30B-A3B-Base							
SFT	59.6	66.3	72.9	38.6	38.3	55.1	–
LoopRes	**65.4**	**72.5**	**76.4**	**44.9**	**44.0**	**60.6**	+5.5

well below LoopRes. In contrast, LoopRes converts the extra loops into further gains from an already high score, showing that its advantage holds as loops deepen.

Overhead. Table 2 compares the training cost of LoopRes with a plain residual baseline that loops the same layers. LoopRes adds little step time and almost no peak activation memory, and the loop-aware key itself is nearly free. Since the residual window limits the number of sources, the cost of forming each input stays bounded as loops deepen. We fuse these operations into a single kernel; without this fusion, the naive PyTorch implementation is far more expensive.

3.3. Scalability

Scaling to deeper loops. We train LoopRes with different numbers of loops on Qwen3-1.7B-Base and evaluate each across thinking budgets. As shown in Figure 6, the score improves as loops deepen and gradually saturates beyond $L = 5$. The best trade-off between score and compute is at $L = 3$.

Scaling to larger backbones. We apply LoopRes to Qwen3-{4B,8B}-Base and the MoE Qwen3-30B-A3B-Base at 2× compute. As shown in Table 3, LoopRes improves every benchmark on all three backbones. Figure 1 and Figure 11 further show that its score rises with a clearly steeper slope than SFT as decode FLOPs grow, especially once reasoning begins to saturate. This advantage is more evident on smaller models. Notably, LoopRes on the 4B backbone even surpasses SFT on the 8B one in average score.

Table 4. **Results on Llama-3.2-1B and Qwen3.5-2B-Base at 2× compute.** *avg.* averages the seven math and code scores, and Δ gives the change of LoopRes relative to SFT in points. Table 11 reports the per-benchmark scores.

backbone	avg.		Δ
	SFT	**LoopRes**	
Llama-3.2-1B	14.3	**19.6**	+5.4
Qwen3.5-2B-Base	52.1	**64.1**	+12.0

Table 5. **Ablation study of LoopRes on Qwen3-1.7B-Base.** Each group varies one component from LoopRes (default in gray).

variant	2× compute		4× compute	
	avg.	Δ	avg.	Δ
LoopRes	55.8	–	57.8	–
Huginn + AttnRes	53.9	−1.9	55.0	−2.8
Residual window (LoopRes: $W = 1$)				
$W = 2$	55.2	−0.6	57.5	−0.3
$W = 3$	–	–	57.7	−0.1
Key (LoopRes: loop-aware + content)				
loop-aware only	54.9	−0.9	57.5	−0.3
content only	54.9	−0.9	55.8	−2.0

Generalizing across model families. We further apply LoopRes to Llama-3.2-1B and Qwen3.5-2B-Base, whose architectures differ from Qwen3. As shown in Table 4 and Table 11, LoopRes improves every benchmark on both. Across all backbones, LoopRes consistently delivers gains, demonstrating its generalizability.

3.4. Ablation Study

We ablate the components of LoopRes on Qwen3-1.7B-Base at both compute budgets in Table 5.

Naive AttnRes. We first apply AttnRes to the core of Huginn without loop-aware keys or the residual window. This naive combination performs clearly worse than LoopRes at both compute budgets, with the gap widening as loops deepen.

Residual window. We vary the window size W, which bounds how many previous loops each layer takes as input. A larger window reduces the average score at both compute budgets. Thus, we set $W = 1$, as it achieves the best score at the lowest cost. However, adding this window alone to naive AttnRes (i.e., the content-only variant) brings only a slight improvement and remains worse than LoopRes.

Loop-aware keys. We then examine the contribution of

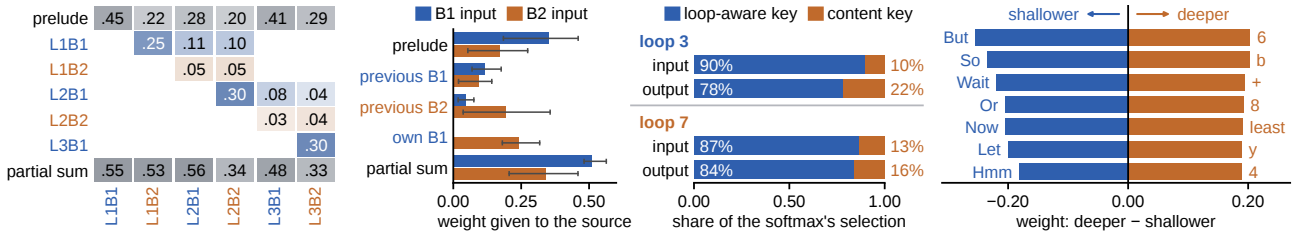

(*a*) What each block reads, loop 3. (*b*) What B1 and B2 read, loop 7. (*c*) Selection by key type. (*d*) What the content key reacts to, loop 7.

Figure 7. **What LoopRes learns.** LoopRes on Qwen3-1.7B-Base, averaged over 128 samples. LtBj denotes the output of block j in loop t. (a) Source weights for each block at L=3. (b) Source weights of B1 and B2 at L=7, averaged over loops with min–max bars. (c) Contribution of the loop-aware key and the content key to source selection, measured by Shapley values. (d) Tokens ranked by whether their content key steers B2 toward its shallower or deeper output at L=7.

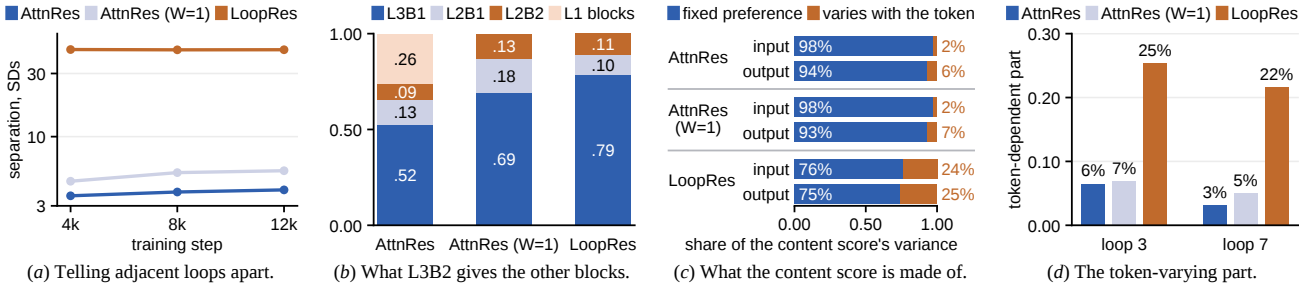

(*a*) Telling adjacent loops apart. (*b*) What L3B2 gives the other blocks. (*c*) What the content score is made of. (*d*) The token-varying part.

Figure 8. **How LoopRes differs from AttnRes.** AttnRes, and AttnRes with W=1 and LoopRes on Qwen3-1.7B-Base over 128 samples at L=3. (a) Output difference between adjacent loops for each block, measured in standard deviations across tokens. (b) Source weights of L3B2 over earlier blocks, renormalized to exclude the prelude and partial sum. (c) Content-score variance across sources, decomposed into a fixed-preference component and a token-varying component. (d) Token-varying fraction of the content-score variance at loop outputs, for L=3 and L=7.

the loop-aware and content keys. Either key alone performs worse than the two combined at both compute budgets. As loops deepen, dropping the loop-aware key hurts more, while using it alone approaches LoopRes. This is likely because iterates become harder to tell apart by content as loops deepen, making structural information more important.

3.5. Analysis

What LoopRes learns. We visualize how LoopRes weighs its sources, revealing three patterns.

- **Iterative loop updates.** In Figure 7a,b, each layer balances its input across the prelude, the loop outputs, and the partial sum. Within the loop outputs, the first block mainly updates its own output from the previous loop, with a small share from the other block of that loop. The second block on average relies most on the first block of the current loop, indicating that LoopRes learns the iterative structure of the loop.
- **Dominant loop-aware keys.** In Figure 7c, the loop-aware key accounts for most of the selection at both compute budgets, with the content key taking a smaller share.
- **Adaptive content keys.** Despite its smaller share, the content key lets each token choose its own depth. In Figure 7d, simple tokens such as “But” and “Wait” favor shallower outputs, while difficult tokens such as digits and math symbols favor deeper ones. This suggests that the content key recognizes the difficulty of each token and allocates computation accordingly.

How LoopRes differs from AttnRes. We further compare LoopRes with AttnRes, with and without the residual window, revealing three differences.

- **Distinguishable loop iterates.** In Figure 8a, AttnRes struggles to tell adjacent loops apart even with the window, while LoopRes separates them clearly from early in finetuning.
- **Recency-focused inputs.** In Figure 8b, the last block of AttnRes gives a sizable share of its weight to the first loop. The window removes these shortcuts, and LoopRes further concentrates the weight on the preceding block.
- **Less is more for content keys.** In Figure 8c,d, the content score of AttnRes is almost a fixed preference over sources and barely varies with the token. Although the content key takes a smaller share in LoopRes, its score varies far more with the token, even as loops deepen, likely because the loop-aware key takes over the fixed preference.

4. Related Work

Residual connections. Residual connections (He et al., 2016) add each layer's output to a shared stream with fixed unit weights and remain standard in modern LLMs. DenseNet (Huang et al., 2017) instead feeds each layer the outputs of all preceding layers. mHC (Xie et al., 2025) expands the residual stream into parallel streams and constrains their mixing to preserve the identity mapping. AttnRes (Kimi Team et al., 2026) replaces the fixed accumulation with softmax attention over preceding layer outputs, allowing each layer to aggregate them with learned, input-dependent weights. However, these designs target non-looped networks, where the sources differ in content and are fixed in number. LoopRes is instead designed for loops, where the sources are iterates of the same layers and grow with the loop count.

Layer loops. Layer loops repeat a layer or a block of consecutive layers before advancing to the next. Sequence sharing (Takase & Kiyono, 2023) ties adjacent layers, so that each unique layer runs several times in a row. BRT (Marques et al., 2026) extends this to blocks, partitioning the layers into contiguous blocks and recursively reusing each block before moving to the next. However, both keep the standard residual stream and learn no connections across loops.

Stack loops. Stack loops repeat all layers or a contiguous subset end to end. Universal Transformers (Dehghani et al., 2019) apply a shared transition recurrently with a dynamic halting mechanism, and looped Transformers have since been shown to emulate programmable computers (Giannou et al., 2023) and to nearly match deeper non-looped models on synthetic reasoning tasks (Saunshi et al., 2025).

Direct looping feeds each loop the output of the previous loop. WideNet (Xue et al., 2022) shares attention and mixture-of-experts layers across blocks with individual layer normalization, Ouro (Zhu et al., 2025) loops the full stack with a learned early exit for adaptive computation, and TaH (Fu et al., 2026) feeds back logit-weighted embeddings for the tokens selected by a neural decider, with duo-causal attention over the key-value states of earlier loops.

Input injection additionally feeds the input into each loop. Huginn (Geiping et al., 2025) joins the latent state and the input embedding through a linear adapter, Parcae (Prairie et al., 2026) constrains the spectral norm of the injection parameters for stability, CloseLoop (Marchenko et al., 2026) mixes the token embedding into each later loop with one learned scalar, and TaH2 (You et al., 2026) extends TaH with an injection updater and an iteration decider that focuses extra iterations on the tokens that benefit.

Cross-loop aggregation combines the states of multiple loops. Hyperloop (Zeitoun et al., 2026) maintains parallel residual streams and mixes them with hyper-connections after each loop, WMRE (Zhuang et al., 2026) conditions each loop on its differences from recent states, and InfiLoop (Li et al., 2026) keeps a running summary of earlier outputs with content-based weights and a learned temporal decay.

Nevertheless, the input of each loop in these designs remains one or a few states compressed from earlier loops. Different from them, LoopRes keeps each block output as a separate source and lets every sub-layer build its input by attending directly to the sources of the latest W loops, with a loop-aware key that tells the iterates apart.

5. Conclusion

In this work, we presented LoopRes, an effective, scalable, and simple approach that lets each looped layer attend directly to the outputs of earlier loops. LoopRes assigns each source a trainable loop-aware key that encodes the structural relationship of loops and layers, and restricts each layer to a residual window of the latest W loops so that every loop builds incrementally on recent ones. Across 6 pretrained LLMs from 3 model families and 9 math and code benchmarks, LoopRes delivered significant and consistent gains, more than doubling the improvement of the strongest baseline at $2\times$ compute. Its score rose with a steeper slope than SFT as compute grew, reaching a higher ceiling and needing less compute to match SFT on long outputs. These advantages benefit both inference and post-training rollouts. We hope LoopRes serves as a simple and effective foundation for scaling LLMs along the loop-depth axis.

Limitations. Our experiments cover math and code benchmarks, and the gains on broader tasks such as general instruction following and open-ended generation remain to be validated. Additionally, LoopRes has been evaluated with supervised finetuning only, and its interaction with reinforcement learning has not yet been explored.

Future work. One key direction is to combine LoopRes with reinforcement learning, where the lower cost of matching SFT on long outputs could reduce the rollout budget. Another direction is to explore adaptive loop depth, allowing the model to decide how many loops each token or sequence needs. Furthermore, our findings suggest that the loop-depth scaling curve has not yet saturated on larger backbones, motivating further study of looping at greater scale. Finally, our analysis reveals that the learned loop dependencies follow clear, structured patterns, suggesting that a simpler, fixed connection scheme may capture most of the benefit without the overhead of learned attention weights.

Acknowledgements

We thank Fuzhao Xue, Tong Li, and Yong Luo for helpful discussions and feedback on this work. Yang You's research group is being sponsored by ByteDance grant, Alibaba grant, Meta Research Award, AliCloud grant, Apple grant, Google Research Award, Microsoft Research grant, and Adobe gift. We are grateful to NUSIT's Research Computing group for the computing support with numbers CFP03-CF-088 and CFP05-CF-062.

Impact Statement

LoopRes improves the math and code performance of pretrained LLMs by finetuning them into loops, rather than pretraining larger models from scratch. This may lower the cost of obtaining capable models and broaden access for researchers and practitioners. However, stronger reasoning and coding capabilities may also be misused, for example to generate malicious code. We therefore encourage safety evaluations and safeguards before deployment. Finally, looping increases the computation per token, which may raise energy consumption at scale.

References

Dao, T. FlashAttention-2: Faster attention with better parallelism and work partitioning. In *International Conference on Learning Representations*, 2024.

Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, Ł. Universal transformers. In *International Conference on Learning Representations*, 2019.

Fu, T., You, Y., Chen, Z., Dai, G., Yang, H., and Wang, Y. Think-at-Hard: Dynamic looped transformers for improved reasoning. In *International Conference on Machine Learning*, pp. 32129–32154, 2026.

Geiping, J., McLeish, S., Jain, N., Kirchenbauer, J., Singh, S., Bartoldson, B. R., Kailkhura, B., Bhatele, A., and Goldstein, T. Scaling up test-time compute with latent reasoning: A recurrent depth approach. *arXiv preprint arXiv:2502.05171*, 2025.

Giannou, A., Rajput, S., Sohn, J.-y., Lee, K., Lee, J. D., and Papailiopoulos, D. Looped transformers as programmable computers. In *International Conference on Machine Learning*, 2023.

Grattafiori, A., Dubey, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

He, C., Luo, R., Bai, Y., Hu, S., Thai, Z. L., Shen, J., Hu, J., Han, X., Huang, Y., Zhang, Y., Liu, J., Qi, L., Liu, Z., and Sun, M. OlympiadBench: A challenging benchmark for promoting AGI with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, 2024.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

Jain, N., Han, K., Gu, A., Li, W.-D., Yan, F., Zhang, T., Wang, S., Solar-Lezama, A., Sen, K., and Stoica, I. LiveCodeBench: Holistic and contamination free evaluation of large language models for code. In *International Conference on Learning Representations*, 2025.

Kimi Team, Chen, G., Zhang, Y., Su, J., Xu, W., Pan, S., Wang, Y., Wang, Y., Chen, G., Yin, B., Chen, Y., Yan, J., Wei, M., Zhang, Y., Meng, F., Hong, C., Xie, X., Liu, S., Lu, E., Tai, Y., Chen, Y., Men, X., Guo, H., Charles, Y., Lu, H., Sui, L., Zhu, J., Zhou, Z., He, W., Huang, W., Xu, X., Wang, Y., Lai, G., Du, Y., Wu, Y., Yang, Z., and Zhou, X. Attention residuals, 2026. URL https://arxiv.org/abs/2603.15031.

Li, P., Muhtar, D., He, D., Su, G., Yin, L., and Liu, S. Scaling to tens of thousands of test-time iterations with loop-native attention residuals. *arXiv preprint arXiv:2610.11570*, 2026.

Lightman, H., Kosaraju, V., Burda, Y., Edwards, H., Baker, B., Lee, T., Leike, J., Schulman, J., Sutskever, I., and Cobbe, K. Let's verify step by step. In *International Conference on Learning Representations*, 2024.

Liu, J., Xia, C. S., Wang, Y., and Zhang, L. Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems*, 2023.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Luong, T., Hwang, D., Nguyen, H. H., Ghiasi, G., Chervonyi, Y., Seo, I., Kim, J., Bingham, G., Lee, J., Mishra, S., Zhai, A., Hu, H., Michalewski, H., Kim, J., Ahn, J., Bae, J., Song, X., Trinh, T. H., Le, Q. V., and Jung, J. Towards robust mathematical reasoning. In *Empirical Methods in Natural Language Processing*, 2025.

Marchenko, A., Bezrukov, V., Kashurin, O., Fedorova, I., Bocharov, D., Shakhvalieva, Y., Tikhonova, M., and Ternovskii, V. Closing the loop: Practical training recipes for looped language models. *arXiv preprint arXiv:2610.00673*, 2026.

Marques, J., Cho, H. A., Kim, Y., Zhang, Y., Jung, A., Shim, K., and Hong, S. Block recursive transformers for structured parameter sharing. In *Findings of the Association for Computational Linguistics: EMNLP 2026*, 2026.

OpenBMB. UltraData-SFT-2605, 2026. URL https://huggingface.co/datasets/openbmb/UltraData-SFT-2605.

Prairie, H., Novack, Z., Berg-Kirkpatrick, T., and Fu, D. Y. Parcae: Scaling laws for stable looped language models. *arXiv preprint arXiv:2604.12946*, 2026.

Qwen Team. Qwen3.5: Towards native multimodal agents. https://qwen.ai/blog?id=qwen3.5, 2026.

Saunshi, N., Dikkala, N., Li, Z., Kumar, S., and Reddi, S. J. Reasoning with latent thoughts: On the power of looped transformers. In *International Conference on Learning Representations*, 2025.

Takase, S. and Kiyono, S. Lessons on parameter sharing across layers in transformers. In *Proceedings of The Fourth Workshop on Simple and Efficient Natural Language Processing (SustaiNLP)*, pp. 78–90, 2023.

Tian, X., Ji, Y., Wang, H., Chen, S., Zhao, S., Peng, Y., Zhao, H., and Li, X. Not all correct answers are equal: Why your distillation source matters. *arXiv preprint arXiv:2505.14464*, 2025.

Tillet, P., Kung, H.-T., and Cox, D. Triton: An intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*, pp. 10–19, 2019.

Wang, S., Zhang, G., Luo, K., Wu, Y., Liu, S., Liu, J., Huang, W., Yan, S., and Li, J. SMELT: Scaling laws for compute-matched MoE looped transformers, 2026. URL https://arxiv.org/abs/2609.01343.

Xie, Z., Wei, Y., Cao, H., Zhao, C., Deng, C., Li, J., Dai, D., Gao, H., Chang, J., Yu, K., Zhao, L., Zhou, S., Xu, Z., Zhang, Z., Zeng, W., Hu, S., Wang, Y., Yuan, J., Wang, L., and Liang, W. mhc: Manifold-constrained hyper-connections. *arXiv preprint arXiv:2512.24880*, 2025.

Xue, F., Shi, Z., Wei, F., Lou, Y., Liu, Y., and You, Y. Go wider instead of deeper. In *AAAI Conference on Artificial Intelligence*, pp. 8779–8787, 2022.

Yang, A. et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

You, Y., Fu, T., Feng, A., Lv, X., Ning, X., Ding, N., and Wang, Y. Improving test-time scaling with adaptive looped transformers. *arXiv preprint arXiv:2609.35748*, 2026.

Zeitoun, A., Torroba-Hennigen, L., and Kim, Y. Hyperloop transformers, 2026. URL https://arxiv.org/abs/2604.21254.

Zhu, R.-J., Wang, Z., Hua, K., Zhang, T., Li, Z., Que, H., Wei, B., Wen, Z., Yin, F., Xing, H., Li, L., Shi, J., Ma, K., Li, S., Kergan, T., Smith, A., Qu, X., Hui, M., Wu, B., Min, Q., Huang, H., Zhou, X., Ye, W., Liu, J., Yang, J., Shi, Y., Lin, C., Zhao, E., Cai, T., Zhang, G., Huang, W., Bengio, Y., and Eshraghian, J. Scaling latent reasoning via looped language models. *arXiv preprint arXiv:2510.25741*, 2025.

Zhuang, X., Wang, S., Razzak, I., and Liu, W. What makes recurrence effective in looped language models?, 2026. URL https://arxiv.org/abs/2609.36636.

Appendix

We organize our appendix as follows:

Experimental Details:

- Section A.1: Model configurations
- Section A.2: Baselines
- Section A.3: Data
- Section A.4: Training
- Section A.5: Evaluation
- Section A.6: Pseudo code

Supplementary Results:

- Section B.1: Per-benchmark scores
- Section B.2: Validation loss over finetuning
- Section B.3: Test-time scaling

Additional Experiments:

- Section C.1: Looping all layers with AttnRes
- Section C.2: Doubling the SFT data

A. Experimental Details

A.1. Model Configurations

Table 6 lists the configuration of LoopRes on each backbone. We split each backbone into a prelude, a core, and a coda of one quarter, one half, and one quarter of its layers, and cut the core into $B = 2$ blocks. We measure compute by the number of layers each token passes through relative to the backbone, so running the core $L = 3$ and $L = 7$ times costs 2 and 4 full-model passes, respectively.

Table 6. **Configuration of LoopRes on each backbone.** L and W are given at $2\times$ / $4\times$ compute, and – marks a budget not run.

Backbone	Layers	Prelude / Core / Coda	B	L	W
Qwen3-1.7B-Base	28	7 / 14 / 7	2	3 / 7	1 / 1
Qwen3-4B-Base	36	9 / 18 / 9	2	3 / –	1 / –
Qwen3-8B-Base	36	9 / 18 / 9	2	3 / –	1 / –
Qwen3-30B-A3B-Base	48	12 / 24 / 12	2	3 / –	1 / –
Qwen3.5-2B-Base	24	6 / 12 / 6	2	3 / –	1 / –
Llama-3.2-1B	16	4 / 8 / 4	2	3 / –	1 / –

A.2. Baselines

We reimplement every baseline on the same backbone and train it with its own objective, the cross-entropy of the last loop unless stated otherwise. Since two of them are unnamed, we call them WMRE (Zhuang et al., 2026) and CloseLoop (Marchenko et al., 2026) after their titles. The originals of Ouro, InfiLoop, BRT, Huginn, Hyperloop, and WMRE pretrain from scratch; we finetune the pretrained backbone throughout. Every token runs all L loops, in training and at inference, and the three core-looped designs run the coda once, on the last loop's output.

Ouro (Zhu et al., 2025). Our implementation loops all layers, each loop reading the previous loop's normed output, and averages the cross-entropy over all loops. The original weights the loops by a learned exit distribution under an entropy-regularized objective; we remove the exit gate and weight the loops uniformly.

TaH2-fixed (You et al., 2026). The paper's fixed-depth variant, following its released recipe. It keeps TaH2's updater, a SwiGLU block as wide as the model that merges the token embeddings with the previous loop's output, and removes the iteration decider. Attention runs in the released code's causal mode: earlier tokens expose their loop-1 keys, and each token reads its own keys of the later loops. The model predicts from the mean of the loops' next-token distributions, and the objective is the cross-entropy of this mean.

CloseLoop (Marchenko et al., 2026). Loop $t > 1$ reads $w\boldsymbol{e} + (1 - w)\boldsymbol{h}_{t-1}$, with w the sigmoid of one trained scalar at the paper's $w = 0.99$, and the LM head is untied, starting as a copy of the input embeddings. An exit gate scores each loop, and the objective weights the loops' cross-entropy by the exit distribution, with the paper's smoothing 0.05 and entropy coefficient 0.15. We omit the paper's optional loop-specific LoRA adapters and always predict from loop L.

InfiLoop (Li et al., 2026). Each loop reads a running average of the earlier loops' outputs, weighted by a trained decay and a content score, $\exp(\boldsymbol{w}^\top \mathrm{RMSNorm}(\boldsymbol{z}_t))$ clamped to $[-8, 8]$, with the released code's initialization ($\boldsymbol{w}$ at zero, the decay at 0.55). The original runs a two-level recurrence with one average per level, restarts the inner one each cycle, seeds the outer one with a learned state, and backpropagates only through the last cycle; we keep a single average over the L loops, seed it with the token embeddings, and backpropagate through all of them. Its post-norm layers also keep the averaged states at unit scale, while we average the raw residual stream of the pre-norm backbone.

BRT (Marques et al., 2026). Consecutive layers are grouped into blocks, and each block runs L times on the standard residual stream before the next. We use seven blocks of four layers on the 28-layer backbone, where the original defaults to two blocks.

Huginn (Geiping et al., 2025). Retrofitted with the same prelude, core, and coda as LoopRes: each loop the core reads Parcae's diagonal injection (Prairie et al., 2026) of the previous core output and the prelude output, initialized at its retention $\sqrt{1/5}$ with the injection the identity, so that one loop reproduces the pretrained model. The original samples the loop count from a log-normal Poisson distribution, draws the initial state from a truncated normal, and joins the two inputs by a linear map of their concatenation; we fix the loop count, start the state at zero, and use the diagonal adapter.

Hyperloop (Zeitoun et al., 2026). The core runs on four parallel residual streams, joined after each loop by its own sigmoid-gated hyper-connection, with a per-loop vector trained from zero added to the core output, and the coda reads the stream mean. The gate scales start at mHC's 0.01 (Xie et al., 2025), and the biases start the pre, post, and res gates at $1/4$, 1, and $1/2$. The stream count and the begin–middle–end split match the original.

WMRE (Zhuang et al., 2026). The paper's H+LG configuration, on the same split as LoopRes: loop $t + 1$ reads $\boldsymbol{h}_t + \sum_j \boldsymbol{b}_j \odot (\boldsymbol{h}_{t-j} - \boldsymbol{h}_t)$ over the lags $j \le 2$ that reach back to the first loop, with $\boldsymbol{h}_t$ the state after loop t, and a scalar gate over the paper's eight time features scales each loop's update. Both sets of weights train from zero.

A.3. Data

We take the math and code questions of AM-Qwen3-Distilled (Tian et al., 2025) and UltraData-SFT-2605 (OpenBMB, 2026) and answer each with Qwen3-8B in thinking mode, sampling with temperature 0.6, top-p 0.95, and top-k 20 up to 16,384 tokens. A judge model Qwen3-235B-A22B checks whether each reply reaches the same final answer as the original reply, and we keep only the consistent replies. We further drop malformed replies and every question that shares a run of 13 words with evaluation benchmarks. Every other backbone trains on the same questions with their original Qwen3-235B-A22B replies, Llama-3.2-1B in the Llama 3 chat format. The validation set holds 1,000 samples of the source datasets with their original replies, disjoint from the training questions and shared by all backbones.

A.4. Training

Table 7 lists the hyperparameters. Each step packs several conversations per sequence, FlashAttention-2 (Dao, 2024) restricts attention to within each conversation, and only the reply tokens contribute to the loss. We sweep the learning rate over $\{3 \times 10^{-6}, 1 \times 10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 1 \times 10^{-4}\}$ and keep 3×10^{-5} for every model.

Table 7. **Training hyperparameters**, shared by all models.

Hyperparameter	Value
Optimizer	AdamW (Loshchilov & Hutter, 2019), $\beta_1 = 0.9$, $\beta_2 = 0.95$
Weight decay	0.01 on matrices, 0 on vectors
Learning rate	3×10^{-5}, linear warmup over 3% of steps, cosine decay to 3×10^{-6}
Gradient clipping	1.0
Training steps	12,000
Batch size	16 sequences
Sequence length	16,384 tokens
Sequence packing	Random order, attention within each conversation
Precision	bfloat16 mixed precision

A.5. Evaluation

We run the benchmarks on the final checkpoint at step 12,000. Replies are generated in thinking mode under the chat template, with temperature 0.6, top-p 0.95, and top-k 20 up to 32,768 new tokens. HumanEval+ and MBPP+ run both the base and the extended tests of EvalPlus (Liu et al., 2023). We report pass@1 averaged over the samples of each problem, and Table 8 lists the number of problems, the samples per problem, and the backbones evaluated on each benchmark.

For the test-time scaling curves, we truncate the thinking at each budget and force the model to start its reply, so their scores differ from those in the tables, which use the full replies.

Table 8. **Problems, samples per problem, and evaluated backbones on each benchmark.** We average pass@1 over the samples of each problem. HE+ denotes HumanEval+, IMO-AB IMO-AnswerBench, and LCB LiveCodeBench v6.

benchmark	math						code		
	AIME25	AIME26	AMC23	Olympiad	MATH500	IMO-AB	HE+	MBPP+	LCB
problems	30	30	40	675	500	400	164	378	175
samples per problem	16	16	16	1	1	2	4	2	2
backbones	all	all	<4B	all	<4B	≥4B	<4B	<4B	≥4B

A.6. Pseudo Code

Figure 9 provides PyTorch-style pseudo code for LoopRes.

B. Supplementary Results

B.1. Per-Benchmark Scores

Ablation study. Table 9 reports the per-benchmark scores of the variants in Table 5.

Loop depths. Table 10 reports the per-benchmark scores of the loop depths in Figure 6.

Llama-3.2-1B and Qwen3.5-2B-Base. Table 11 reports the per-benchmark scores of Table 4.

B.2. Validation Loss over Finetuning

Figure 10 reports the validation loss of each design relative to SFT over finetuning.

```python
def loop_res(values: list[Tensor], loop_keys: list[Tensor], loop_query: Tensor, content_query: Tensor,
             norm: RMSNorm) -> Tensor:
    """
    AttnRes input with loop-aware keys.
    values:
        S sources of shape [N, T, D]
    loop_keys:
        S loop-aware keys of shape [D], one trainable vector per source
    """
    values = torch.stack(values)        # [S, N, T, D]
    loop_keys = torch.stack(loop_keys)  # [S, D]
    content_keys = norm(values)         # [S, N, T, D], as in AttnRes
    logits = torch.einsum('d, s n t d -> s n t', content_query, content_keys)
    logits = logits + torch.einsum('d, s d -> s', loop_query, loop_keys)[:, None, None]  # loop-aware keys
    weights = logits.softmax(0)
    return torch.einsum('s n t, s n t d -> n t d', weights, values)

def forward(self, prelude_out: Tensor) -> Tensor:
    """
    prelude_out:
        [N, T, D]: output of the prelude; the returned core output goes to the coda
    """
    # residual window: the blocks of the previous W loops and their loop-aware keys
    window_values, window_keys = [], []
    for t in range(self.num_loops):
        values = [prelude_out] + window_values
        loop_keys = [self.prelude_key] + window_keys

        # block AttnRes over the core; every sub-layer forms its input by loop_res, with its own queries in loop t
        blocks = []
        for i, block in enumerate(self.core_blocks):
            blocks.append(block(values + blocks, loop_keys + self.block_keys[t][:i], loop=t))

        # end of loop t: once the window is full, drop all blocks of its oldest loop, then add loop t
        if len(window_values) == self.window_size * self.num_blocks:
            del window_values[:self.num_blocks], window_keys[:self.num_blocks]
        window_values += blocks
        window_keys += self.block_keys[t]

    # output: the window and the blocks of the last loop
    return loop_res(values + blocks, loop_keys + self.block_keys[t],
                    self.output_loop_query, self.output_content_query, self.output_norm)
```

Figure 9. **PyTorch-style pseudo code for LoopRes.** Lines in blue mark its two changes to Block AttnRes: (i) `loop_res` adds a loop-aware key to each source. (ii) Each loop takes a residual window of the previous W loops as input, from which older loops are dropped.

B.3. Test-Time Scaling

Figures 11 to 13 report the test-time scaling of LoopRes and SFT on Qwen3-30B-A3B-Base, Llama-3.2-1B, and Qwen3.5-2B-Base as the thinking budget grows from 1K to 30K tokens.

C. Additional Experiments

C.1. Looping All Layers with AttnRes

We run Huginn + AttnRes with all layers looped in place of the prelude, core, and coda, at $L = 2$ so that it matches the $2\times$ compute of the split version at $L = 3$. Table 12 reports the results.

C.2. Doubling the SFT Data

The runs of Table 1 finetune for 12,000 steps on a store of 11,834 steps, so we also train SFT for 24,000 steps on a store of 23,513 steps and read its final checkpoint at step 24,000. Table 13 reports the per-benchmark scores, and Figure 14 the test-time scaling of both SFT models.

Table 9. **Per-benchmark scores of the ablation study on Qwen3-1.7B-Base.** The variants are those of Table 5, and LoopRes uses $W = 1$ at both compute budgets. HE+ denotes HumanEval+, and both code benchmarks report the plus tests. *avg.* averages the seven math and code scores.

variant	math					code		avg.
	AIME25	AIME26	AMC23	Olympiad	MATH500	HE+	MBPP+	
2× compute								
LoopRes	27.9	19.4	64.7	53.8	84.0	71.8	68.9	55.8
Huginn + AttnRes	25.4	19.0	62.3	52.0	82.6	71.2	65.1	53.9
window: $W = 2$	28.1	18.3	62.2	53.5	83.0	72.6	68.9	55.2
key: loop-aware only	25.8	21.5	62.3	49.9	84.8	71.0	68.8	54.9
key: content only	25.8	20.0	62.2	51.9	83.6	72.7	68.0	54.9
4× compute								
LoopRes	30.2	24.6	67.5	52.9	84.2	75.3	69.8	57.8
Huginn + AttnRes	27.9	21.0	63.1	52.4	83.6	71.2	66.0	55.0
window: $W = 2$	31.0	22.3	65.0	54.8	85.6	74.7	68.8	57.5
window: $W = 3$	29.0	24.2	68.0	52.4	87.2	74.2	68.8	57.7
key: loop-aware only	31.0	23.1	65.9	55.1	84.2	73.2	69.7	57.5
key: content only	27.7	22.9	63.0	54.8	83.6	72.6	66.1	55.8

Table 10. **Per-benchmark scores across loop depths on Qwen3-1.7B-Base.** Each row runs the core L times, at the per-token compute in parentheses relative to SFT. HE+ denotes HumanEval+, and both code benchmarks report the plus tests. *avg.* averages the seven math and code scores. **Bold** marks the best score of each column.

depth	math					code		avg.
	AIME25	AIME26	AMC23	Olympiad	MATH500	HE+	MBPP+	
SFT	14.6	11.7	48.6	40.7	73.4	57.3	56.2	43.2
$L = 2$ (1.5×)	22.9	15.6	60.6	49.6	83.0	70.1	63.0	52.1
$L = 3$ (2×)	27.9	19.4	64.7	53.8	84.0	71.8	68.9	55.8
$L = 5$ (3×)	26.9	22.7	64.8	**55.7**	85.4	73.0	68.8	56.8
$L = 7$ (4×)	**30.2**	**24.6**	**67.5**	52.9	84.2	**75.3**	**69.8**	**57.8**
$L = 9$ (5×)	28.5	22.5	67.0	54.7	84.8	73.3	66.0	56.7
$L = 11$ (6×)	29.2	23.5	65.8	55.3	**85.8**	71.3	66.8	56.8

Table 11. **Per-benchmark scores on Llama-3.2-1B and Qwen3.5-2B-Base at 2× compute.** Both code benchmarks report the plus tests. Subscripts give the change of LoopRes relative to SFT in points. *avg.* averages the seven math and code scores.

benchmark	Llama-3.2-1B		Qwen3.5-2B-Base	
	SFT	**LoopRes**	SFT	**LoopRes**
AIME25	0.4	**1.0**$_{+0.6}$	24.6	**37.7**$_{+13.1}$
AIME26	0.2	**0.4**$_{+0.2}$	24.0	**41.2**$_{+17.2}$
AMC23	13.3	**22.7**$_{+9.4}$	64.7	**80.9**$_{+16.2}$
OlympiadBench	9.5	**16.6**$_{+7.1}$	52.0	**63.6**$_{+11.6}$
MATH500	28.4	**43.8**$_{+15.4}$	85.6	**89.6**$_{+4.0}$
HumanEval+	29.1	**32.5**$_{+3.4}$	62.0	**71.0**$_{+9.0}$
MBPP+	18.9	**20.5**$_{+1.6}$	51.6	**64.8**$_{+13.2}$
avg.	14.3	**19.6**$_{+5.4}$	52.1	**64.1**$_{+12.0}$

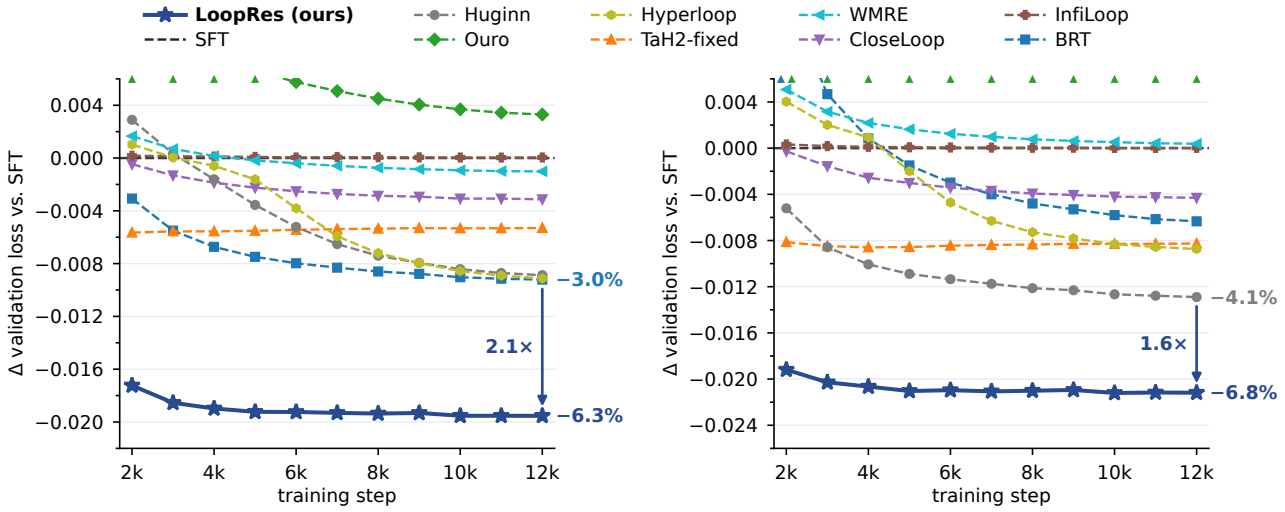

(a) Loop of 2× compute. *(b)* Loop of 4× compute.

Figure 10. **Validation loss relative to SFT.** We finetune Qwen3-1.7B-Base on code and math with each looped design and subtract the loss of the same model finetuned without looping (SFT), each point averaging 1,000 steps. Panels (a) and (b) give every design the compute of two and four full passes through the model, respectively. Triangles on the top edge mark values above the plotted range.

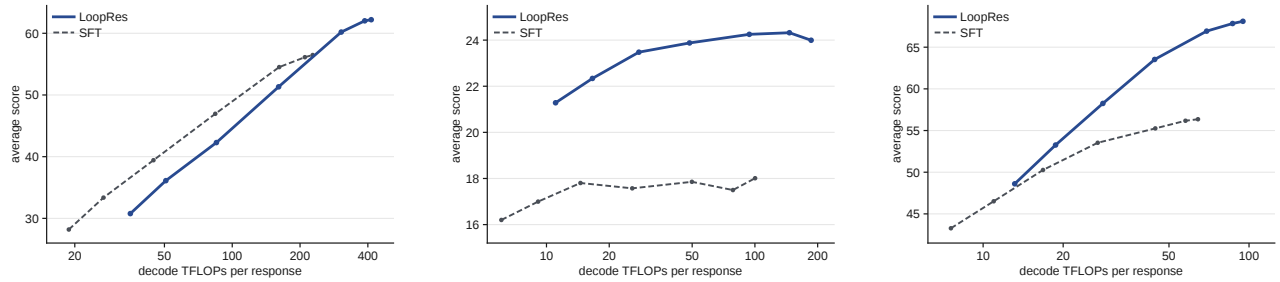

Figure 11. **Test-time scaling on Qwen3-30B-A3B-Base.** Average score on the five benchmarks of Table 3 against mean decode compute, as the thinking budget grows from 1K to 30K tokens. LoopRes runs at 2× compute, so it spends twice SFT's compute per token.

Figure 12. **Test-time scaling on Llama-3.2-1B.** Average score on the seven benchmarks of Table 11 against mean decode compute, as the thinking budget grows from 1K to 30K tokens. LoopRes runs at 2× compute, so it spends twice SFT's compute per token.

Figure 13. **Test-time scaling on Qwen3.5-2B-Base.** Average score on the seven benchmarks of Table 11 against mean decode compute, as the thinking budget grows from 1K to 30K tokens. LoopRes runs at 2× compute, so it spends twice SFT's compute per token.

Table 12. **Huginn variants at 2× compute on Qwen3-1.7B-Base.** *core* loops the core of a prelude, core, and coda, as in Tables 1 and 5, and *all* loops the whole model. *avg.* averages the seven math and code scores.

variant	looped layers	L	avg.
Huginn	core	3	49.0
Huginn + AttnRes	all	2	50.9
Huginn + AttnRes	core	3	**53.9**

Table 13. **Doubling the SFT data against looping on Qwen3-1.7B-Base.** Every row is scored under the decoding of Section A.5 with no thinking budget. SFT at 1× data is read from the final checkpoint at step 12,000 and at 2× data from the final checkpoint at step 24,000, both at 1× compute, and the LoopRes rows are those of Table 1, trained on 1× data. HE+ denotes HumanEval+, and both code benchmarks report the plus tests. *avg.* averages the seven math and code scores, and Δ gives the change relative to SFT at 1× data.

setting	math					code		avg.	Δ
	AIME25	AIME26	AMC23	Olympiad	MATH500	HE+	MBPP+		
SFT, 1× data	14.6	11.7	48.6	40.7	73.4	57.3	56.2	43.2	–
SFT, 2× data	20.8	15.2	56.9	45.9	80.2	66.6	60.7	49.5	+6.3
LoopRes, 2× compute	27.9	19.4	64.7	53.8	84.0	71.8	68.9	55.8	+12.6
LoopRes, 4× compute	30.2	24.6	67.5	52.9	84.2	75.3	69.8	57.8	+14.6

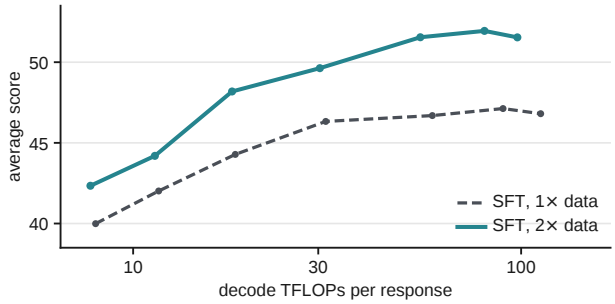

Figure 14. **Test-time scaling after doubling the SFT data.** Average score on seven math and code benchmarks against decode compute as the thinking budget grows from 1K to 30K tokens, on Qwen3-1.7B-Base. Both curves are SFT, trained on 1× and 2× the finetuning data.